

Introduction

Introduction to Databases

CompSci 316 Spring 2020



DUKE
COMPUTER SCIENCE

Welcome to

CompSci 316: Introduction to Database Systems!!
Spring 2020

About us: instructor

- Instructor: [Sudeepa Roy](#)
 - At Duke CS since Fall 2015
 - Member of “Duke Database Devils”
a.k.a. the database research group
 - PhD. UPenn, Postdoc: U. of Washington
 - Research interests:
 - “data”
 - data management, database theory, data analysis, data science, causality and explanations, uncertain data, data provenance, crowdsourcing,




Meet your grad TAs



Tiangang Chen

- Duke CS MS student
- Interested in computer systems
- Runs a half-marathon every year!

- Duke CS MS student
- Interested in data science and machine learning, currently working on image processing
- Loves cats!



Xiangchen Shen




Yuchao Tao

- Duke CS PhD student
- Interested in databases and privacy research
- Enjoys cooking!

* All CompSci 516 veterans

Meet your UTAs



David Chen

- Duke CS major
- Interested in applying computational models and algorithms to biological systems
- Loves to downhill ski!

- Duke CS major, minor in Visual Media studies
- Interested in UI/UX, front-end development, and project management
- Has two dogs and a hamster!



Jane Li



Runxin (Rebecca) Wang

- Duke CS major
- Interested in coding, machine learning, and theory
- Loves hiking and bubble tea!

* All CompSci 316 veterans

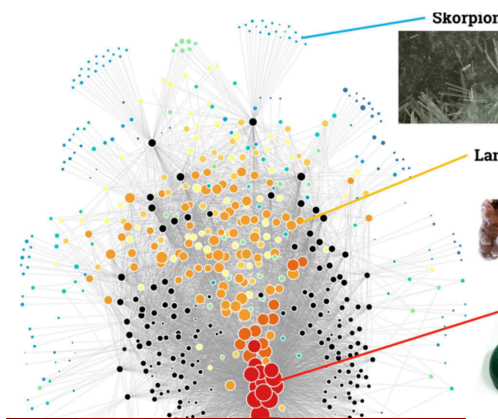
What are the goals of this course?

- Learn about “databases” or data management

Why do we care about data? (easy)

How big data can help find new mineral deposits

Valentina Ruiz Leotaud | Aug. 2, 2018, 4:11 PM |



Skorpion

Lar

Cal



Cambridge Analytica whistleblower Chris Wylie speaks during a press conference at the Frontline Club on March 26, 2018 in London | Dan Kitwood/Getty Images

Cambridge Analytica helped 'cheat' Brexit vote and US election, claims whistleblower

Giving evidence to MPs, Chris Wylie claimed the company's actions during the Brexit campaign were "a breach of the law."

By MARK SCOTT | 3/27/18, 5:46 PM CET | Updated 3/29/18, 9:18 PM CET

The New York Times


When Sports Betting Is Legal, the Value of Game Data Soars



A trader working at William Hill, an international sports betting book, in Las Vegas. Bridget Bennett for The New York Times



Transportation Researchers are trying to teach computers to forecast traffic like the weather



... The three years of gathering and analyzing data culminated in what U.S. Sailing calls their "Rio Weather Playbook," a body of critical information about each of the seven courses only available to the U.S. team...

— FiveThirtyEight, "Will Data Help U.S. Sailing Get Back On The Olympic Podium?"
Aug 15, 2016

Data =
Money
Information
Power
Fun
in
Science, Business,
Politics, Security
Sports, Education, ...



Wait.. don't we need to take a Machine Learning or Stat course for those things?

Yes, but..



... we also need to manage this (huge or not-so-huge) data!

Also think about building a new App or website based on data from scratch

- E.g., **your own version of book purchase platform** (like a mini-Amazon)
- Large data! (think about all books in the world or even in English)

• **How do we start?**

* You are going to do something similar in the course project!

Who are the key people?

Who are the key people?

- At least two types:
 - Database admin (assuming they own all copies of all the books)
 - Users who purchase books
 - Let's proceed with these two only

- Other people:
 - Sellers
 - HR
 - Finance
 - Who deal with the warehouse of the books
 -

What should the user be able to do?

