

Relational Databases

Charles Severance




Unless otherwise noted, the content of this course material is licensed under a Creative Commons Attribution 3.0 License.
<http://creativecommons.org/licenses/by/3.0/>.

Copyright 2009- Charles Severance



SQLite Manager For Firefox



<https://addons.mozilla.org/en-US/firefox/addon/sqlite-manager/>


Relational Databases

Relational databases model data by storing rows and columns in tables. The power of the relational database lies in its ability to efficiently retrieve data from those tables and in particular where there are multiple tables and the relationships between those tables involved in the query.

http://en.wikipedia.org/wiki/Relational_database

Terminology

- Database - Contains many tables
- Relation (or table) - contains tuples and attributes
- Tuple (or row) - is a set of fields it generally represents an “object” like a person or a music track
- Attribute (also column or field) - One of possibly many elements of data corresponding to the object represented by the row



A relation is defined as a set of tuples that have the same attributes. A tuple usually represents an object and information about that object. Objects are typically physical objects or concepts. A relation is usually described as a table, which is organized into rows and columns. All the data referenced by an attribute are in the same domain and conform to the same constraints. (wikipedia)

The screenshot shows a spreadsheet application window titled 'SIS02 - Database'. The spreadsheet has columns labeled A, B, C, and D, and rows numbered 1 through 8. A table is displayed with the following data:

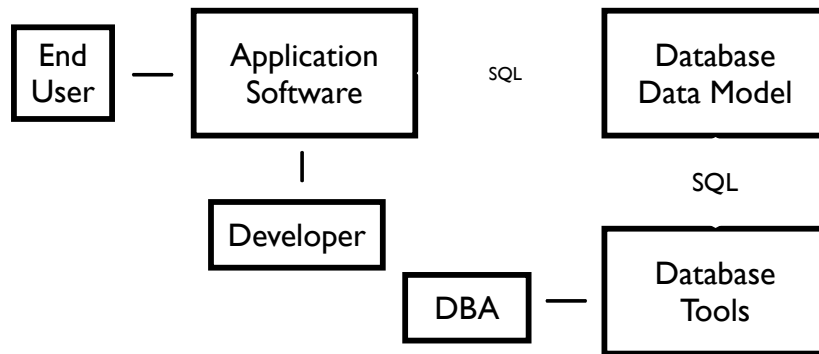
Columns / Attributes			
TITLE	RATING	LEN	
About to Rock	3	354	Rows /
Who Made Who	4	252	Tuples

At the bottom of the spreadsheet, there is a section labeled 'Tables / Relations' with tabs for 'Tracks', 'Albums', 'Artists', and 'Genres'.

Two Roles in Large Projects

- Application Developer - Builds the logic for the application, the look and feel of the application - monitors the application for problems
- Database Administrator - Monitors and adjusts the database as the program runs in production
- Often both people participate in the building of the “Data model”

Application Structure



Database Administrator (dba)

A database administrator (DBA) is a person responsible for the design, implementation, maintenance and repair of an organization's database. The role includes the development and design of database strategies, monitoring and improving database performance and capacity, and planning for future expansion requirements. They may also plan, co-ordinate and implement security measures to safeguard the database.

http://en.wikipedia.org/wiki/Database_administrator

Database Model

A database model or database schema is the structure or format of a database, described in a formal language supported by the database management system. In other words, a "database model" is the application of a data model when used in conjunction with a database management system.

http://en.wikipedia.org/wiki/Database_model

SQL

- Structured Query Language is the language we use to issue commands to the database
- Create a table
- Retrieve some data
- Insert data
- Delete data

<http://en.wikipedia.org/wiki/SQL>

Common Database Systems

- Three Major Database Management Systems in wide use
 - Oracle - Large, commercial, enterprise-scale, very very tweakable
 - MySQL - Simpler but very fast and scalable - commercial open source
 - SqlServer - Very nice - from Microsoft (also Access)
- Many other smaller projects, free and open source
 - HSQL, SQLite, Postgress, ...

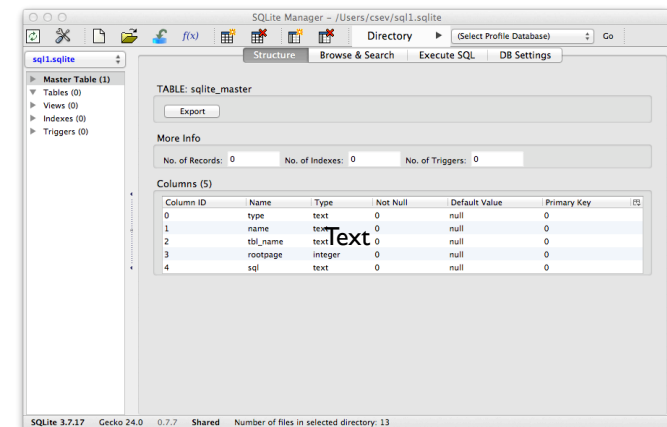
SQLite Database Manager

- SQLite is a very popular database - it is free and fast and small
- We have a FireFox plugin to manipulate SQLite databases
 - <https://addons.mozilla.org/en-US/firefox/addon/sqlite-manager/>
- SQLite is embedded in Python and a number of other languages

SQLite is in lots of software...




<http://www.sqlite.org/famous.html>



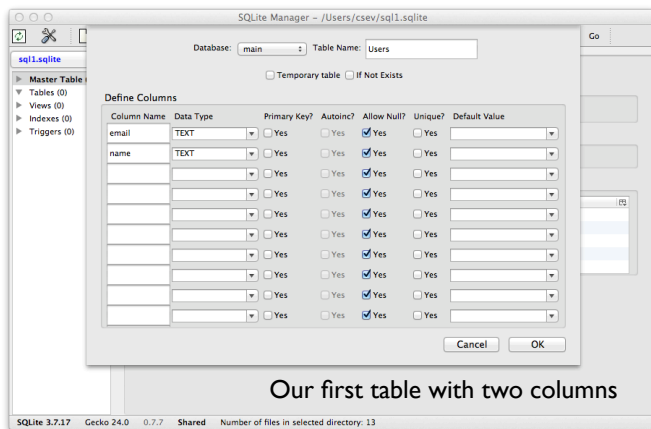
<https://addons.mozilla.org/en-US/firefox/addon/sqlite-manager/>

