

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.

You assume all responsibility for use and potential liability associated with any use of the material. Material contains copyrighted content, used in accordance with U.S. law. Copyright holders of content included in this material should contact open.michigan@umich.edu with any questions, corrections, or clarifications regarding the use of content. The Regents of the University of Michigan do not license the use of third party content posted to this site unless such a license is specifically granted in connection with particular content. Users of content are responsible for their compliance with applicable law. Mention of specific products in this material solely represents the opinion of the speaker and does not represent an endorsement by the University of Michigan. For more information about how to cite these materials visit <http://michigan.educommons.net/about/terms-of-use>.

Any medical information in this material is intended to inform and educate and is not a tool for self-diagnosis or a replacement for medical evaluation, advice, diagnosis or treatment by a healthcare professional. You should speak to your physician or make an appointment to be seen if you have questions or concerns about this information or your medical condition. Viewer discretion is advised: Material may contain medical images that may be disturbing to some viewers.

Relational Databases

Charles Severance

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

SQLite Database Browser

- SQLite is a very popular browser - it is free and fast and small
- We have a program to manipulate SQLite databases
- <http://sqlitebrowser.sourceforge.net/>
- SQLite is embedded in Python and a number of other languages

SQLite is in lots of software...

Symbian

Python

Philips

Skype

GE

Microsoft

McAfee

Apple

Adobe

Firefox

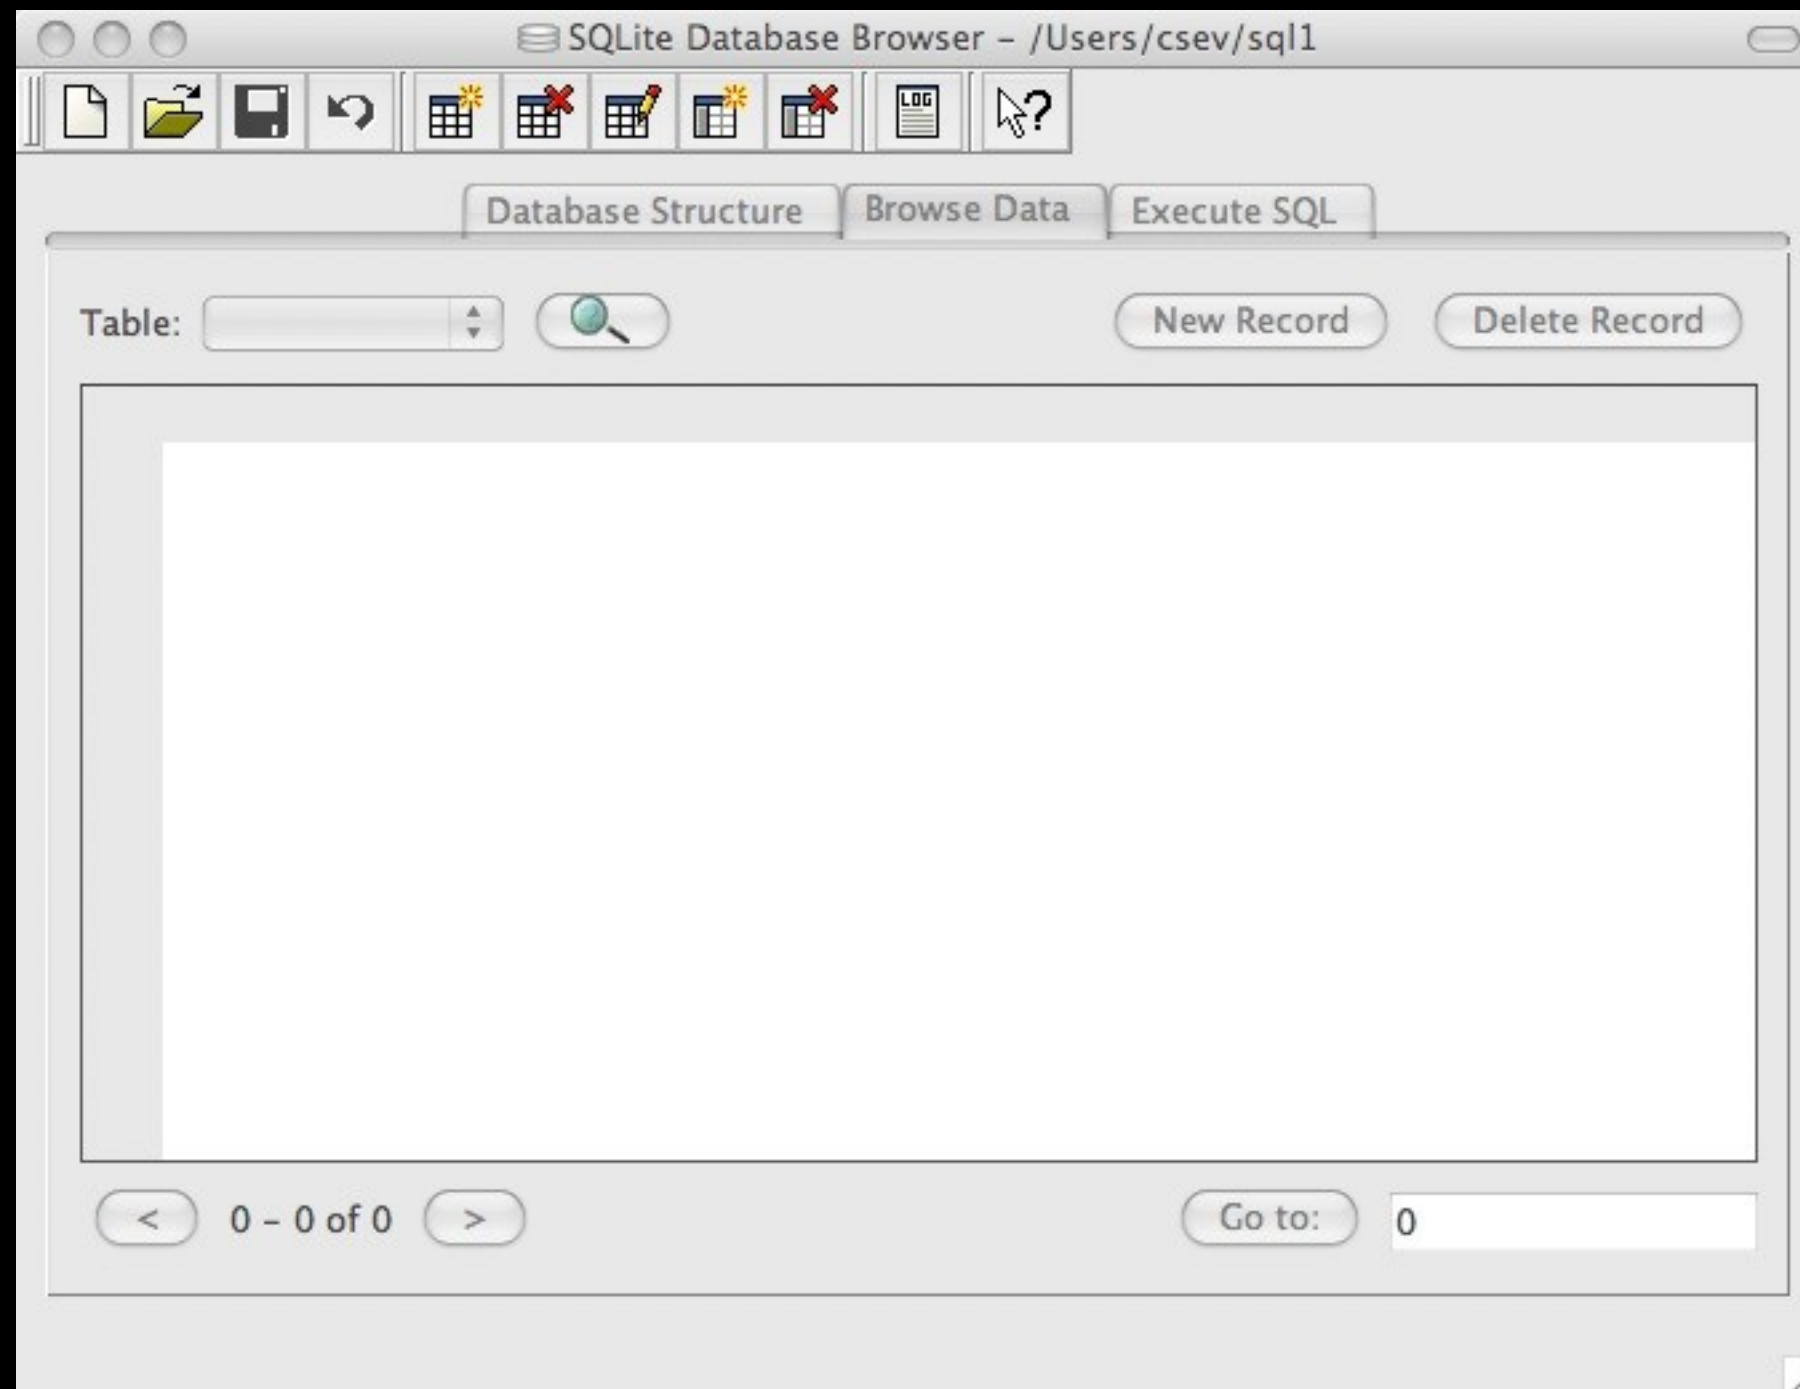
PHP

Toshiba

Sun Microsystems

Google

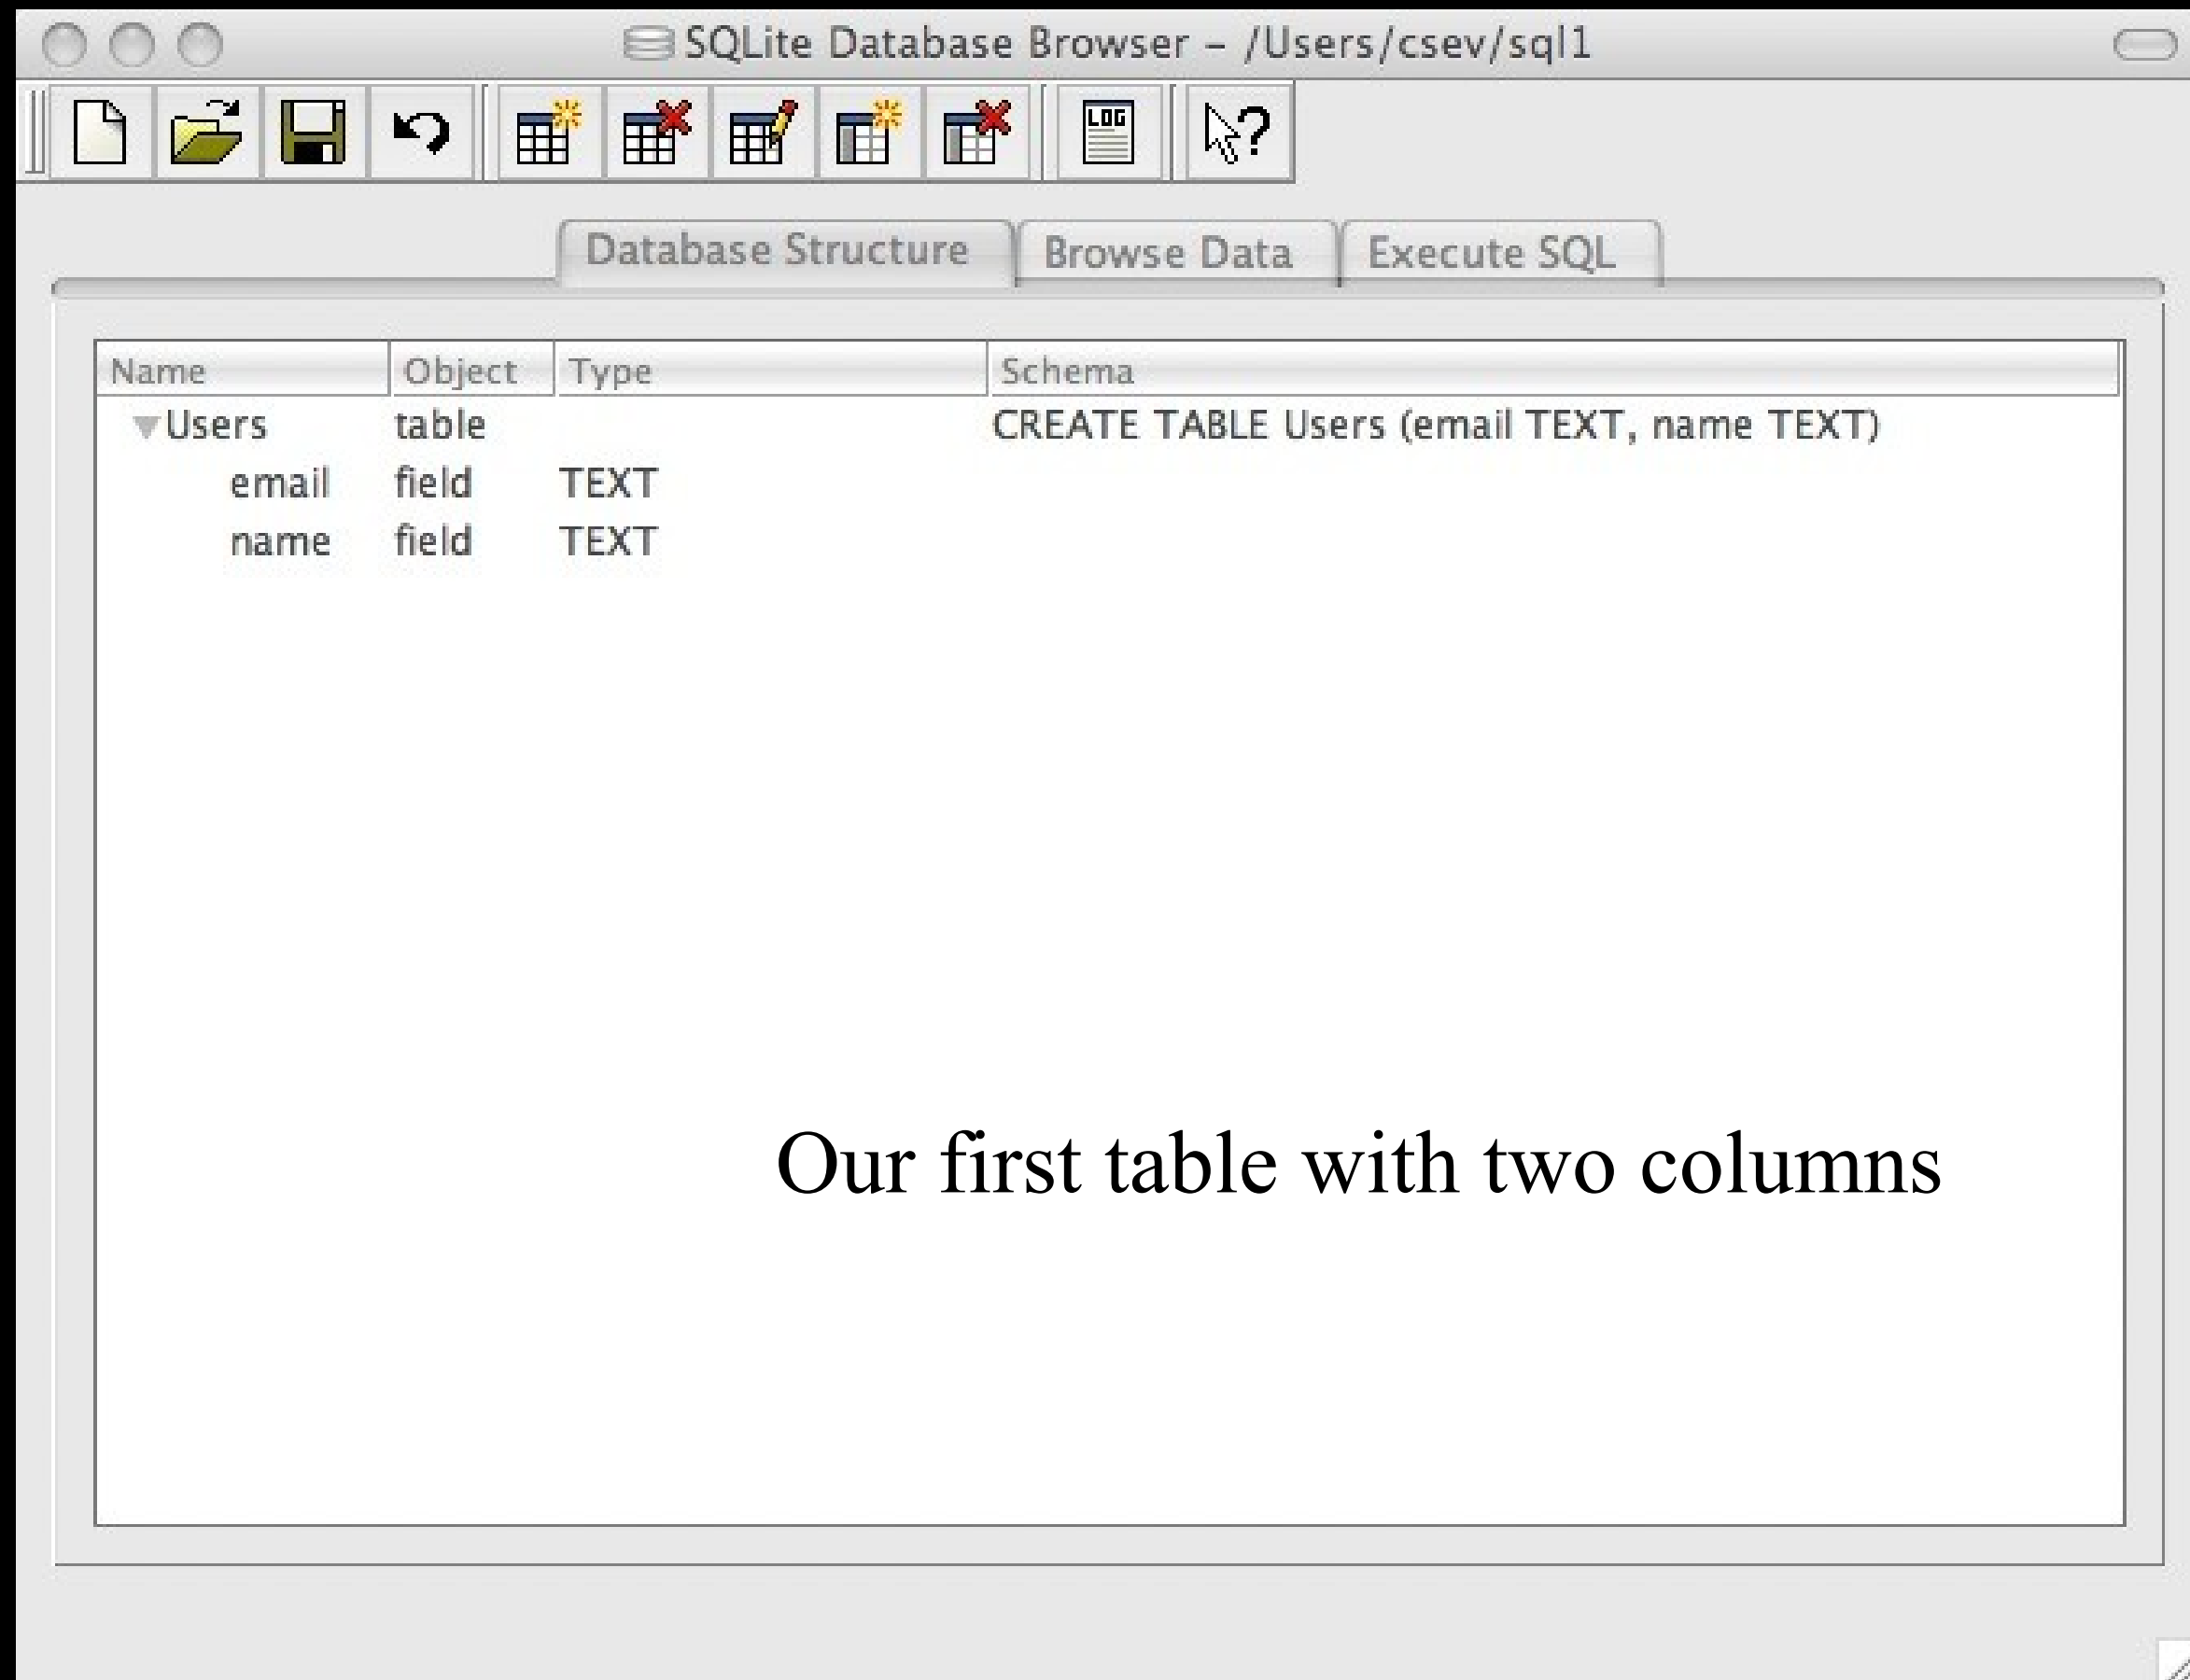
<http://www.sqlite.org/famous.html>



- <http://sqlitebrowser.sourceforge.net/>

Start Simple - A Single Table

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



Our first table with two columns

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

Table: Users