- i.e. what the interface look like? (think about Amazon)

What should the user be able to do?

- i.e. what the interface look like? (think about Amazon)
1. Search for books
 - With author, title, topic, price range,
 2. Purchase books
 3. Bookmark/add to wishlist

What should the platform do?

What should the platform do?

1. Returns books as searched by the authors
2. Check that the payment method is valid
3. Update no. of copies as books are sold
4. Manage total money it has
5. Add new books as they are published
6.

What are the desired and necessary properties of the platform?

What are the desired and necessary properties of the platform?

- Should be able to handle a **large amount of data**
- Should be **efficient** and **easy to use** (e.g., search with **authors as well as title**)
- If there is a crash or loss of power, **information should not be lost or inconsistent**
 - Imagine a user was in the middle of a transaction when a crash happened, paid the money, but the book has not been purchased
- No surprises with **multiple users** logged in at the same time
 - Imagine one last copy of a book that two users are trying to purchase at the same time
- Easy to **update and program**
 - For the admin

That was the design phase
(a basic one though)



How about C++, Java, or Python?
On data stored in large files

Sounds simple!

James Morgan#Durham, NC

... ..

A tale of two cities#Charles Dickens#3.50#7

To Kill a Mockingbird#Harper Lee#7.20#1

Les Miserables#Victor Hugo#12.80#2

... ..

- Text files – for books, customer, ...
- Books listed with title, author, price, and no. of copies
- Fields separated by #'s

Query by programming

James Morgan#Durham, NC

... ..

A tale of two cities#Charles Dickens#3.50#7

To Kill a Mockingbird#Harper Lee#7.20#1

Les Miserables#Victor Hugo#12.80#2

... ..

- James Morgan wants to buy “To Kill a Mockingbird”

- A simple script

Better idea than scanning?

- Scan through the books file
- Look for the line containing “To Kill a Mockingbird”
- Check if the no. of copies is ≥ 1
- Bill James \$7.20 and reduce the no. of copies by 1

Binary search! Keep
file sorted on titles

What if he changes the “query” and wants to buy a book by Victor Hugo?

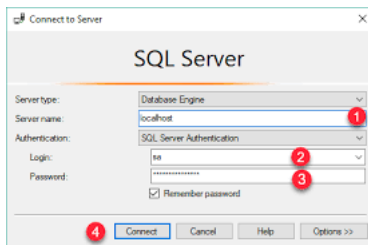
Revisit: What are the desired and necessary properties of the platform?

- Should be able to handle a **large amount of data**
 - Try to open a 10-100 GB file
- Should be **efficient** and **easy to use** (e.g., search with authors as well as title)
 - Try to search both on a large flat file
- If there is a crash or loss of power, **information should not be lost or inconsistent**
 - Imagine a user was in the middle of a transaction when a crash happened, paid the money, but the book has not been purchased
 - Imagine programmer's task
- No surprises with **multiple users** logged in at the same time
 - Imagine one last copy of a book that two users are trying to purchase at the same time
 - Imagine adding a new book or updating Copies (+ allow search) on a 10-100 GB text file
- Easy to **update and program**
 - For the admin

Solution?



- DBMS = Database Management System



A DBMS takes care of all of the following (and more):

In an easy-to-code, efficient, and robust way

- Should be able to handle a large amount of data
- Should be efficient and easy to use (e.g., search with authors as well as title)
- If there is a crash or power, information should not be lost or inconsistent
- If a user was in the middle of a transaction when a crash happened, paid the money, but the book has not been purchased
- No surprises with multiple users logged in at the same time
 - Imagine one last copy of a book that two users are trying to purchase at the same time
- Easy to update and program
 - For the admin

Optimization

Index

Recovery

Consistency

Declarative

* We will learn these in the course!

DBMS helps the big ones!

The screenshot shows the MySQL website's 'MySQL Customer: Facebook' page. The header includes the MySQL logo and navigation links like 'MYSQL.COM', 'DOWNLOADS', 'DOCUMENTATION', and 'DEVELOPER ZONE'. The main content area features the Facebook logo and a quote: "We are one of the largest MySQL web sites in production. MySQL has been a revolution for young entrepreneurs." A sidebar on the left lists navigation options like 'Customer Overview' and 'Case Studies'.