Application Structure

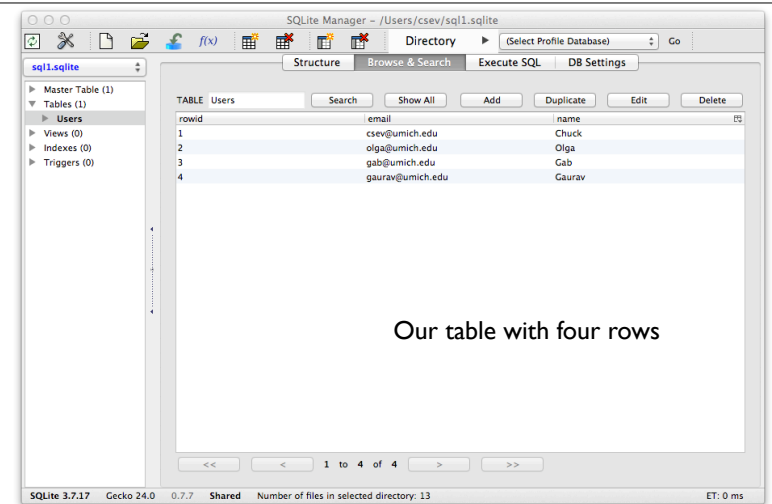


Start Simple - A Single Table

- Lets make a table of People - with a Name and an E-Mail



Our first table with two columns



Our table with four rows

SQL

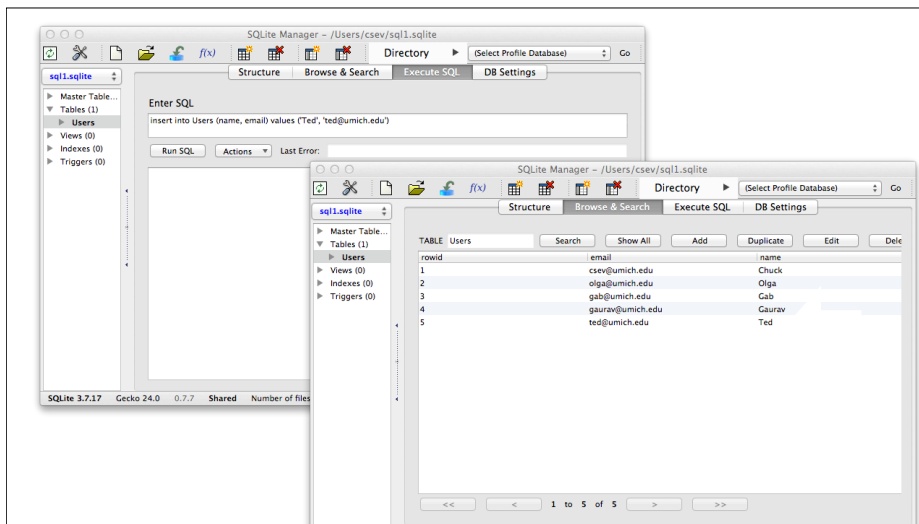
- Structured Query Language is the language we use to issue commands to the database
- Create a table
- Retrieve some data
- Insert data
- Delete data

<http://en.wikipedia.org/wiki/SQL>

SQL Insert

- The Insert statement inserts a row into a table

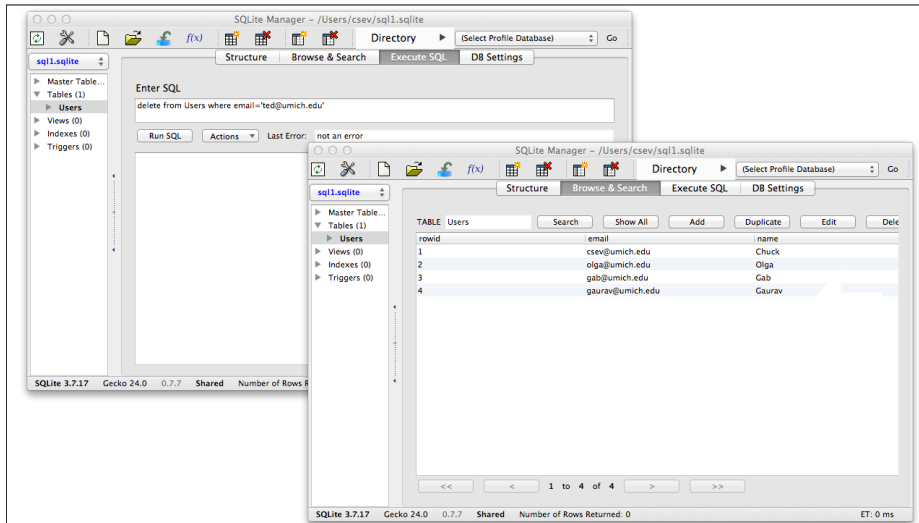
`insert into Users (name, email) values ('Ted', 'ted@umich.edu')`



SQL Delete

- Deletes a row in a table based on a selection criteria

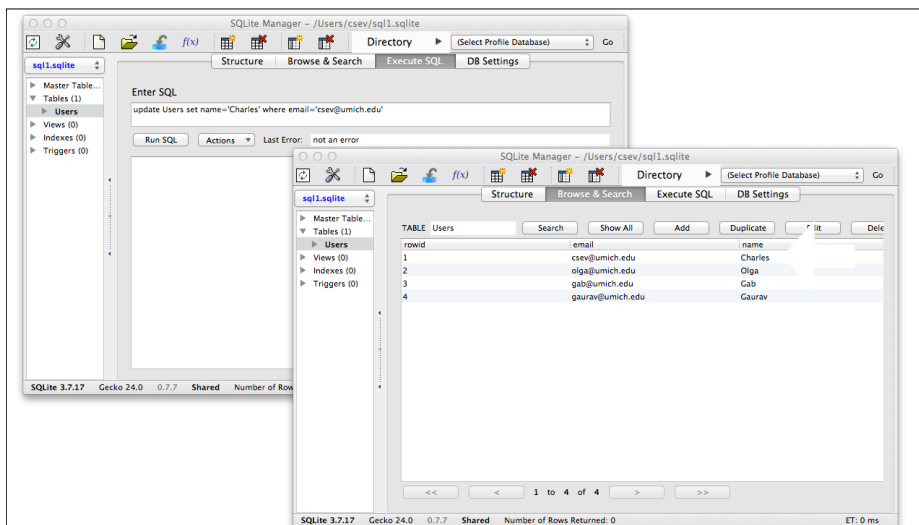
`delete from Users where email='ted@umich.edu'`



SQL: Update

- Allows the updating of a field with a where clause

update Users set name='Charles' where email='csev@umich.edu'