	email	name
1	csev@umich.edu	Chuck
2	botimer@umich.edu	Noah
3	mrzhou@umich.edu	Daniel
4	ksnam@umich.edu	Kevin

1 - 4 of 4 0

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')
```

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

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

Execute query

Error message from database engine:

No error

Data returned:

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

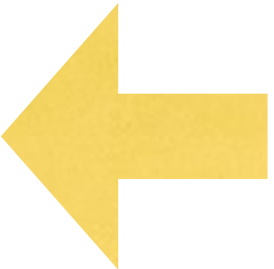
Table: Users

New Record Delete Record

	email	name
1	csev@umich.edu	Chuck
2	botimer@umich.edu	Noah
3	mrzhou@umich.edu	Daniel
4	ksnam@umich.edu	Kevin
5	ted@umich.edu	Ted

< 1 - 5 of 5 >

Go to: 0



SQL Delete

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

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

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

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

Execute query

Error message from database engine:

No error

Data returned:

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

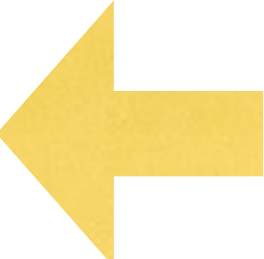
Table: Users

New Record Delete Record

	email	name
1	csev@umich.edu	Chuck
2	botimer@umich.edu	Noah
3	mrzhou@umich.edu	Daniel
4	ksnam@umich.edu	Kevin

1 - 4 of 4

Go to: 0



SQL: Update

- Allows the updating of a field with a where clause

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

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

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

Execute query

Error message from database engine:

No error

Data returned:

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL


Table: Users

New Record Delete Record

	email	name
1	csev@umich.edu	Charles
2	botimer@umich.edu	Noah
3	mrzhou@umich.edu	Daniel
4	ksnam@umich.edu	Kevin

1 - 4 of 4

Go to: 0



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'
```

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

```
select * from Users
```


Execute query

Error message from database engine:

No error

Data returned:

csev@umich.edu	Charles
botimer@umich.edu	Noah
mrzhou@umich.edu	Daniel
ksnam@umich.edu	Kevin



SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

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


Execute query

Error message from database engine:

No error

Data returned:

csev@umich.edu	Charles
----------------	---------



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
```

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

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

Execute query

Error message from database engine:

No error

Data returned:

csev@umich.edu	Charles
mrzhou@umich.edu	Daniel
ksnam@umich.edu	Kevin
botimer@umich.edu	Noah

SQLite Database Browser - /Users/csev/sql1

Database Structure Browse Data Execute SQL

SQL string:

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

Execute query

Error message from database engine:

No error

Data returned:

botimer@umich.edu	Noah
csev@umich.edu	Charles
ksnam@umich.edu	Kevin
mrzhou@umich.edu	Daniel

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)

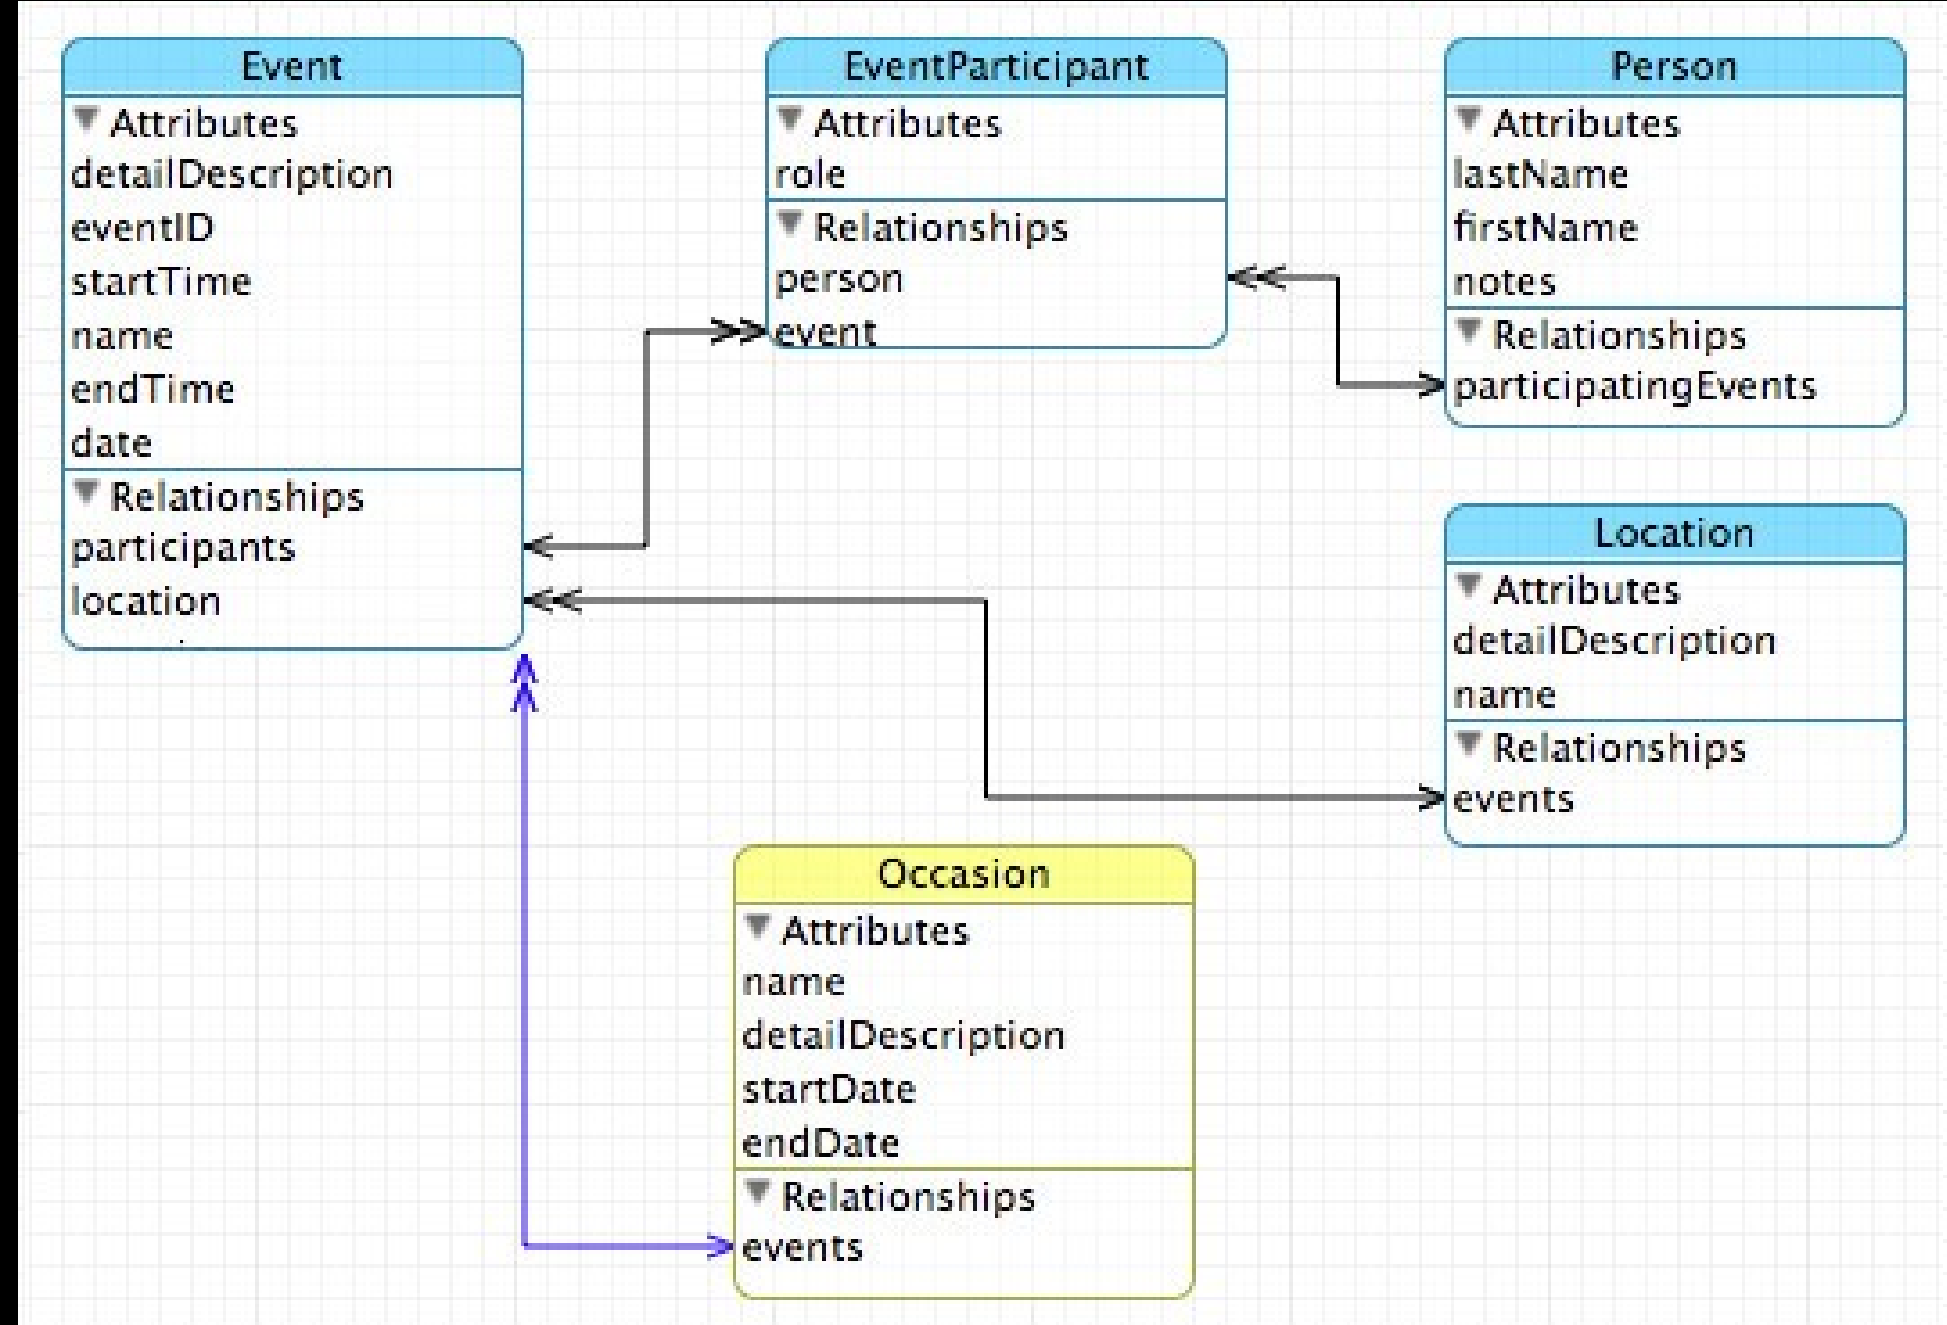
- Tables pretty much look like big fast programmable spreadsheet with rows, columns, and commands
- The power comes when we have more than one table and we can exploit the relationships between the tables

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...



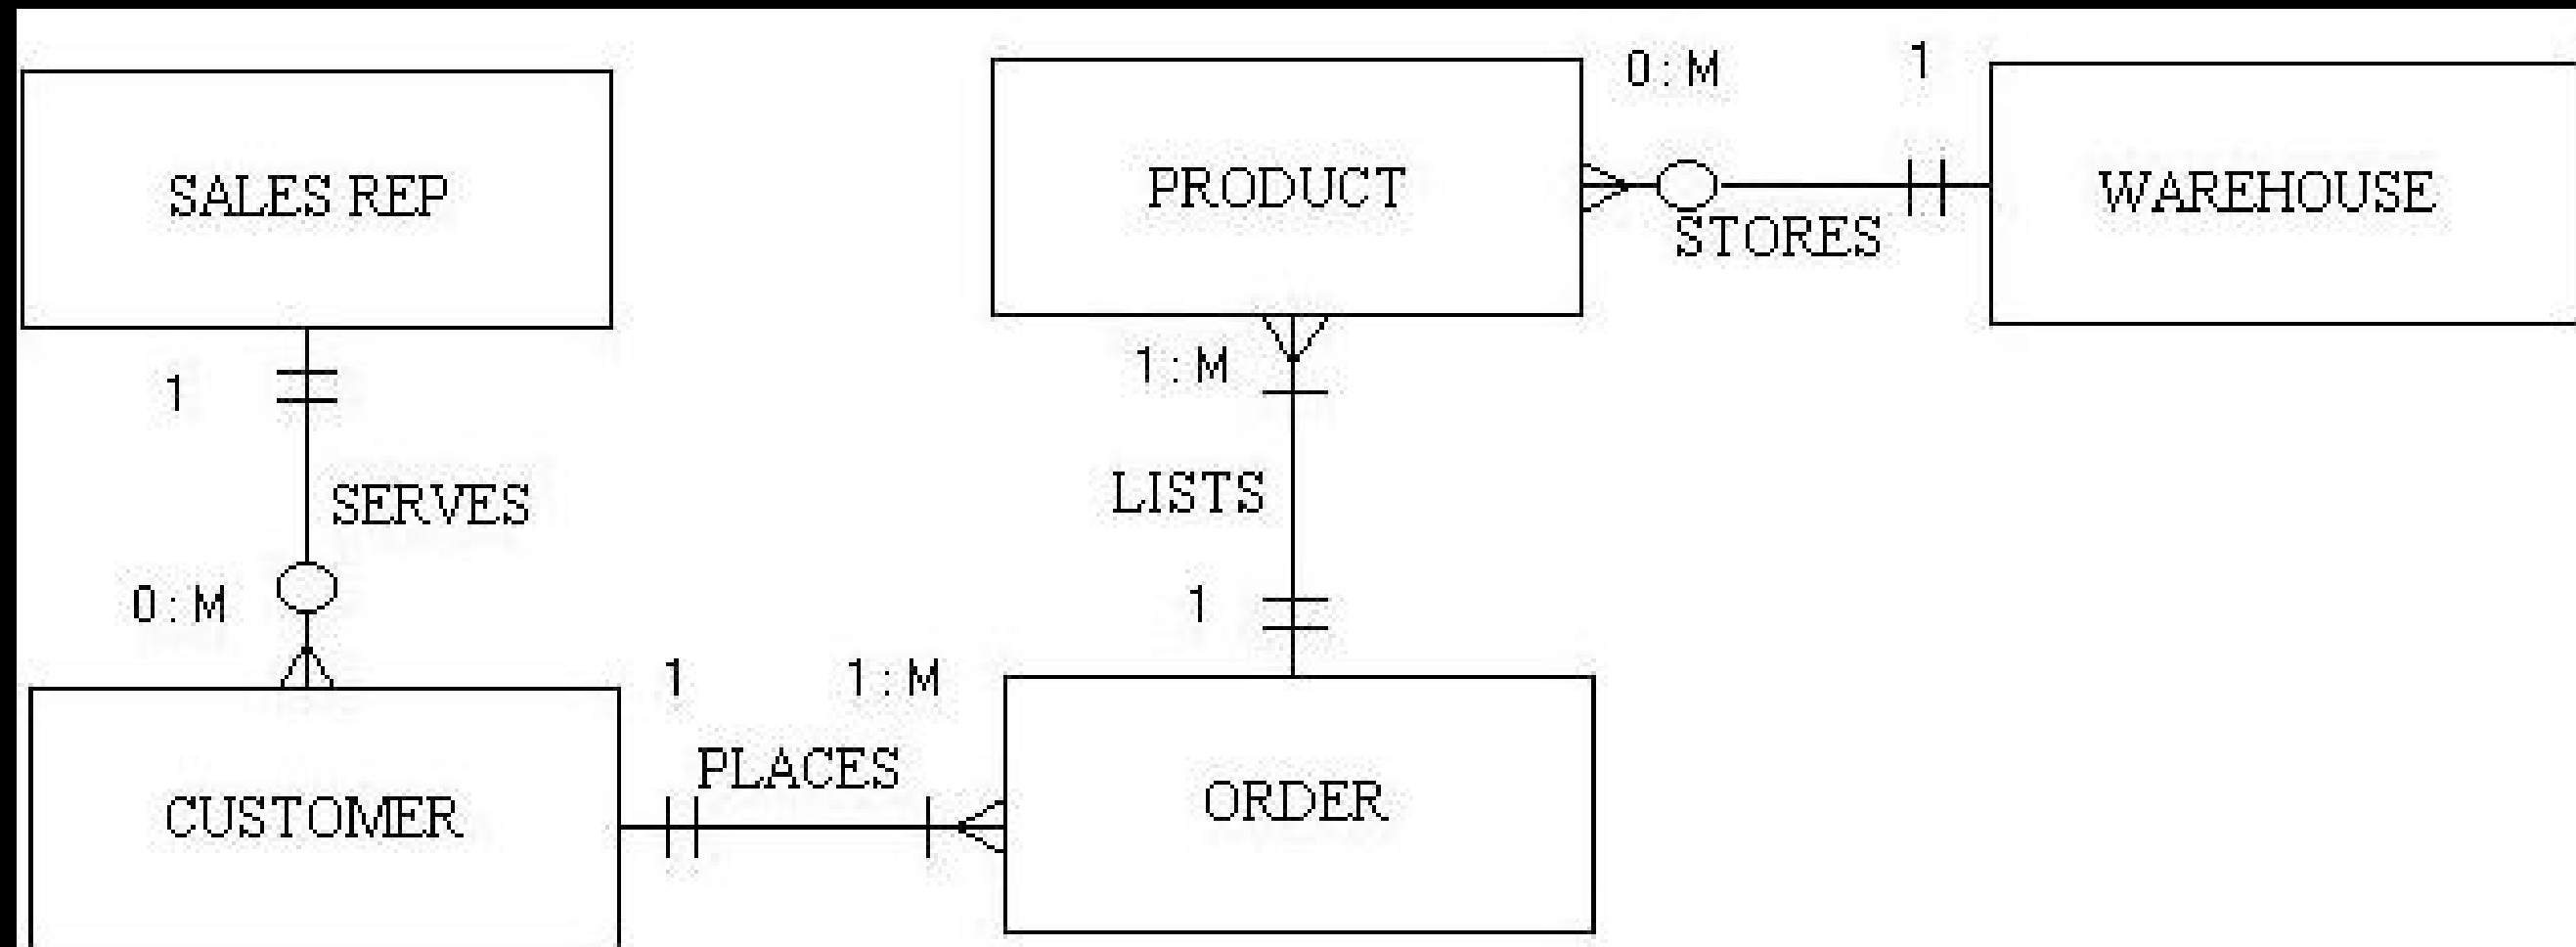
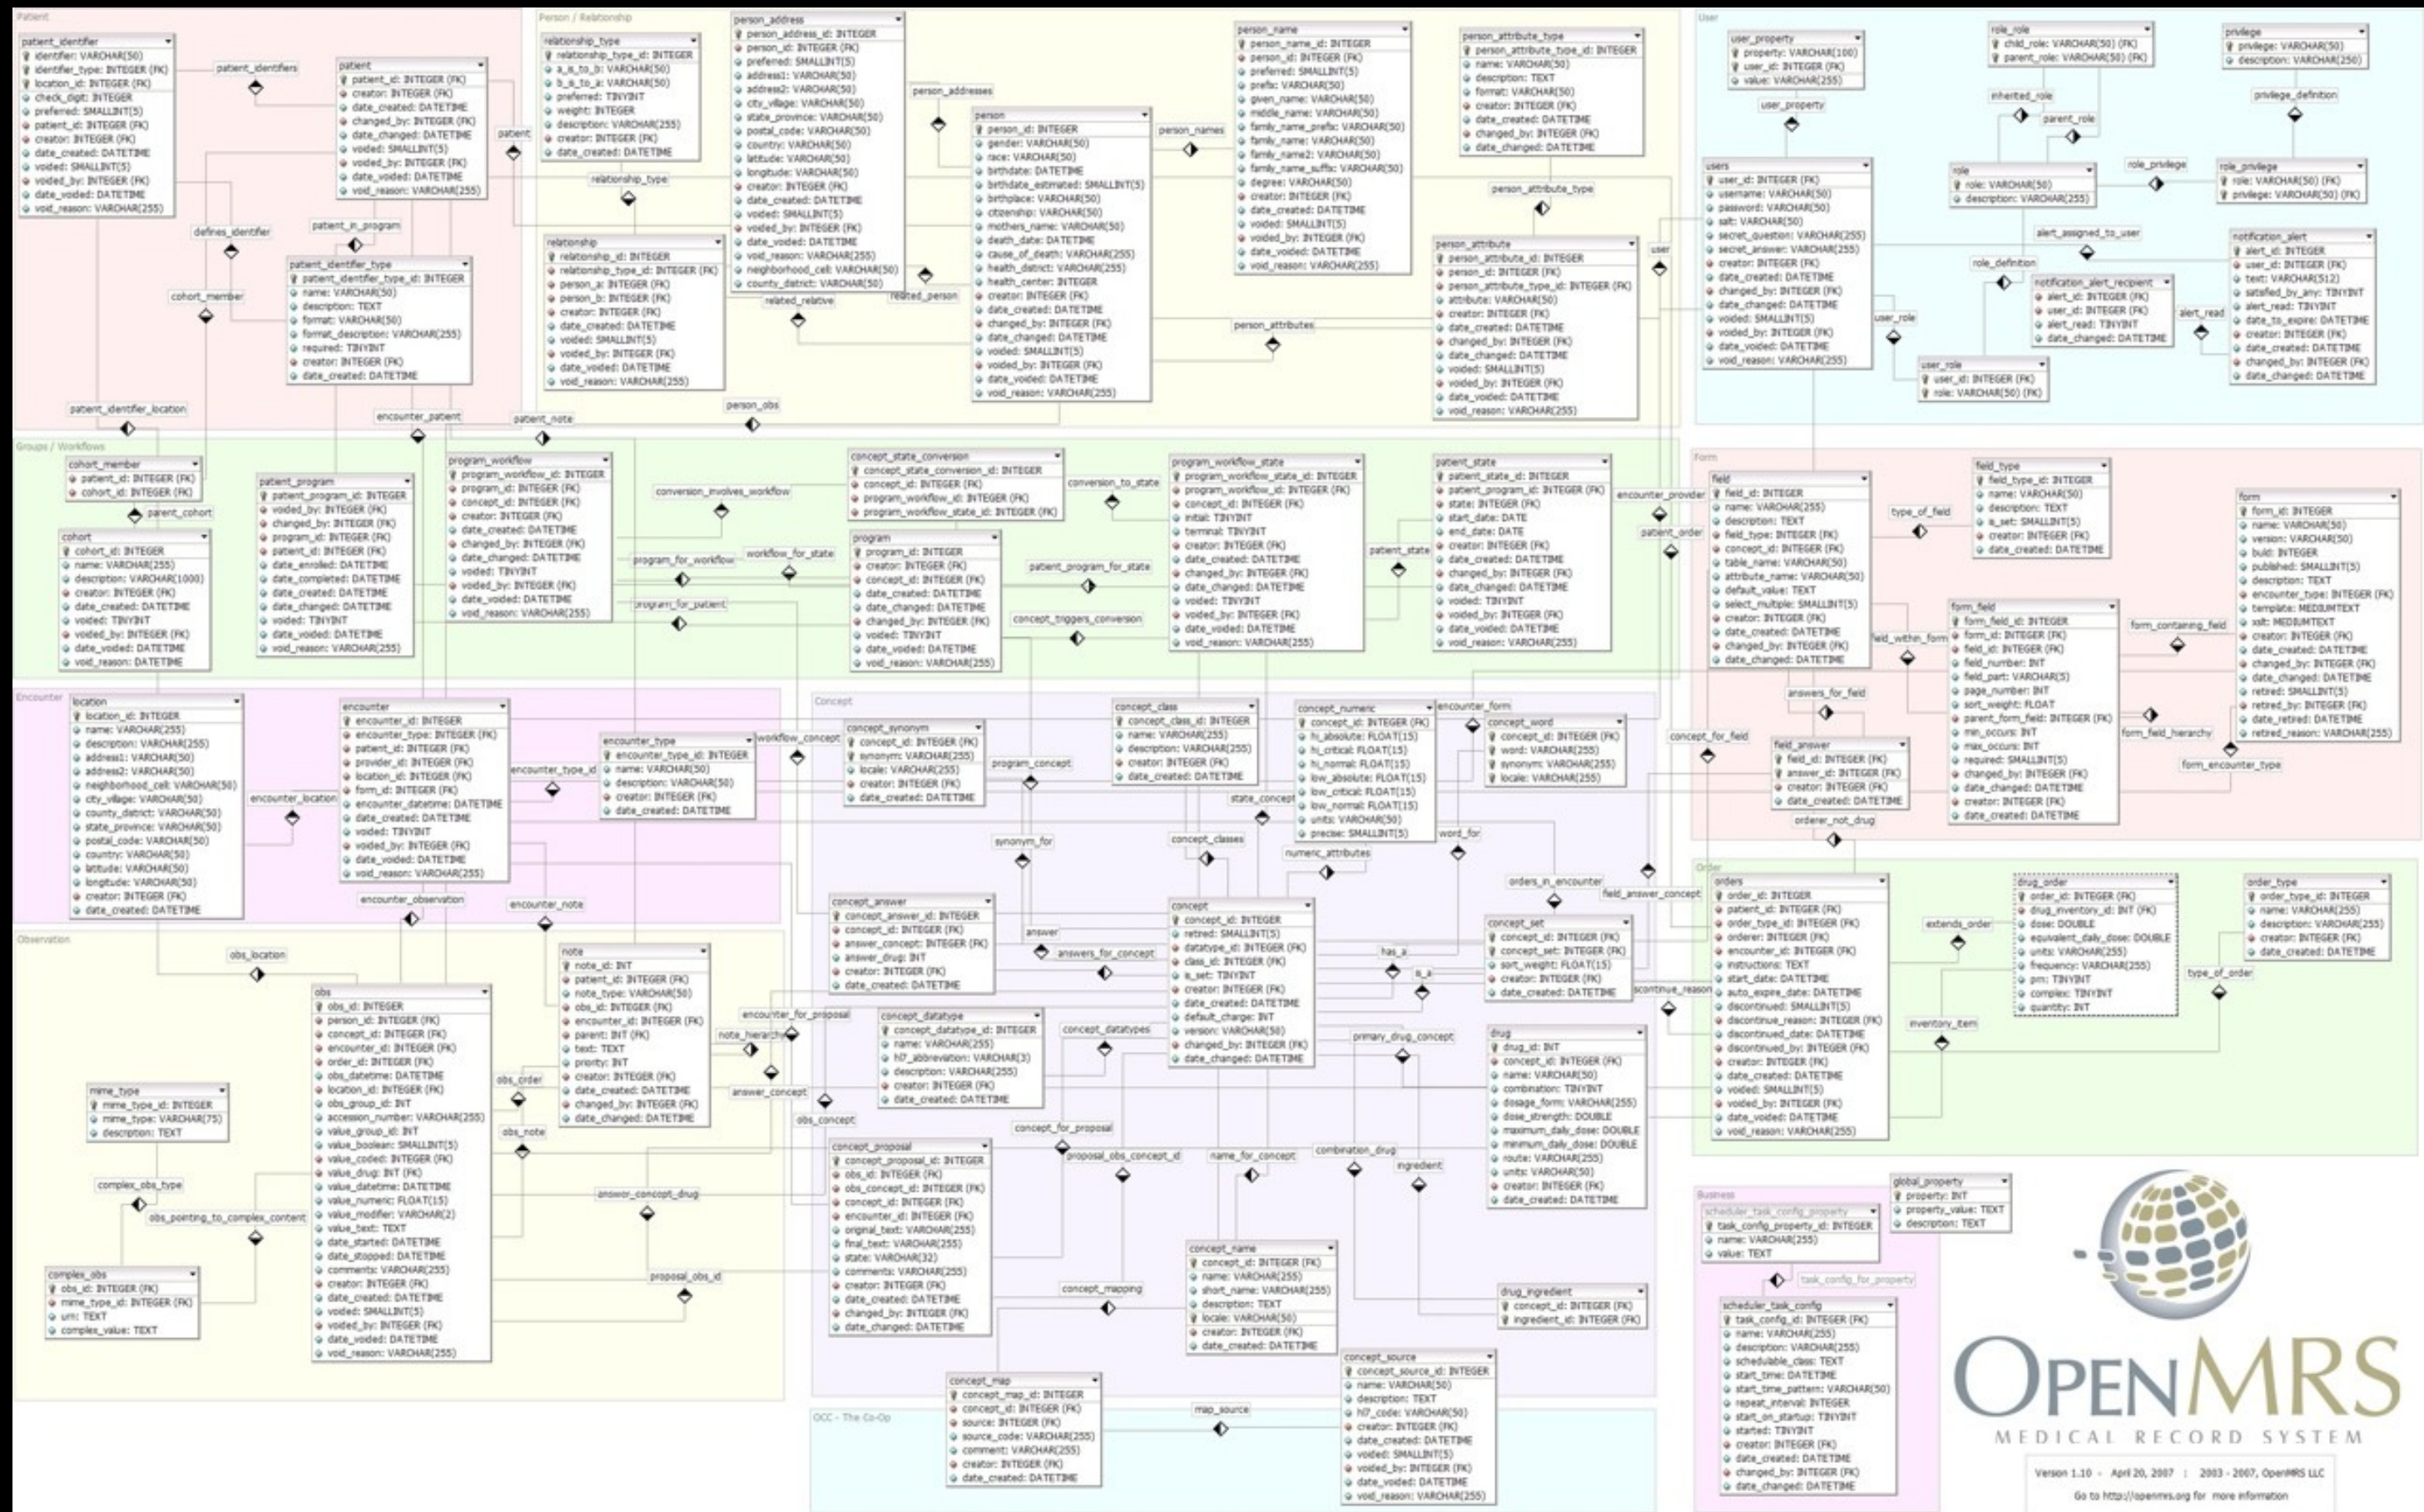


Figure 1. Entity-Relationship Diagram

- * 1 INSTANCE OF A SALES REP SERVES 1 TO MANY CUSTOMERS
- * 1 INSTANCE OF A CUSTOMER PLACES 1 TO MANY ORDERS
- * 1 INSTANCE OF AN ORDER LISTS 1 TO MANY PRODUCTS