The screenshot shows a Google Cloud documentation page titled 'Querying Cloud Bigtable data'. The page is part of the 'Data Analytics Products' documentation. It includes a table of contents with items like 'Labeling BigQuery resources', 'Loading data into BigQuery', and 'Querying external data sources'. The main content area describes how to use BigQuery to query data stored in Cloud Bigtable, mentioning that Cloud Bigtable is a sparsely populated NoSQL database that can scale to billions of rows.

The screenshot shows Mark Zuckerberg's Facebook profile. It includes his profile picture, name, and a bio: "We just took down networks of coordinated information campaigns from Iran and Russia as part of our efforts to protect against election interference. Here's what I said on a press call we held to share more details about these actions." Below the bio, there are statistics for likes and shares, and a 'Continue Reading' link. The page also shows a 'Photos' section with several images.

The screenshot shows the MySQL website's 'MySQL Customers' page. The page features a grid of customer logos and quotes. The logos include GitHub, Facebook, and ITALTEL. Each logo is accompanied by a short quote about their use of MySQL. For example, GitHub says: "MySQL is our core data store that we used for storing all data that powers the site as well as the metadata around the users." The page also includes navigation links and a 'Learn More' button.

Note: Not always the “standard” DBMS (called Relational DBMS), but we need to know pros and cons of all alternatives

CompSci 316 gives an intro to DBMS

- How can a user use a DBMS (programmer's/designer's perspective)
 - Run queries, update data (SQL, Relational Algebra)
 - Design a good database (ER diagram, normalization)
 - Use different types of data (Mostly relational, also XML/JSON)
- How does a DBMS work (system's or admin's perspective)
 - Storage, index
 - Query processing, join algorithms, query optimizations
 - Transactions: recovery and concurrency control
- Glimpse of advance topics and other DBMS
 - NOSQL, Spark (big data)
 - Data mining
- Hands-on experience in class projects by building an end-to-end website or an app that runs on a database

Misc. course info

- Website:
<https://www2.cs.duke.edu/courses/spring20/compsci316/>
 - Course info; tentative schedule and reference sections in the book; lecture slides, assignments, help docs, ...
- Book: *Database Systems: The Complete Book*, by H. Garcia-Molina, J. D. Ullman, and J. Widom. 2nd Ed.
- **Programming**: VM required, need significant programming on different platforms and languages
- Prerequisite: **CompSci 201** - or you would have to learn some concepts yourself
- Q&A on **Piazza**
- Grades, sample solutions on **Sakai**
- Submissions on **Gradescope** and **Gradiance**
- Watch your email for announcements

Important: Grading

Absolute but adjustable grading

Guarantees:

[90%, 100%] A- / A / A+

[80%, 90%) B- / B / B+

[70%, 80%) C- / C / C+

[60%, 70%) D

Class topper gets A+

- Scale will not go upwards but can get downwards (e.g., based on the class performance in the exams)
- We will give you a feedback on your approximate standing after the midterm.

Duke Community Standard

- See course website for link
- Group discussion for assignments is okay (and encouraged), but
 - Acknowledge any help you receive from others
 - Make sure you “own” your solution
- All suspected cases of violation will be aggressively pursued

Course load

- (See course webpage for full details)
- **Weekly (short) homework assignments (25%)**
 - Each homework has same weight
 - Released on Tuesdays and due next Tuesday night (mostly)
 - **Gradiance**: immediately and automatically graded
 - **Gradescope**: programming problems, immediate feedback, later also manual grading
 - **Gradescope**: written solution, manual grading
- **Midterm and final (20% each)**
 - Open book, open notes
 - No communication/Internet whatsoever
 - Final is comprehensive, but emphasizes the second half of the course

Course load (contd.)