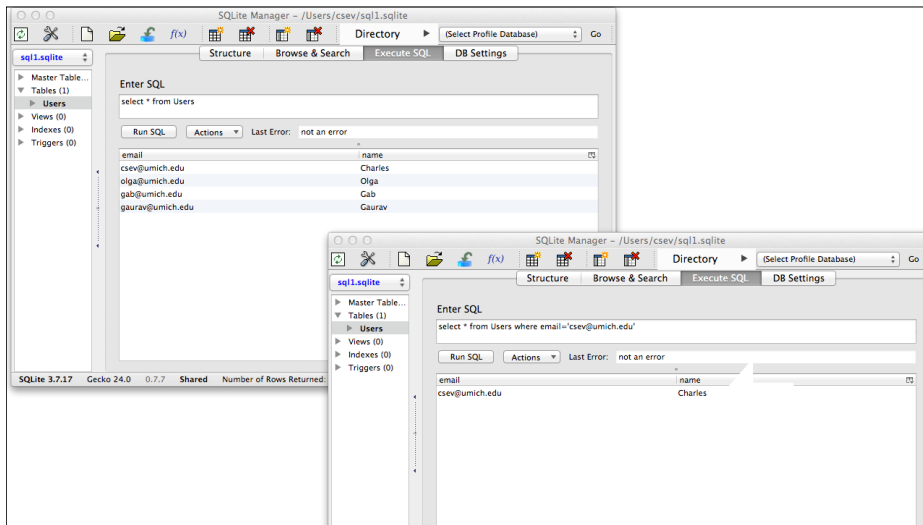


Retrieving Records: Select

- The select statement retrieves a group of records - you can either retrieve all the records or a subset of the records with a WHERE clause

select * from Users

select * from Users where email='csev@umich.edu'

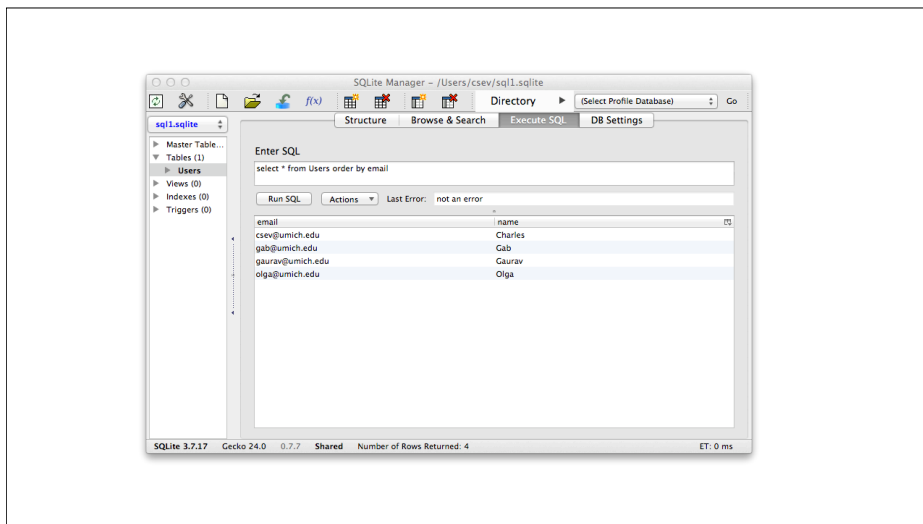


Sorting with ORDER BY

- You can add an ORDER BY clause to SELECT statements to get the results sorted in ascending or descending order

```
select * from Users order by email
```

```
select * from Users order by name
```



SQL Summary

```
insert into Users (name, email) values ('Ted', 'ted@umich.edu')
delete from Users where email='ted@umich.edu'
update Users set name="Charles" where email='csev@umich.edu'
select * from Users
select * from Users where email='csev@umich.edu'
select * from Users order by email
```


This is not too exciting (so far)

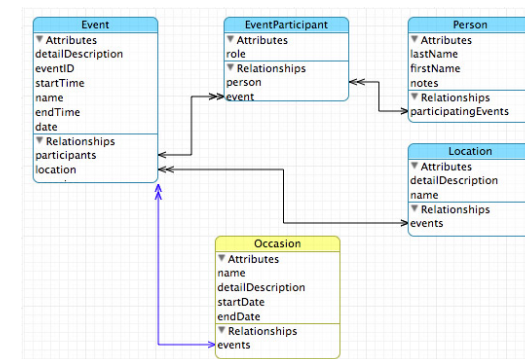
- Tables pretty much look like big fast programmable spreadsheet with rows, columns, and commands
- The power comes when we have more than one table and we can exploit the relationships between the tables

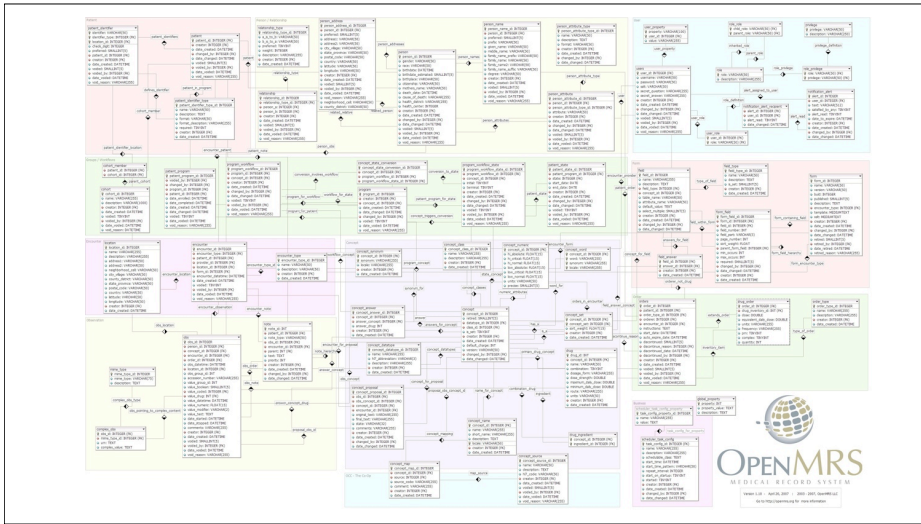
Complex Data Models and Relationships

http://en.wikipedia.org/wiki/Relational_model

Database Design

- Database design is an art form of its own with particular skills and experience
- Our goal is to avoid the really bad mistakes and design clean and easily understood databases
- Others may performance tune things later
- Database design starts with a picture...