- * 1 INSTANCE OF A WAREHOUSE STORES 0 TO MANY PRODUCTS



OPENMRS

MEDICAL RECORD SYSTEM

Version 1.10 · April 20, 2007 · 2003 - 2007, OpenMRS LLC
 Go to <http://openmrs.org> for more information

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.

Len Album

Genre

Artist

Rating

Track

Count

Source: Apple iTunes Terminal

<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"/> Tie Man	3:20	America	Greatest Hits	Easy Listen...	★★★★★	22

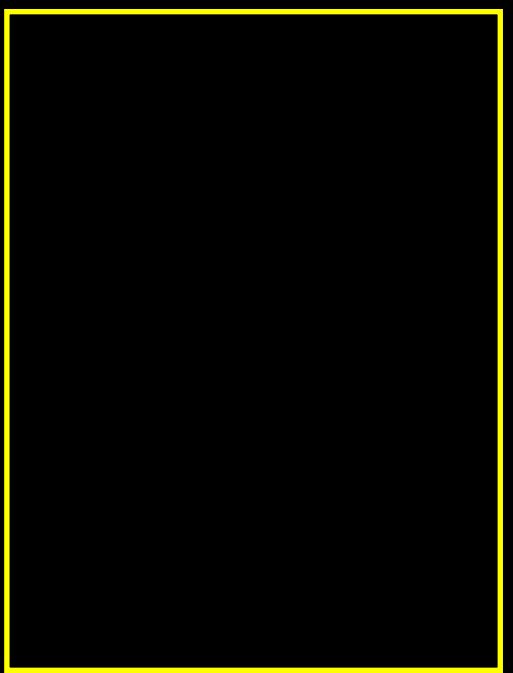
Track
 Artist
 Album
 Genre
 Rating
 Len
 Count



belongs-to



belongs-to

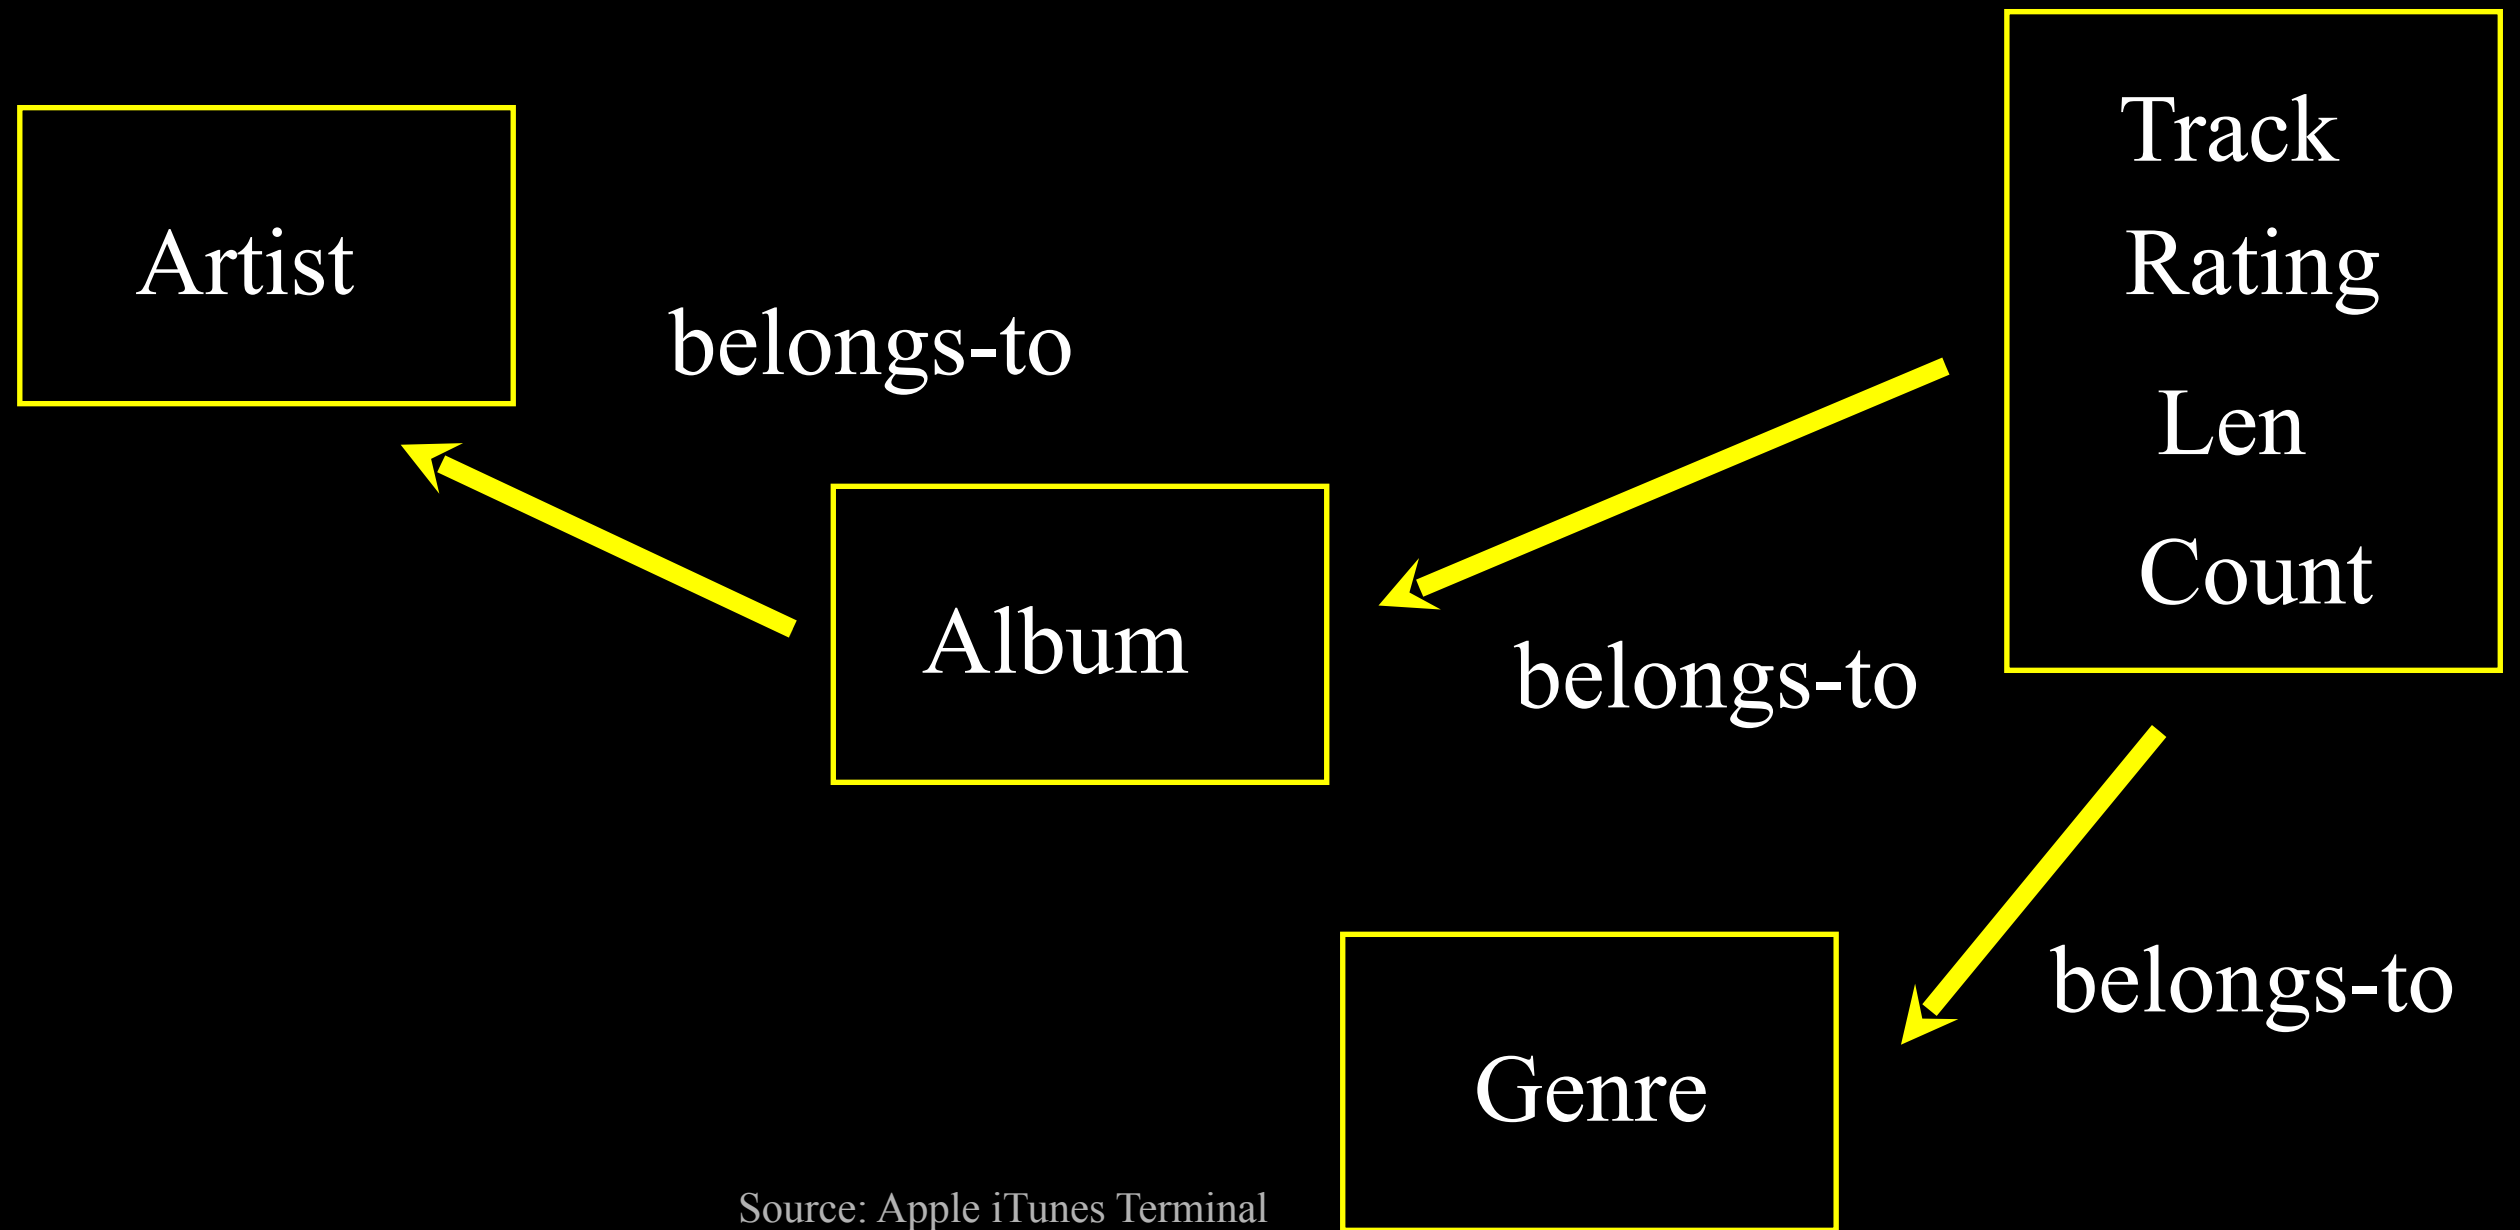


belongs-to



Source: Apple iTunes Terminal