- **Course project (20%)**
 - Details to be given in the next 1-2 weeks
- **In-class quiz (5%)**
 - To review concepts right away in class – will be open for 5-10 mins
 - Will be announced at least one class in advance and on piazza
 - Each quiz: 50% for attempt on time and 50% for correct solution
 - Lowest score will be dropped (each quiz has same weight)
- **In-class labs (5%)**
 - Practice problems in class (both programming and conceptual) – each lab has the same weight
 - Will be announced at least one class in advance and on piazza
 - Due by the next day after class, 10% bonus points for finishing all problems in class correctly
 - TAs will be around to help you

Tentative office hours schedules

- Locations: TBD.
- See the updated info on the webpage
- More office hours around Tuesday (hws due), but good to start early!

	MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY	SUNDAY
8:00 AM							
8:30 AM							
9:00 AM							
9:30 AM							
10:00 AM				Tiangang (10 - 11)	Tiangang (10 - 11)		
10:30 AM							
11:00 AM							
11:30 AM							
12:00 PM	Xiangcheng (12-1)						
12:30 PM							
1:00 PM							
1:30 PM		Yuchao (1:30 - 2:30)					
2:00 PM	Sudeepa (2-3)				Sudeepa (2-3)		
2:30 PM							
3:00 PM							
3:30 PM			Xiangcheng (3:30 - 4:30)				
4:00 PM							
4:30 PM							
5:00 PM							
5:30 PM							
6:00 PM							
6:30 PM							
7:00 PM	Jane (7-9)	Rebecca (7-9)					David (7-9)
7:30 PM							
8:00 PM							
8:30 PM							

Projects from past years

- RA: next-generation relational algebra interpreter
 - You may get to try it out for Homework #1!
- *Managing tent shifts and schedules!*
- *Tutor-tutee matching*
- *What's in my fridge and what can I cook?*
- *Hearsay: manage your own musics*
- *Dining at Duke (and deliver meals to students)*
- *National Parklopedia: a website to find information about national parks*

- *More examples later - but we expect you to be creative with a new idea!*


Let's get started!

Relational Data Model

What is a good model to store data?
Tree? Nested data? Graph?