Building a Data Model

- Drawing a picture of the data objects for our application and then figuring out how to represent the objects and their relationships
- Basic Rule: Don't put the same string data in twice - use a relationship instead
- When there is one thing in the "real world" there should be one copy of that thing in the database

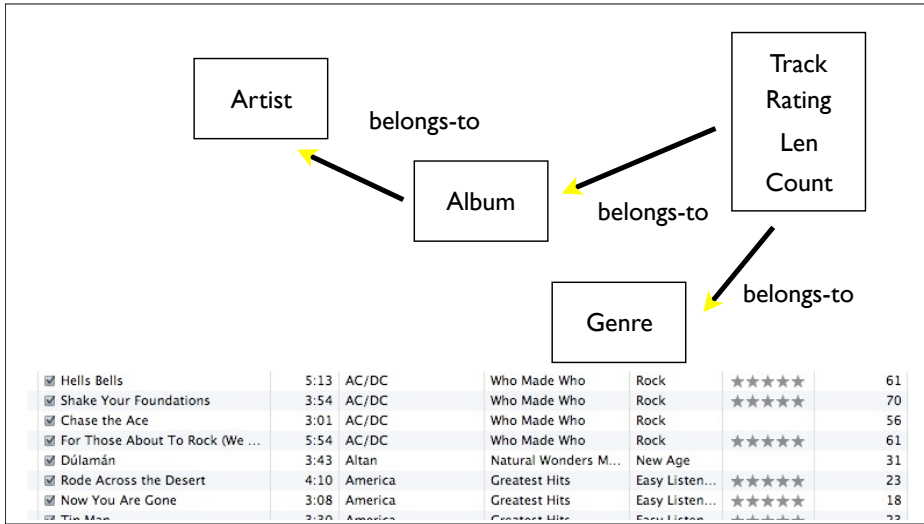
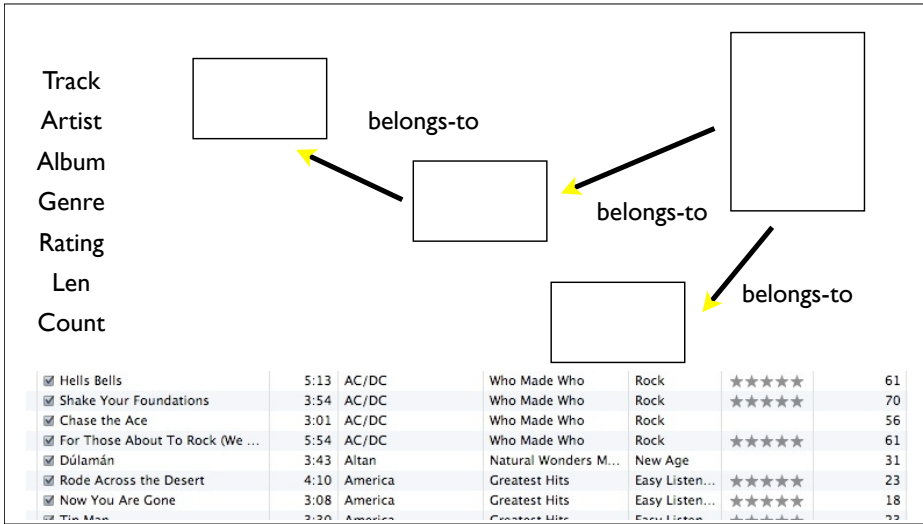
Track	Len	Artist	Album	Genre	Rating	Count
<input checked="" type="checkbox"/> Hells Bells	5:13	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/> Shake Your Foundations	3:54	AC/DC	Who Made Who	Rock	★★★★★	70
<input checked="" type="checkbox"/> Chase the Ace	3:01	AC/DC	Who Made Who	Rock	★★★★★	56
<input checked="" type="checkbox"/> For Those About To Rock (We ...	5:54	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/> Dúlamán	3:43	Altan	Natural Wonders M...	New Age	★★★★★	31
<input checked="" type="checkbox"/> Rode Across the Desert	4:10	America	Greatest Hits	Easy Listen...	★★★★★	23
<input checked="" type="checkbox"/> Now You Are Gone	3:08	America	Greatest Hits	Easy Listen...	★★★★★	18
<input checked="" type="checkbox"/> Tin Man	3:30	America	Greatest Hits	Easy Listen...	★★★★★	23
<input checked="" type="checkbox"/> Sister Golden Hair	3:22	America	Greatest Hits	Easy Listen...	★★★★★	24
<input checked="" type="checkbox"/> Track 01	4:22	Billy Price	Danger Zone	Blues/R&B	★★★★★	26
<input checked="" type="checkbox"/> Track 02	2:45	Billy Price	Danger Zone	Blues/R&B	★★★★★	18
<input checked="" type="checkbox"/> Track 03	3:26	Billy Price	Danger Zone	Blues/R&B	★★★★★	22
<input checked="" type="checkbox"/> Track 04	4:17	Billy Price	Danger Zone	Blues/R&B	★★★★★	18
<input checked="" type="checkbox"/> Track 05	3:50	Billy Price	Danger Zone	Blues/R&B	★★★★★	21
<input checked="" type="checkbox"/> War Pigs/Luke's Wall	7:58	Black Sabbath	Paranoid	Metal	★★★★★	25
<input checked="" type="checkbox"/> Paranoid	2:53	Black Sabbath	Paranoid	Metal	★★★★★	22
<input checked="" type="checkbox"/> Planet Caravan	4:35	Black Sabbath	Paranoid	Metal	★★★★★	25
<input checked="" type="checkbox"/> Iron Man	5:59	Black Sabbath	Paranoid	Metal	★★★★★	26
<input checked="" type="checkbox"/> Electric Funeral	4:53	Black Sabbath	Paranoid	Metal	★★★★★	22
<input checked="" type="checkbox"/> Hand of Doom	7:10	Black Sabbath	Paranoid	Metal	★★★★★	23
<input checked="" type="checkbox"/> Rat Salad	2:30	Black Sabbath	Paranoid	Metal	★★★★★	31
<input checked="" type="checkbox"/> Jack the Stripper/Fairies Wear ...	6:14	Black Sabbath	Paranoid	Metal	★★★★★	24
<input checked="" type="checkbox"/> Bomb Squad (TECH)	3:28	Brent	Brent's Album			1
<input checked="" type="checkbox"/> clay techno	4:36	Brent	Brent's Album			2
<input checked="" type="checkbox"/> Heavy	3:08	Brent	Brent's Album			1
<input checked="" type="checkbox"/> Hi metal man	4:20	Brent	Brent's Album			1
<input checked="" type="checkbox"/> Mistro	2:58	Brent	Brent's Album			1

For each "piece of info"...

- Is the column an object or an attribute of another object?
- Once we define objects we need to define the relationships between objects.