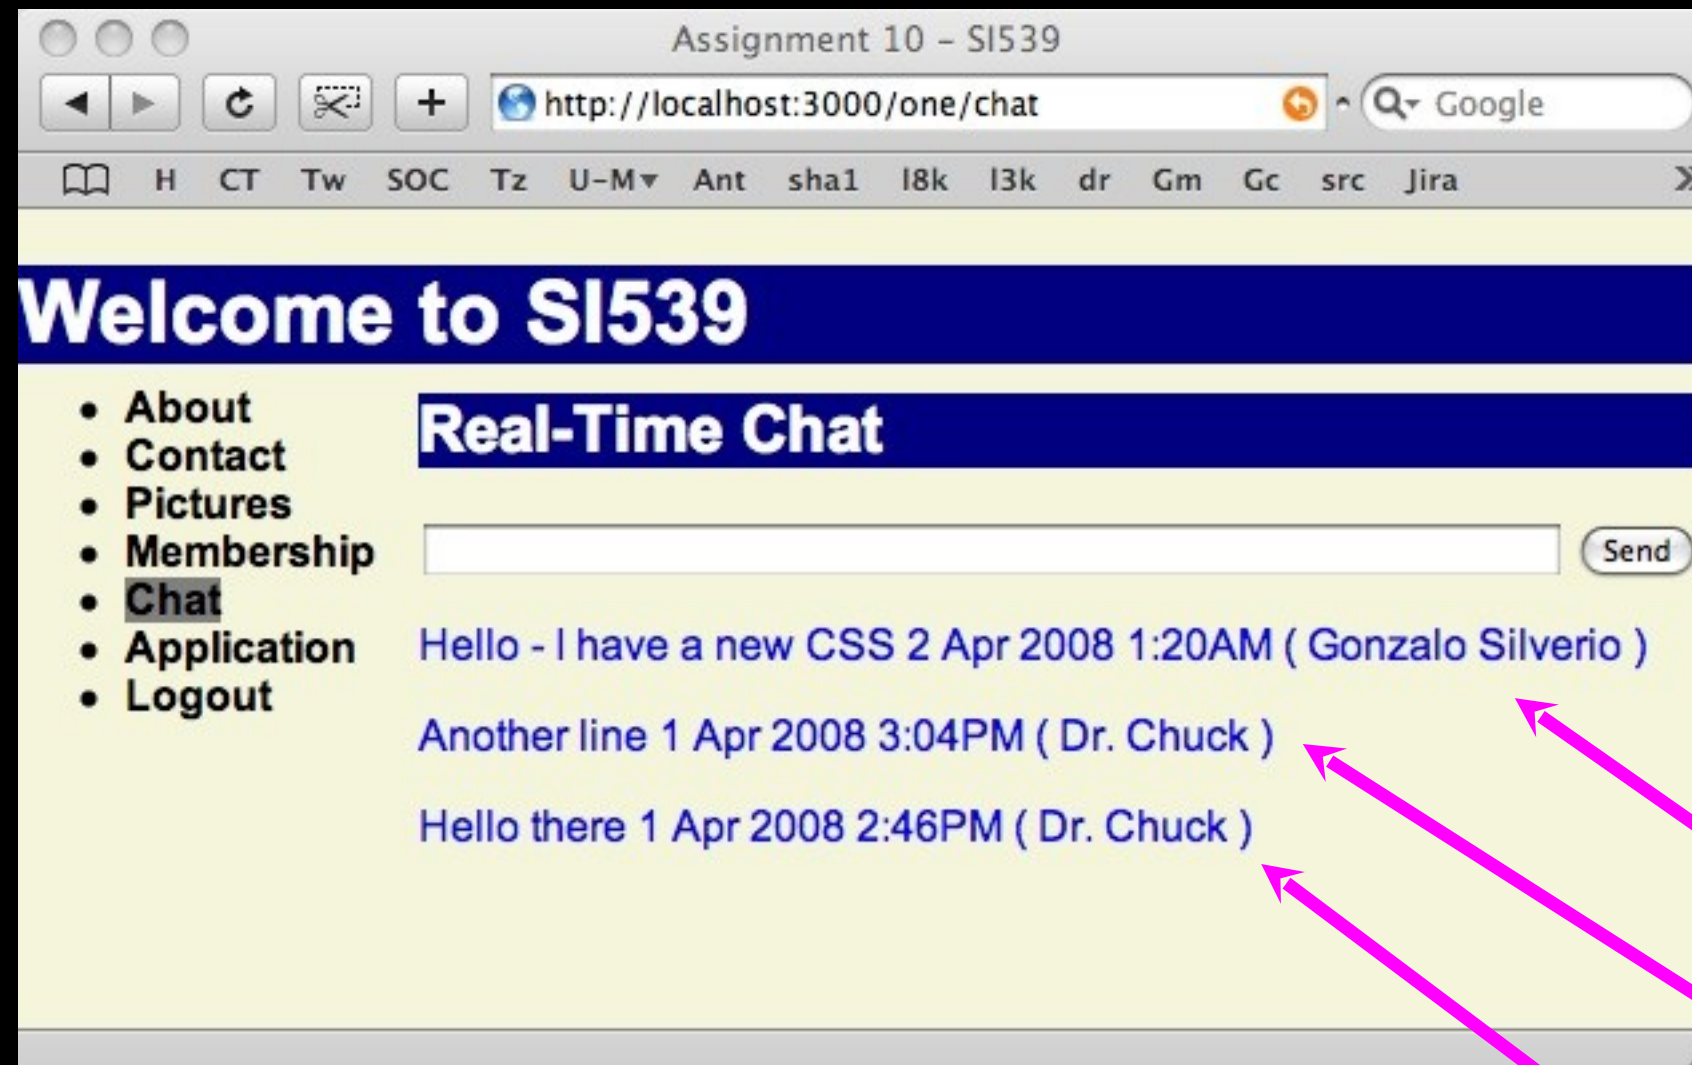
<input checked="" type="checkbox"/>	Hells Bells	5:13	AC/DC	Who Made Who	Rock	★★★★★	61
<input checked="" type="checkbox"/>	Shake Your Foundations	3:54	AC/DC	Who Made Who	Rock	★★★★★	70
<input checked="" type="checkbox"/>	Chase the Ace	3:01	AC/DC	Who Made Who	Rock		56
<input checked="" type="checkbox"/>	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"/>	Tie Man	3:20	America	Greatest Hits	Easy Listen...	★★★★★	22



Source: Apple iTunes Terminal

<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"/>	Tie Man	3:20	America	Greatest Hits	Easy Listen...	★★★★★	22

Representing Relationships in a Database



We want to keep track of who is the “owner” of each chat message...