(just) **Tables!**

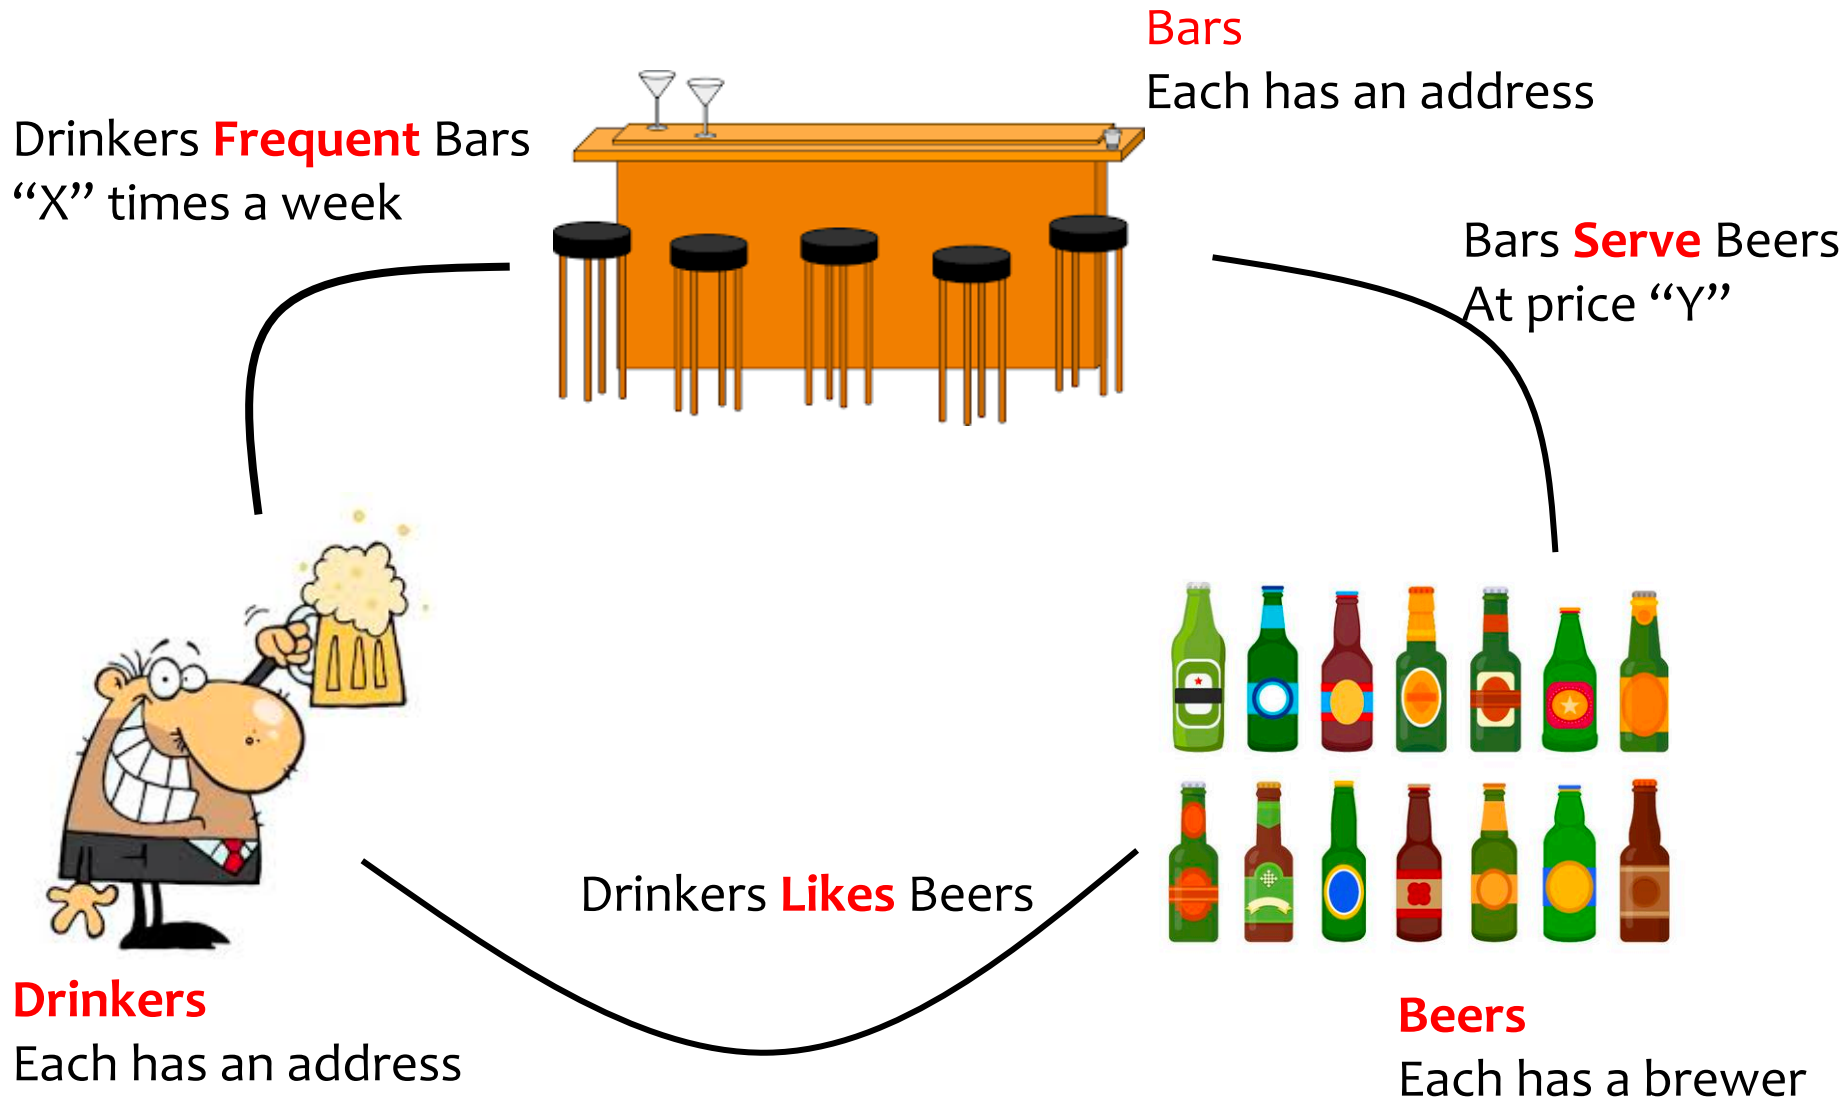
Edgar F. Codd (1923-2003)