<input checked="" type="checkbox"/> Hells Bells	5:13	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/> Shake Your Foundations	3:54	AC/DC	Who Made Who	Rock	★★★★★	70
<input checked="" type="checkbox"/> Chase the Ace	3:01	AC/DC	Who Made Who	Rock	★★★★★	56
<input checked="" type="checkbox"/> For Those About To Rock (We ...	5:54	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/> Dúlamán	3:43	Altan	Natural Wonders M...	New Age	★★★★★	31
<input checked="" type="checkbox"/> Rode Across the Desert	4:10	America	Greatest Hits	Easy Listen...	★★★★★	23
<input checked="" type="checkbox"/> Now You Are Gone	3:08	America	Greatest Hits	Easy Listen...	★★★★★	18
<input checked="" type="checkbox"/> Tin Man	3:30	America	Greatest Hits	Easy Listen...	★★★★★	23



Representing Relationships in a Database

<input checked="" type="checkbox"/>	Hells Bells	5:13	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/>	Shake Your Foundations	3:54	AC/DC	Who Made Who	Rock	★★★★★	70
<input checked="" type="checkbox"/>	Chase the Ace	3:01	AC/DC	Who Made Who	Rock	★★★★★	56
<input checked="" type="checkbox"/>	For Those About To Rock (We ...	5:54	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/>	Dúlaman	3:43	Altan	Natural Wonders M...	New Age	★★★★★	31
<input checked="" type="checkbox"/>	Rode Across the Desert	4:10	America	Greatest Hits	Easy Listen...	★★★★★	23
<input checked="" type="checkbox"/>	Now You Are Gone	3:08	America	Greatest Hits	Easy Listen...	★★★★★	18
<input checked="" type="checkbox"/>	Tie Me	3:20	America	Greatest Hits	Easy Listen...	★★★★★	22

We want to keep track of which band is the “creator” of each music track...
What album does this song “belong to”??

Which album is this song related to?

Database Normalization (3NF)

- There is *tons* of database theory - way too much to understand without excessive predicate calculus
- Do not replicate data - reference data - point at data
- Use integers for keys and for references
- Add a special "key" column to each table which we will make references to. By convention many programmers call this column "id"

http://en.wikipedia.org/wiki/Database_normalization

Integer Reference Pattern

We use integers to reference rows in another table.

id	artist_id	title
1	2	Who Made Who
2	1	IV

id	title
1	Led Zeppelin
2	AC/DC

Artist

Album

Keys

Finding our way around...

Three Kinds of Keys

- Primary key - generally an integer auto-increment field
- Logical key - What the outside world uses for lookup
- Foreign key - generally an integer key point to a row in another table

Site
id
title
user_id
...


Primary Key Rules

- Best practices
- Never use your logical key as the primary key
- Logical keys can and do change albeit slowly
- Relationships that are based on matching string fields are far less efficient than integers performance-wise

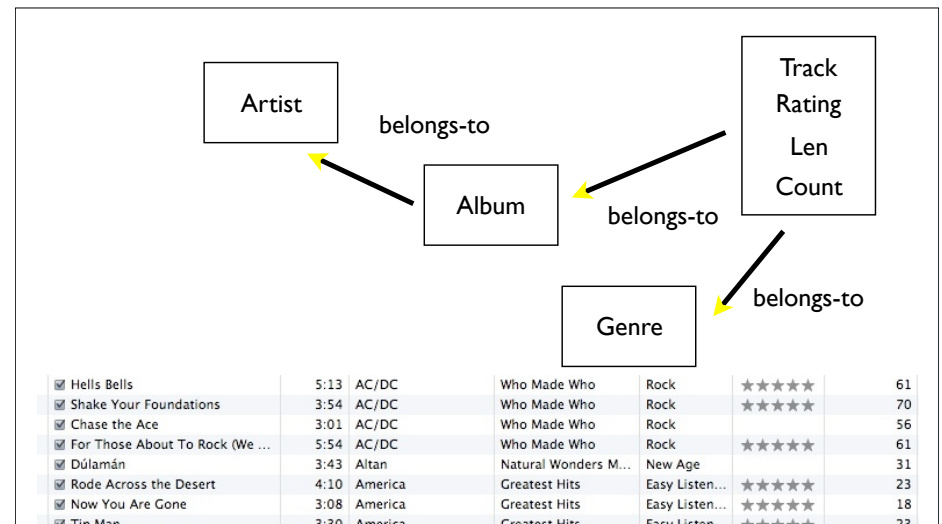
```
User
id
login
password
name
email
created_at
modified_at
login_at
```