Who does this chat message “belong 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” to each table which we will reference - by convention many programmers call this “id”

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

Better Reference Pattern

We use integers to reference rows in another table.

Table: members

	id	name	email
1	1	Dr. Chuck	csev@umich.edu
2	2	Gonzalo Silverio	gsilver@umich.edu

Table: chats

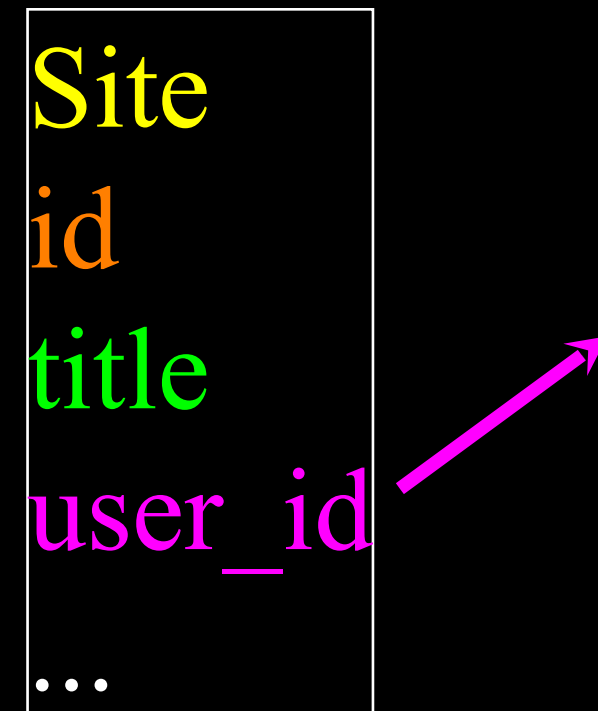
	id	chatmsg	member_id	created_at
1	1	Hello there	1	2008-04-01 14
2	2	Another line	1	2008-04-01 15
3	3	Hello - I have a	2	2008-04-02 01

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