- Pilot in the Royal Air Force in WW2
- Inventor of the relational model and algebra while at IBM
- Turing Award, 1981

RDBMS = Relational DBMS

The famous “Beers” database



(Later in ER diagram – how to design a relational database)

“Beers” as a Relational Database

Bar

name	address
The Edge	108 Morris Street
Satisfaction	905 W. Main Street

Beer

Name	brewer
Budweiser	Anheuser-Busch Inc.
Corona	Grupo Modelo
Dixie	Dixie Brewing

Drinker

name	address
Amy	100 W. Main Street
Ben	101 W. Main Street
Dan	300 N. Duke Street

Serves

bar	beer	price
The Edge	Budweiser	2.50
The Edge	Corona	3.00
Satisfaction	Budweiser	2.25

drinker	bar	times_a_week
Ben	Satisfaction	2
Dan	The Edge	1
Dan	Satisfaction	2

Frequents

drinker	beer
Amy	Corona
Dan	Budweiser
Dan	Corona
Ben	Budweiser

Likes

Relational data model

- A database is a collection of **relations** (or **tables**)
- Each relation has a set of **attributes** (or **columns**)
- Each attribute has a name and a **domain** (or **type**)
 - Set-valued attributes are not allowed
- Each relation contains a “**set**” of **tuples** (or **rows**)
 - Each tuple has a value for each attribute of the relation
 - Duplicate tuples are not allowed (Two tuples are duplicates if they agree on all attributes)
 - Ordering of rows doesn’t matter (even though output is always in some order)
- However, SQL supports “**bag**” or duplicate tuples (why?)

👉 **Simplicity is a virtue**

- **not a weakness!**

Serves

bar	beer	price
The Edge	Budweiser	2.50
The Edge	Corona	3.00
Satisfaction	Budweiser	2.25

Schema vs. instance

• Schema

- *Beer* (name string, brewer string)
- *Serves* (bar string, beer string, price float)
- *Frequents* (drinker string, bar string, times_a_week int)

• Instance

- Actual tuples or records

☞ Compare to **types** vs. collections of **objects of these types** in a programming language

bar	beer	price
The Edge	Budweiser	2.50
The Edge	Corona	3.00
Satisfaction	Budweiser	2.25

Serves

Name	brewer
Budweiser	Anheuser-Busch Inc.
Corona	Grupo Modelo
Dixie	Dixie Brewing

Beer

Frequents

drinker	bar	times_a_week
Ben	Satisfaction	2
Dan	The Edge	1
Dan	Satisfaction	2

SQL: Querying a RDBMS

- SQL: **Structured Query Language**
 - Pronounced “S-Q-L” or “sequel”
 - The standard query language supported by most DBMS
 - First developed at IBM System R
 - Follows ANSI standards

SQL is Declarative:

Programmer specifies **what** answers a query should return, but **not how** the query is executed

DBMS picks the best execution strategy based on availability of indexes, data/workload characteristics, etc.

☞ Provides **physical data independence**

Not a “Procedural” or “Operational” language like C++, Java, Python

Basic queries: SFW statement

- **SELECT** A_1, A_2, \dots, A_n
FROM R_1, R_2, \dots, R_m
WHERE *condition*

- **SELECT, FROM, WHERE** are often referred to as
SELECT, FROM, WHERE “clauses”

Example: reading a table

- **SELECT ***
FROM Serves

Serves

bar	beer	price
The Edge	Budweiser	2.50
The Edge	Corona	3.00
Satisfaction	Budweiser	2.25

- Single-table query
- **WHERE** clause is optional
- ***** is a short hand for “all columns”

Example: selecting few rows

- `SELECT beer AS mybeer`
`FROM Serves`
`WHERE price < 2.75`

Serves

bar	beer	price
The Edge	Budweiser	2.50
The Edge	Corona	3.00
Satisfaction	Budweiser	2.25

- `SELECT beer`
`FROM Serves`
`WHERE bar = 'The Edge'`

What does these return?