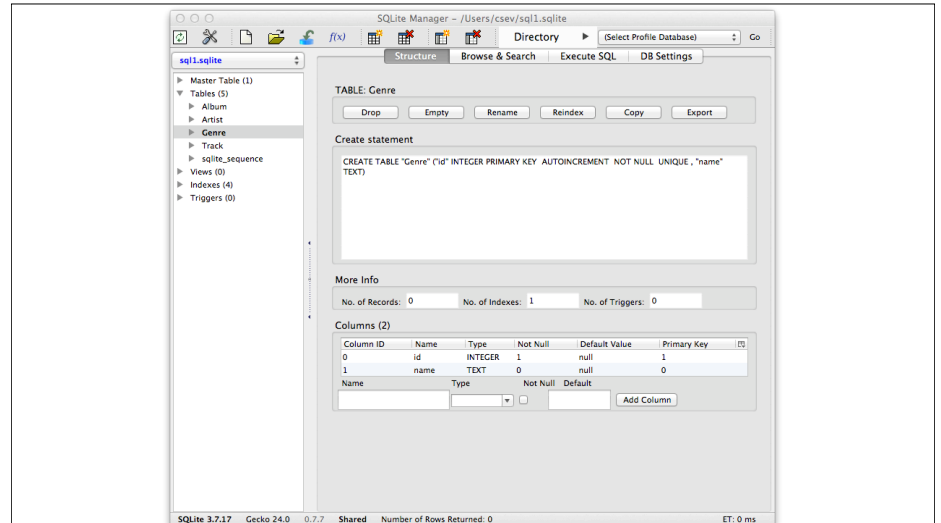
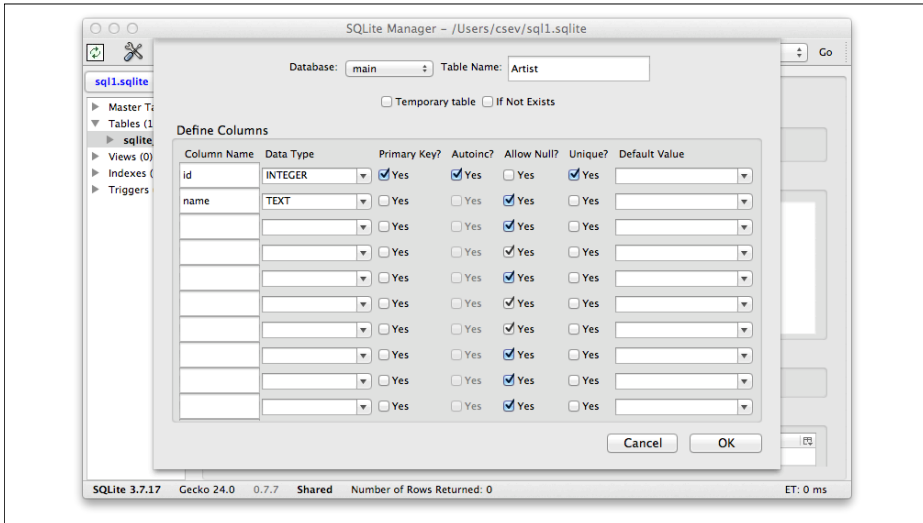
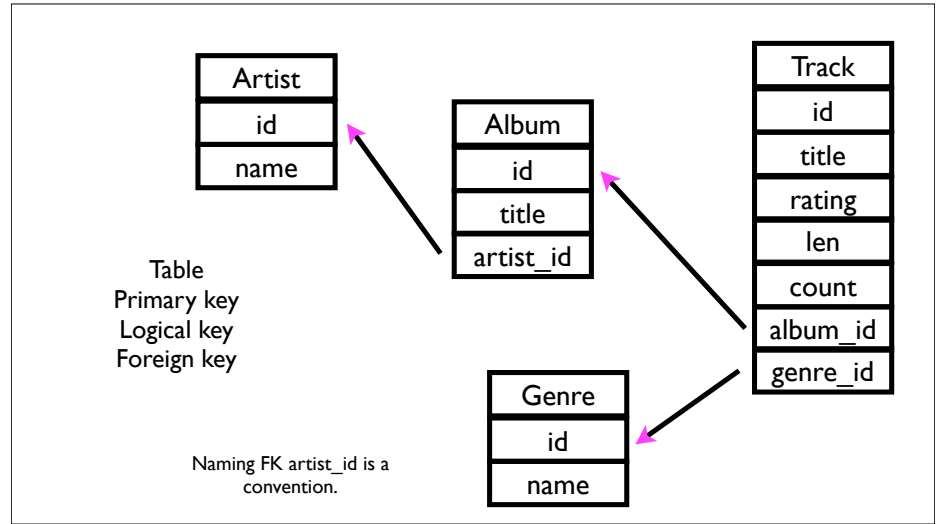
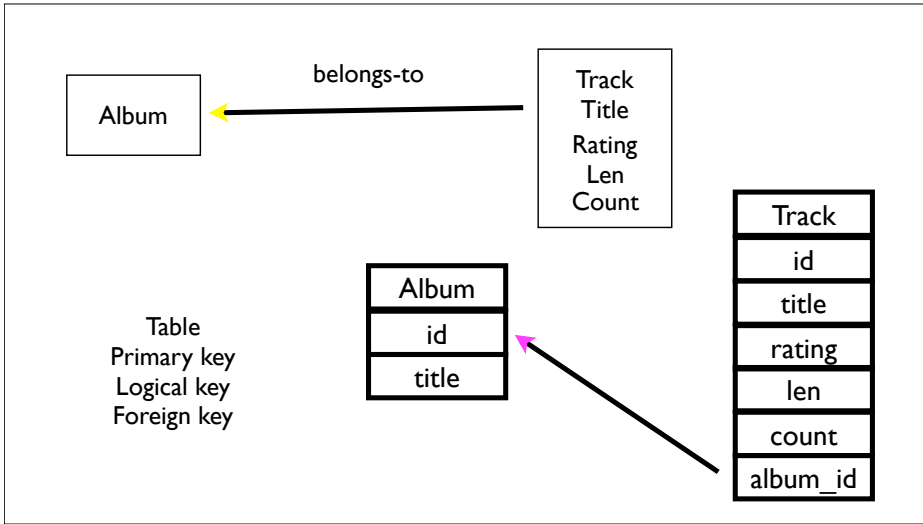
Foreign Keys