- `SELECT` list can contain expressions
Can also use built-in functions such as `SUBSTR`, `ABS`, etc.
- String literals (case sensitive) are enclosed in **single quotes**
- “AS” is optional
- Do not want duplicates? Write `SELECT DISTINCT beer ...`

Example: Join

- Find addresses of all bars that 'Dan' frequents
- Which tables do we need?

Example: Join

- Find addresses of all bars that 'Dan' frequents

Bar

name	address
The Edge	108 Morris Street
Satisfaction	905 W. Main Street

Beer

Name	brewer
Budweiser	Anheuser-Busch Inc.
Corona	Grupo Modelo
Dixie	Dixie Brewing

Drinker

name	address
Amy	100 W. Main Street
Ben	101 W. Main Street
Dan	300 N. Duke Street

bar	beer	price
The Edge	Budweiser	2.50
The Edge	Corona	3.00
Satisfaction	Budweiser	2.25

drinker	bar	times_a_week
Ben	Satisfaction	2
Dan	The Edge	1
Dan	Satisfaction	2

Frequents

drinker	beer
Amy	Corona
Dan	Budweiser
Dan	Corona
Ben	Budweiser

Likes

Which tables do we need?

How do we combine them?

Example: Join

- Find addresses of all bars that ‘Dan’ frequents

- ```
SELECT B.address
FROM Bar B, Frequents F
WHERE B.name = F.bar
 AND F.drinker = 'Dan'
```

- Okay to omit *table\_name* in *table\_name.column\_name* if *column\_name* is unique
- Can use “Aliases” for convenience
  - “Bar as B” or “Bar B”

**Bar**

| name         | address            |
|--------------|--------------------|
| The Edge     | 108 Morris Street  |
| Satisfaction | 905 W. Main Street |

| drinker | bar          | times_a_week |
|---------|--------------|--------------|
| Ben     | Satisfaction | 2            |
| Dan     | The Edge     | 1            |
| Dan     | Satisfaction | 2            |

**Frequents**

# Semantics of SFW

- SELECT  $E_1, E_2, \dots, E_n$   
FROM  $R_1, R_2, \dots, R_m$   
WHERE *condition*

- For each  $t_1$  in  $R_1$ :  
    For each  $t_2$  in  $R_2$ : ... ..  
    For each  $t_m$  in  $R_m$ :
  1. Apply “FROM”  
Form cross-product of  $R_1, \dots, R_m$

If *condition* is true over  $t_1, t_2, \dots, t_m$ :

2. Apply “WHERE”  
Only consider satisfying rows

Compute and output  $E_1, E_2, \dots, E_n$  as a row

3. Apply “SELECT”  
Output the desired columns

# Step 1: Illustration of Semantics of SFW

- NOTE: This is “NOT HOW” the DBMS outputs the result, but “WHAT” is outputs!

Form Cross product of two relations

- SELECT B.address  
FROM Bar B, Frequent F  
WHERE B.name = F.bar  
AND F.drinker = 'Dan'

## Bar

| name         | address            |
|--------------|--------------------|
| The Edge     | 108 Morris Street  |
| Satisfaction | 905 W. Main Street |

## Frequent

| rinker | bar          | times_a_week |
|--------|--------------|--------------|
| Ben    | Satisfaction | 2            |
| Dan    | The Edge     | 1            |
| Dan    | Satisfaction | 2            |

| name         | address            | rinker | bar          | times_a_week |
|--------------|--------------------|--------|--------------|--------------|
| The Edge     | 108 Morris Street  | Ben    | Satisfaction | 2            |
| The Edge     | 108 Morris Street  | Dan    | The Edge     | 1            |
| The Edge     | 108 Morris Street  | Dan    | Satisfaction | 2            |
| Satisfaction | 905 W. Main Street | Ben    | Satisfaction | 2            |
| Satisfaction | 905 W. Main Street | Dan    | The Edge     | 1            |
| Satisfaction | 905 W. Main Street | Dan    | Satisfaction | 2            |



# Step 2: Illustration of Semantics of SFW

- NOTE: This is “NOT HOW” the DBMS outputs the result, but “WHAT” is outputs!