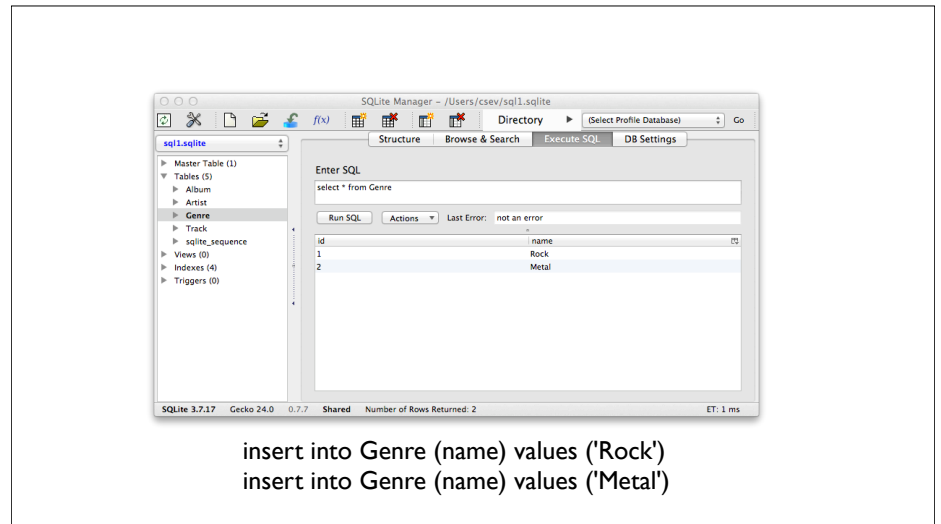
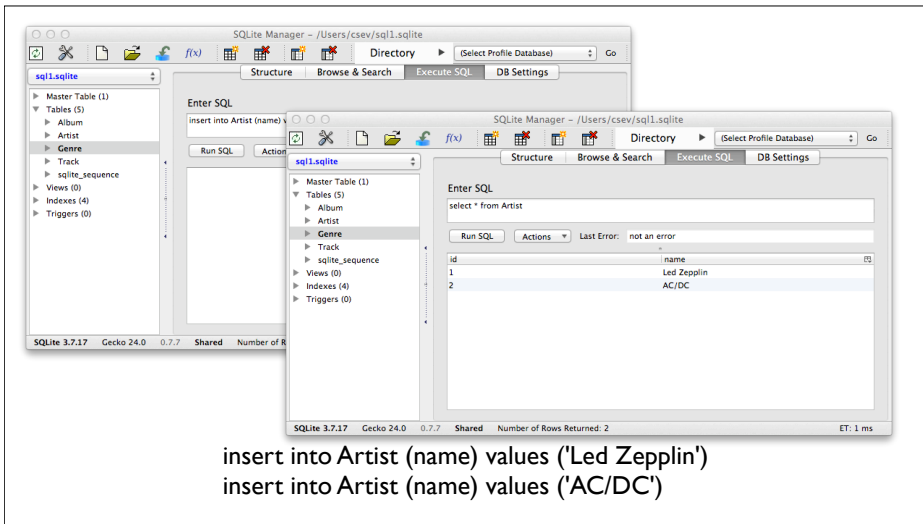
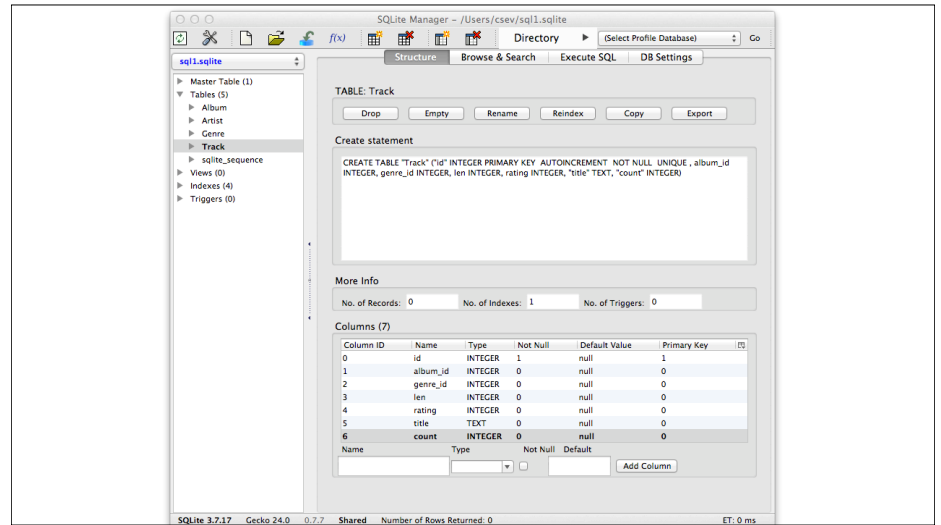
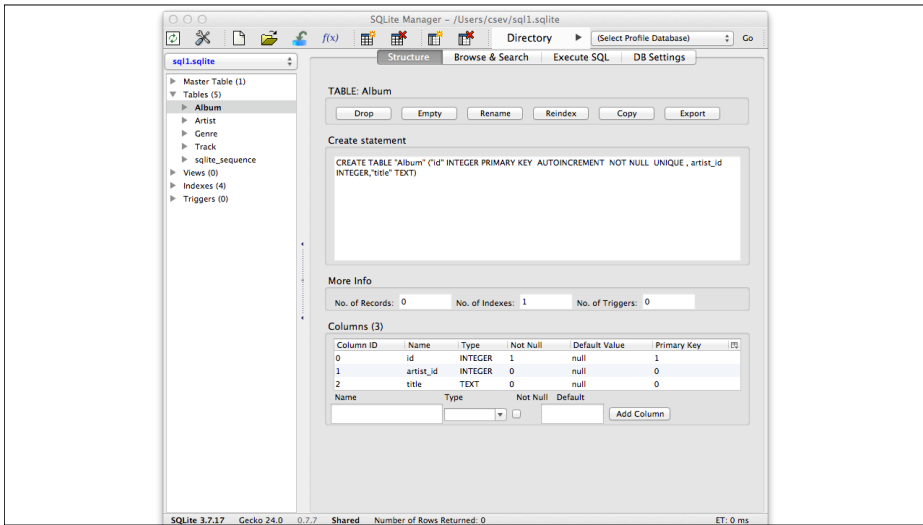
- A foreign key is when a table has a column that contains a key which points the primary key of another table.
- When all primary keys are integers, then all foreign keys are integers - this is good - very good
- If you use strings as foreign keys - you show yourself to be an uncultured swine

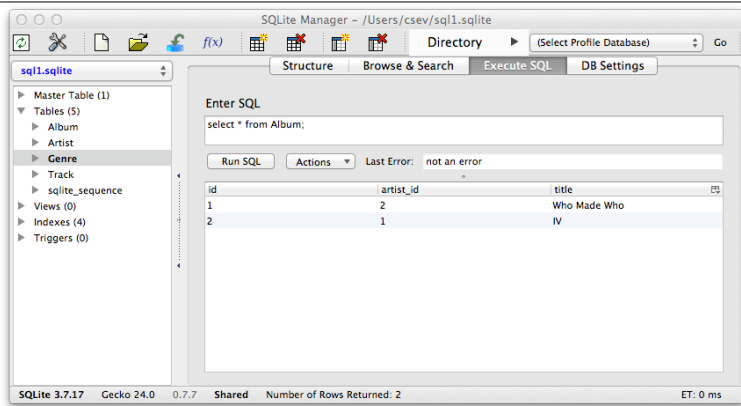


Relationship Building (in tables)







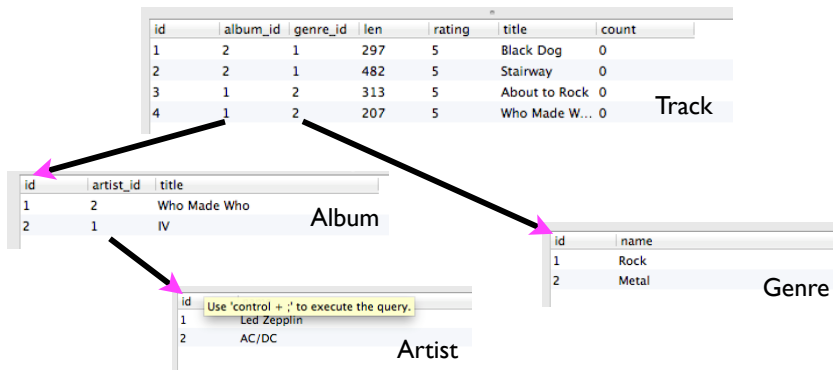


insert into Album (title, artist_id) values ('Who Made Who', 2)
 insert into Album (title, artist_id) values ('IV', 1)

insert into Track (title, rating, len, count, album_id, genre_id)
 values ('Black Dog', 5, 297, 0, 2, 1)
 insert into Track (title, rating, len, count, album_id, genre_id)
 values ('Stairway', 5, 482, 0, 2, 1)
 insert into Track (title, rating, len, count, album_id, genre_id)
 values ('About to Rock', 5, 313, 0, 1, 2)
 insert into Track (title, rating, len, count, album_id, genre_id)
 values ('Who Made Who', 5, 207, 0, 1, 2)

id	album_id	genre_id	len	rating	title	count
1	2	1	297	5	Black Dog	0
2	2	1	482	5	Stairway	0
3	1	2	313	5	About to Rock	0
4	1	2	207	5	Who Made W...	0

We have relationships!



Using Join Across Tables

[http://en.wikipedia.org/wiki/Join_\(SQL\)](http://en.wikipedia.org/wiki/Join_(SQL))

Relational Power

- By removing the replicated data and replacing it with references to a single copy of each bit of data we build a “web” of information that the relational database can read through very quickly - even for very large amounts of data
- Often when you want some data it comes from a number of tables linked by these foreign keys

The JOIN Operation

- The JOIN operation links across several tables as part of a select operation
- You must tell the JOIN how to use the keys that make the connection between the tables using an ON clause

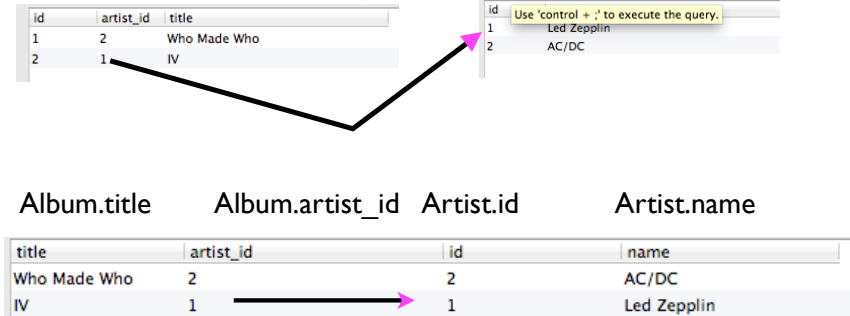


select Album.title, Artist.name from Album join Artist on Album.artist_id = Artist.id

What we want to see

The tables which hold the data

How the tables are linked



select Album.title, Album.artist_id, Artist.id, Artist.name from Album join Artist on Album.artist_id = Artist.id

id	album_id	genre_id	len	rating	title	count
1	2	1	297	5	Black Dog	0
2	2	1	482	5	Stairway	0
3	1	2	313	5	About to Rock	0
4	1	2	207	5	Who Made W...	0

title	name
Black Dog	Rock
Stairway	Rock
About to Rock	Metal
Who Made Who	Metal

id	name
1	Rock
2	Metal

select Track.title, Genre.name from Track join Genre on Track.genre_id = Genre.id

What we want to see

The tables which hold the data

How the tables are linked

It can get complex...

select Track.title, Artist.name, Album.title, Genre.name from Track join Genre join Album join Artist on Track.genre_id = Genre.id and Track.album_id = Album.id and Album.artist_id = Artist.id

What we want to see

title	name	title	name
Black Dog	Led Zeppelin	IV	Rock
Stairway	Led Zeppelin	IV	Rock
About to Rock	AC/DC	Who Made Who	Metal
Who Made Who	AC/DC	Who Made Who	Metal

The tables which hold the data

How the tables are linked

<input checked="" type="checkbox"/>	Hells Bells	5:13	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/>	Shake Your Foundations	3:54	AC/DC	Who Made Who	Rock	★★★★★	70
<input checked="" type="checkbox"/>	Chase the Ace	3:01	AC/DC	Who Made Who	Rock	★★★★★	56
<input checked="" type="checkbox"/>	For Those About To Rock (We ...	5:54	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/>	Dúlamán	3:43	Altan	Natural Wonders M...	New Age	★★★★★	31
<input checked="" type="checkbox"/>	Rode Across the Desert	4:10	America	Greatest Hits	Easy Listen...	★★★★★	23
<input checked="" type="checkbox"/>	Now You Are Gone	3:08	America	Greatest Hits	Easy Listen...	★★★★★	18
<input checked="" type="checkbox"/>	Tin Man	3:30	America	Greatest Hits	Easy Listen...	★★★★★	23
<input checked="" type="checkbox"/>	Sister Golden Hair	3:22	America	Greatest Hits	Easy Listen...	★★★★★	24
<input checked="" type="checkbox"/>	Track 01	4:22	Billy Price	Danger Zone	Blues/R&B	★★★★★	26
<input checked="" type="checkbox"/>	Track 02	2:45	Billy Price	Danger Zone	Blues/R&B	★★★★★	18
<input checked="" type="checkbox"/>	Track 03	3:26	Billy Price	Danger Zone	Blues/R&B	★★★★★	22
<input checked="" type="checkbox"/>	Track 04	4:17	Billy Price	Danger Zone	Blues/R&B	★★★★★	18
<input checked="" type="checkbox"/>	Track 05	3:50	Billy Price	Danger Zone	Blues/R&B	★★★★★	21
<input checked="" type="checkbox"/>	War Pigs/Luke's Wall	7:58	Black Sabbath	Paranoid	Metal	★★★★★	25
<input checked="" type="checkbox"/>	Paranoid	2:53	Black Sabbath	Paranoid	Metal	★★★★★	22
<input checked="" type="checkbox"/>	Planet Caravan						
<input checked="" type="checkbox"/>	Iron Man						
<input checked="" type="checkbox"/>	Electric Funeral						
<input checked="" type="checkbox"/>	Hand of Doom						
<input checked="" type="checkbox"/>	Rat Salad						
<input checked="" type="checkbox"/>	Jack the Stripper/Fal						
<input checked="" type="checkbox"/>	Bomb Squad (TECH)						
<input checked="" type="checkbox"/>	clay techno						
<input checked="" type="checkbox"/>	Heavy						
<input checked="" type="checkbox"/>	Hi metal man	4:20	Brent	Brent's Album			1
<input checked="" type="checkbox"/>	Mistro	2:58	Brent	Brent's Album			1

title	name	title	name
Black Dog	Led Zeppelin	IV	Rock
Stairway	Led Zeppelin	IV	Rock
About to Rock	AC/DC	Who Made Who	Metal
Who Made Who	AC/DC	Who Made Who	Metal

Complexity Enables Speed

- Complexity makes speed possible and allows you to get very fast results as the data size grows.
- By normalizing the data and linking it with integer keys, the overall amount of data which the relational database must scan is far lower than if the data were simply flattened out.
- It might seem like a tradeoff - spend some time designing your database so it continues to be fast when your application is a success

Additional SQL Topics

- Indexes improve access performance for things like string fields
- Constraints on data - (cannot be NULL, etc..)
- Transactions - allow SQL operations to be grouped and done as a unit
- See SI664 - Database Design (All Semesters)

Summary

- Relational databases allow us to scale to very large amounts of data
- The key is to have one copy of any data element and use relations and joins to link the data to multiple places
- This greatly reduces the amount of data which must be scanned when doing complex operations across large amounts of data
- Database and SQL design is a bit of an art-form