Discard rows that do not satisfy WHERE condition

- SELECT B.address  
FROM Bar B, Frequents F

WHERE B.name = F.bar  
AND F.drinker = 'Dan'

## Bar

| name         | address            |
|--------------|--------------------|
| The Edge     | 108 Morris Street  |
| Satisfaction | 905 W. Main Street |

## Frequents

| drinker | bar          | times_a_week |
|---------|--------------|--------------|
| Ben     | Satisfaction | 2            |
| Dan     | The Edge     | 1            |
| Dan     | Satisfaction | 2            |

| name                | address                      | drinker        | bar                     | times_a_week |
|---------------------|------------------------------|----------------|-------------------------|--------------|
| <del>The Edge</del> | <del>108 Morris Street</del> | <del>Ben</del> | <del>Satisfaction</del> | <del>2</del> |
| The Edge            | 108 Morris Street            | Dan            | The Edge                | 1            |
| <del>The Edge</del> | <del>108 Morris Street</del> | <del>Dan</del> | <del>Satisfaction</del> | <del>2</del> |
| Satisfaction        | 905 W. Main Street           | Ben            | Satisfaction            | 2            |
| Satisfaction        | 905 W. Main Street           | Dan            | The Edge                | 4            |
| Satisfaction        | 905 W. Main Street           | Dan            | Satisfaction            | 2            |

# Step 3: Illustration of Semantics of SFW

- NOTE: This is “NOT HOW” the DBMS outputs the result, but “WHAT” is outputs!

Output the “address” output of rows that survived

- SELECT B.address  
FROM Bar B, Frequents F  
WHERE B.name = F.bar  
AND F.drinker = 'Dan'

## Bar

| name         | address            |
|--------------|--------------------|
| The Edge     | 108 Morris Street  |
| Satisfaction | 905 W. Main Street |

## Frequents

| drinker | bar          | times_a_week |
|---------|--------------|--------------|
| Ben     | Satisfaction | 2            |
| Dan     | The Edge     | 1            |
| Dan     | Satisfaction | 2            |

| name                | address                      | drinker        | bar                     | times_a_week |
|---------------------|------------------------------|----------------|-------------------------|--------------|
| <del>The Edge</del> | <del>108 Morris Street</del> | <del>Ben</del> | <del>Satisfaction</del> | <del>2</del> |
| The Edge            | 108 Morris Street            | Dan            | The Edge                | 1            |
| <del>The Edge</del> | <del>108 Morris Street</del> | <del>Dan</del> | <del>Satisfaction</del> | <del>2</del> |
| Satisfaction        | 905 W. Main Street           | Ben            | Satisfaction            | 2            |
| Satisfaction        | 905 W. Main Street           | Dan            | The Edge                | 4            |
| Satisfaction        | 905 W. Main Street           | Dan            | Satisfaction            | 2            |

# Final output: Illustration of Semantics of SFW

- NOTE: This is “NOT HOW” the DBMS outputs the result, but “WHAT” is outputs!

Output the “address” output of rows that survived

- SELECT B.address  
FROM Bar B, Frequents F  
WHERE B.name = F.bar  
AND F.drinker = 'Dan'

## Bar

| name         | address            |
|--------------|--------------------|
| The Edge     | 108 Morris Street  |
| Satisfaction | 905 W. Main Street |

| address            |
|--------------------|
| 108 Morris Street  |
| 905 W. Main Street |

## Frequents

| drinker | bar          | times_a_week |
|---------|--------------|--------------|
| Ben     | Satisfaction | 2            |
| Dan     | The Edge     | 1            |
| Dan     | Satisfaction | 2            |

Try some SQL queries yourself on pgweb!

(See how to access the pgweb interface for a small “Beers” database on the slides posted on the course website)

Next: semantics of SFW statements in SQL

# Announcements (Tue, 01/09)

- You should be on Sakai, Piazza, Gradescope
  - If you are not there or recently enrolled, please contact the instructor
- You will receive instructions on installing the VM
  - Please follow Piazza posts, all notifications will be posted there and you should receive emails right away
- First homework to be released on next class  
Tuesday 01/14, due in a week
  - No in-class quiz or labs unless explicitly announced in the class before (and posted on Piazza)