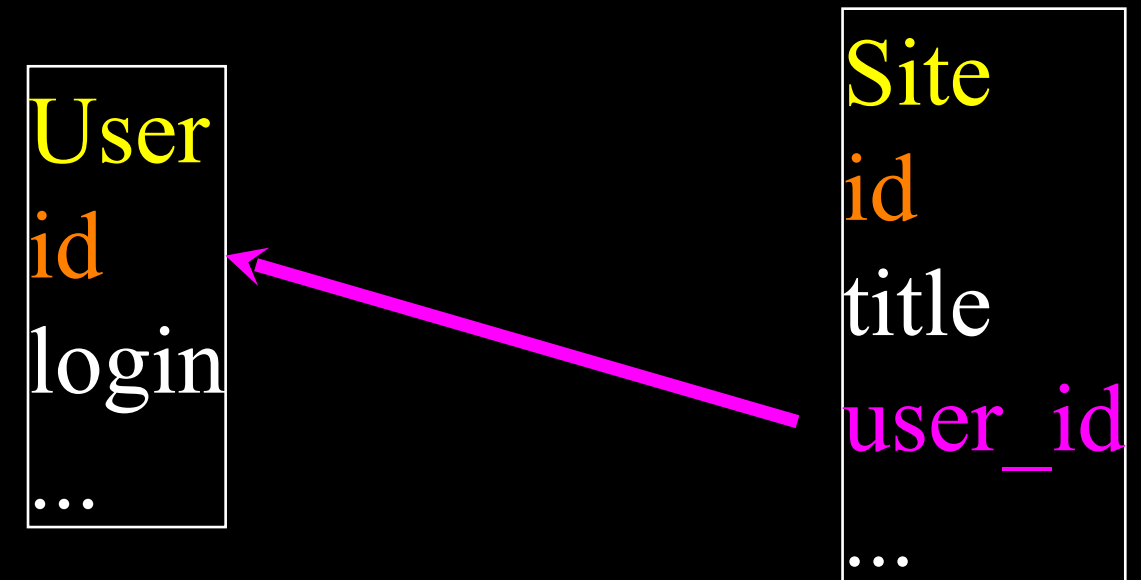
Primary Key Rules

- Rails encourages you to follow 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

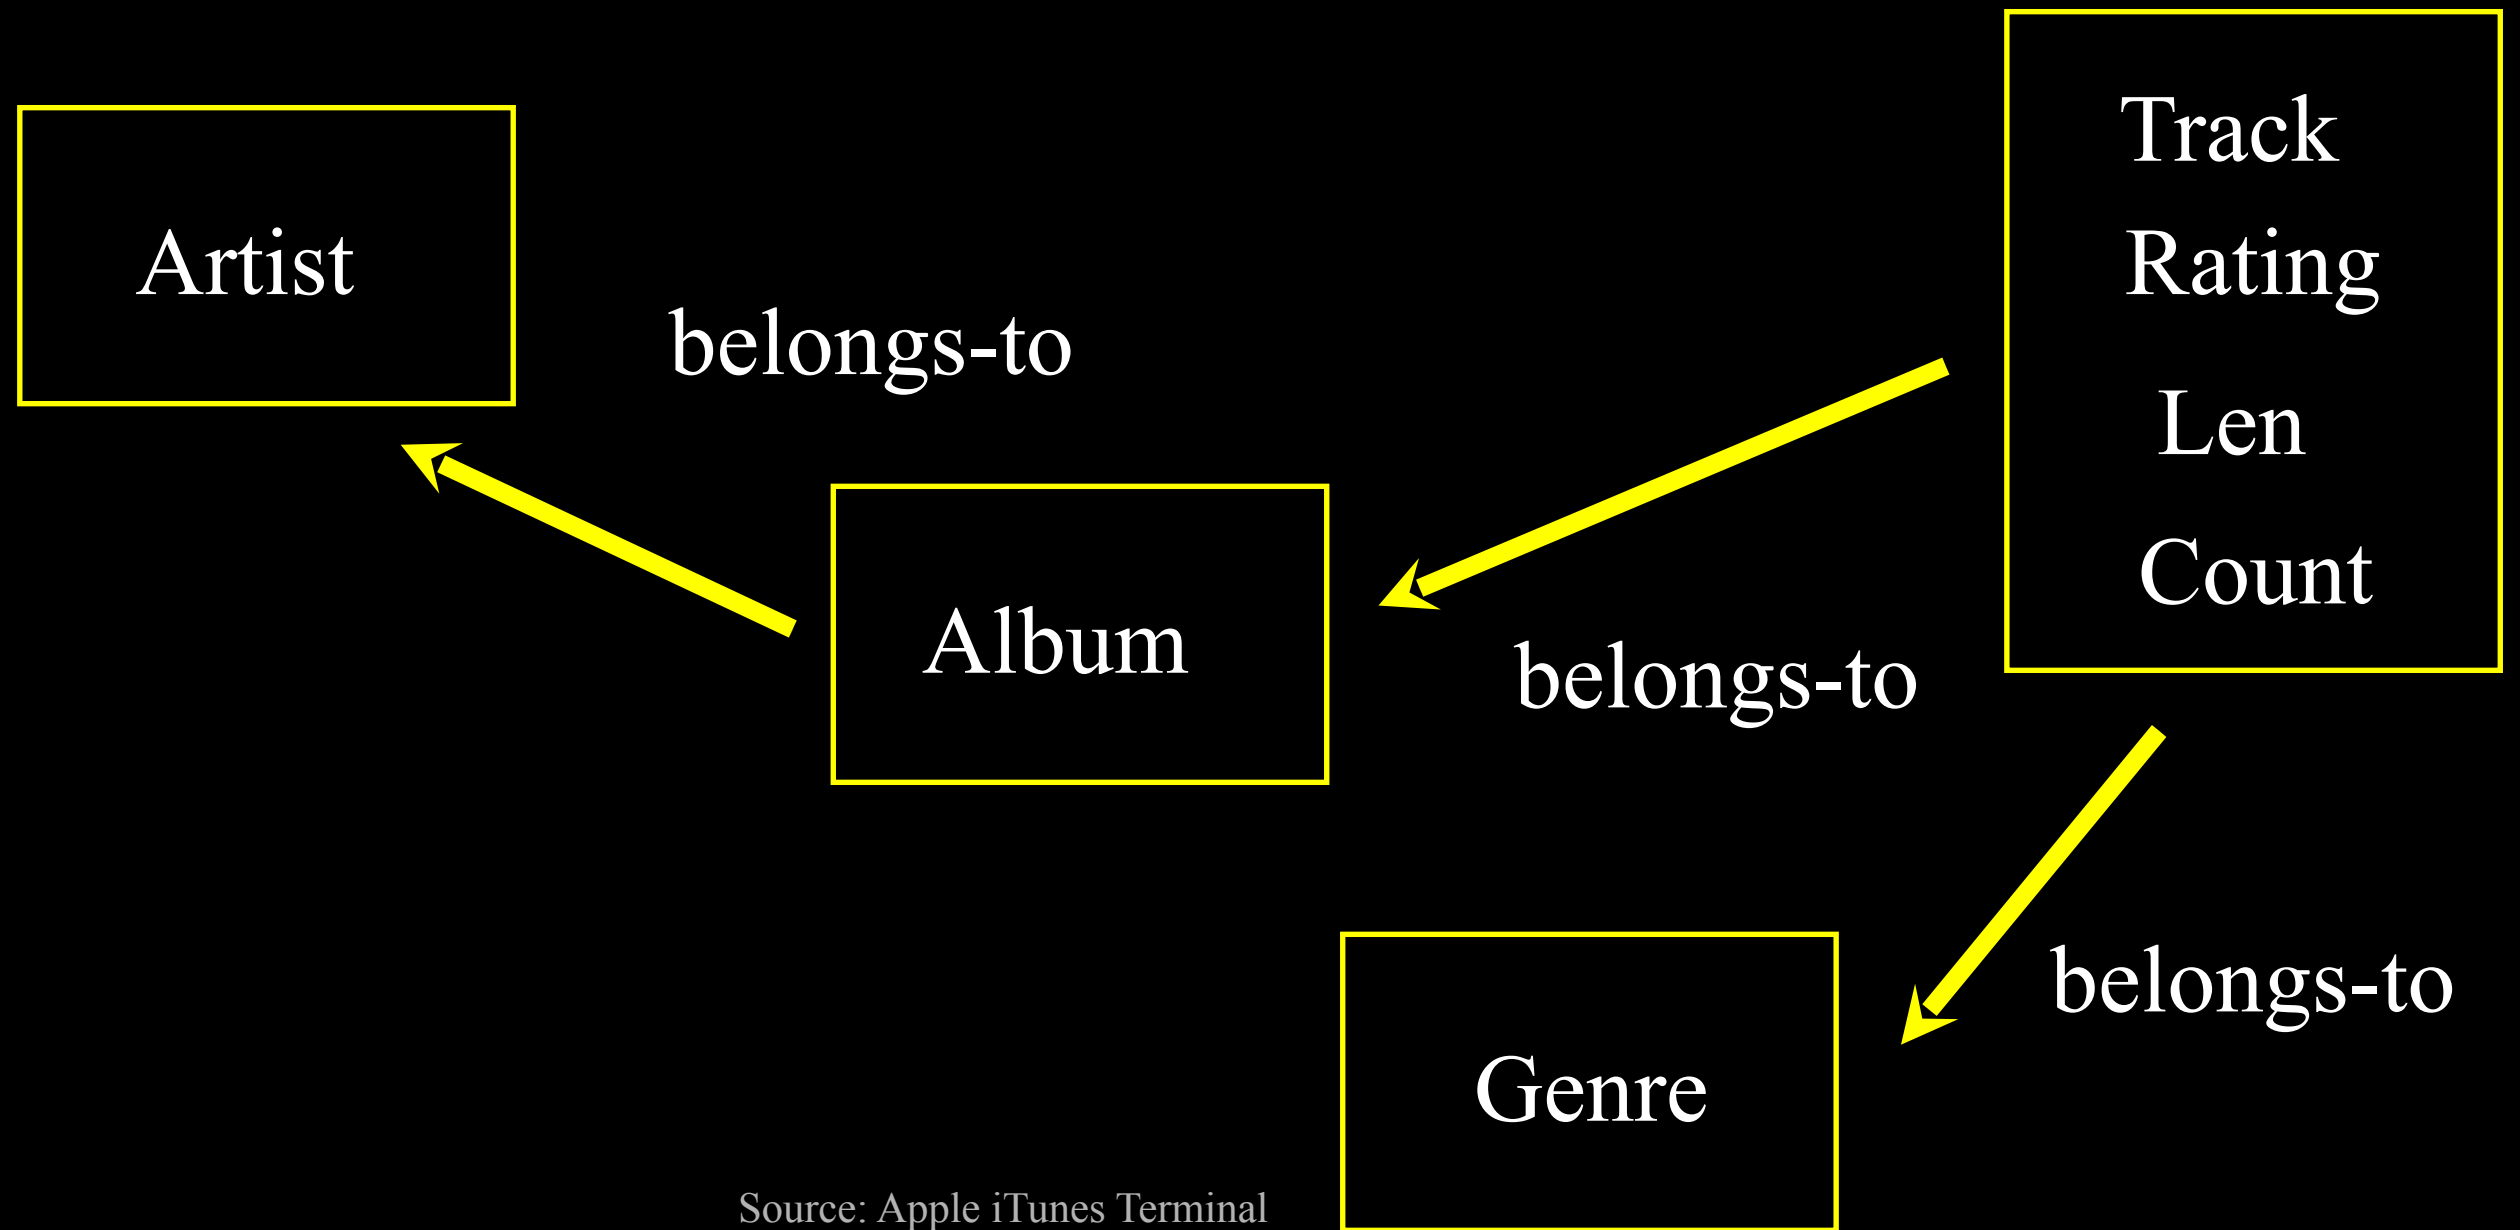
```
User
id
login
password
name
email
created_at
modified_at
login_at
```

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)



Source: Apple iTunes Terminal

<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"/>	Tie Man	3:20	America	Greatest Hits	Easy Listen...	★★★★★	22

Album

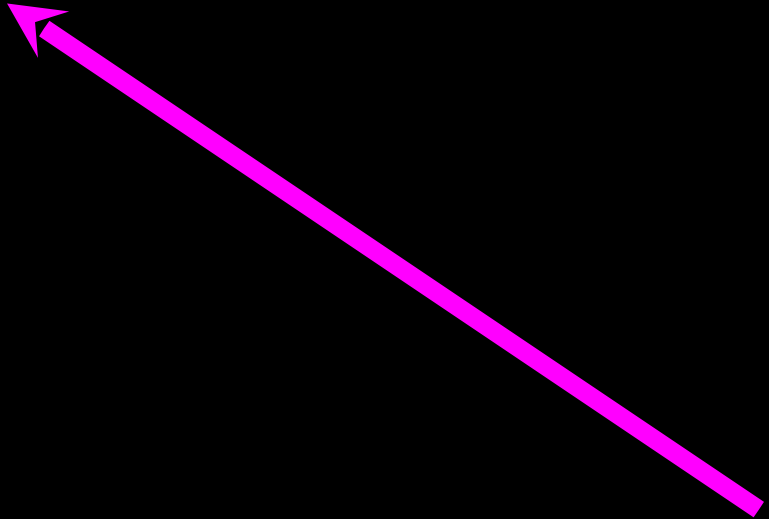
belongs-to



Track
Rating
Len
Count

Table
Primary key
Logical key
Foreign key

Album
id
title



Track
id
title
rating
len
count
album_id

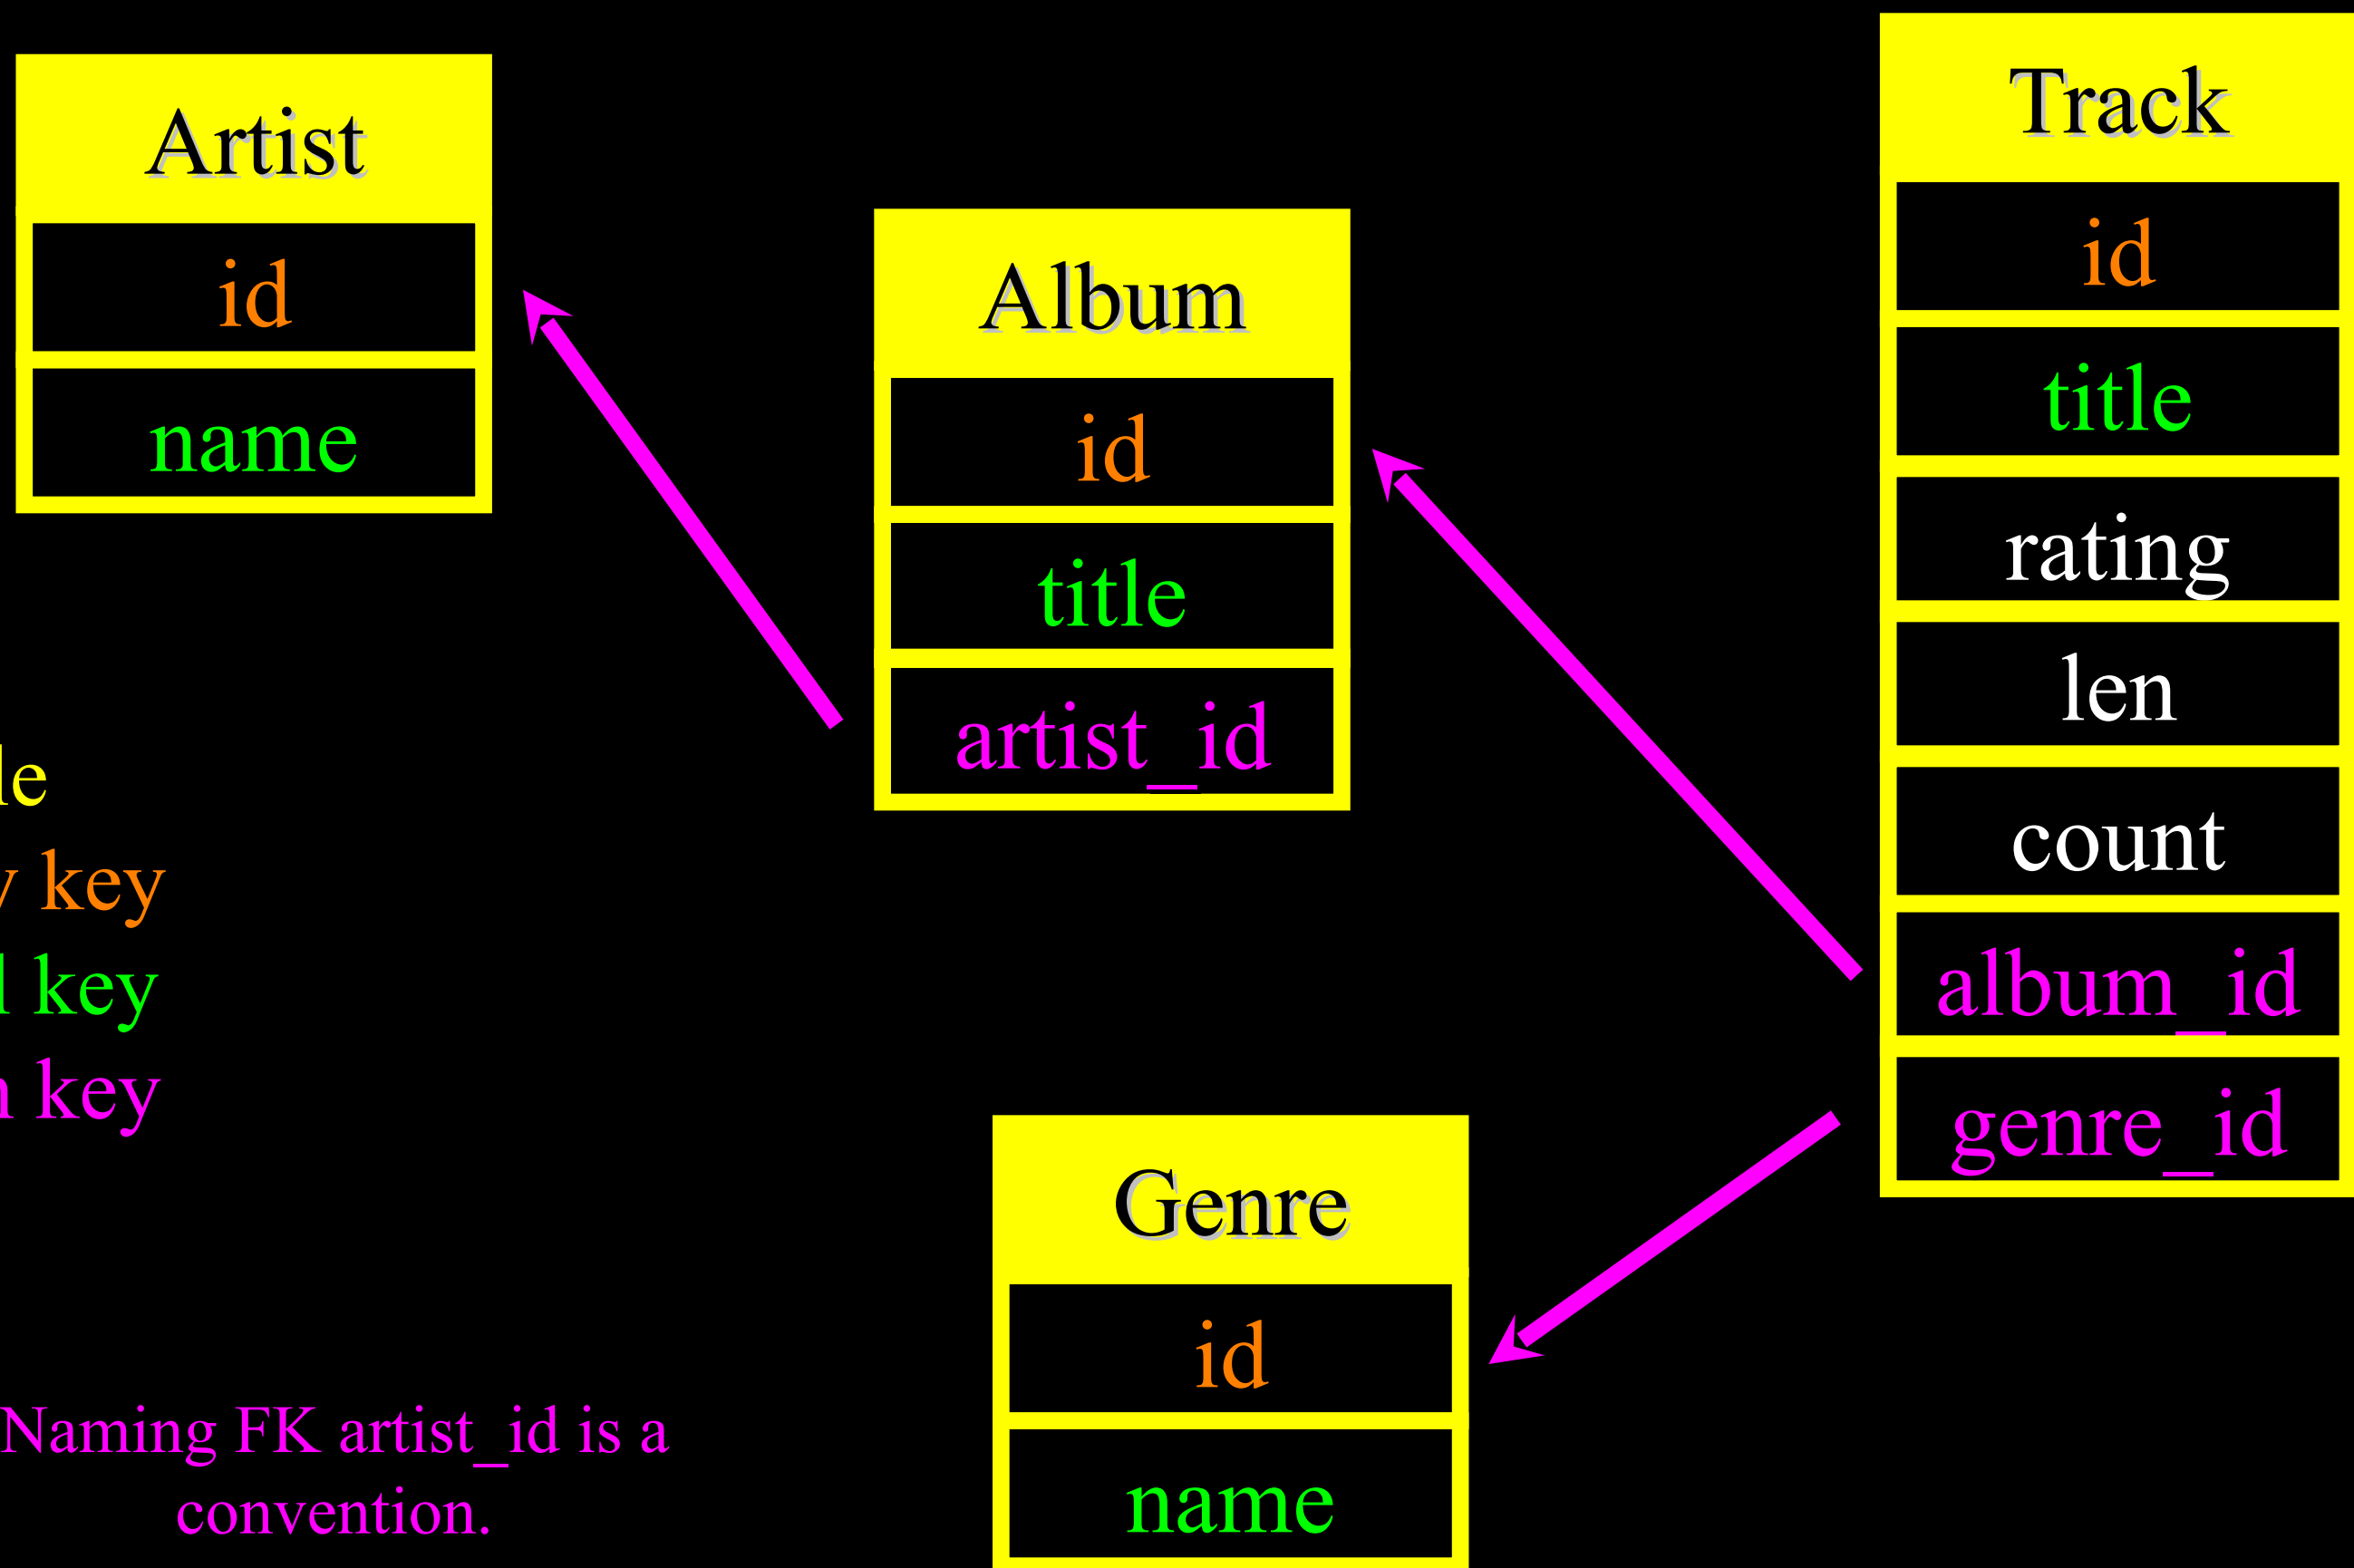
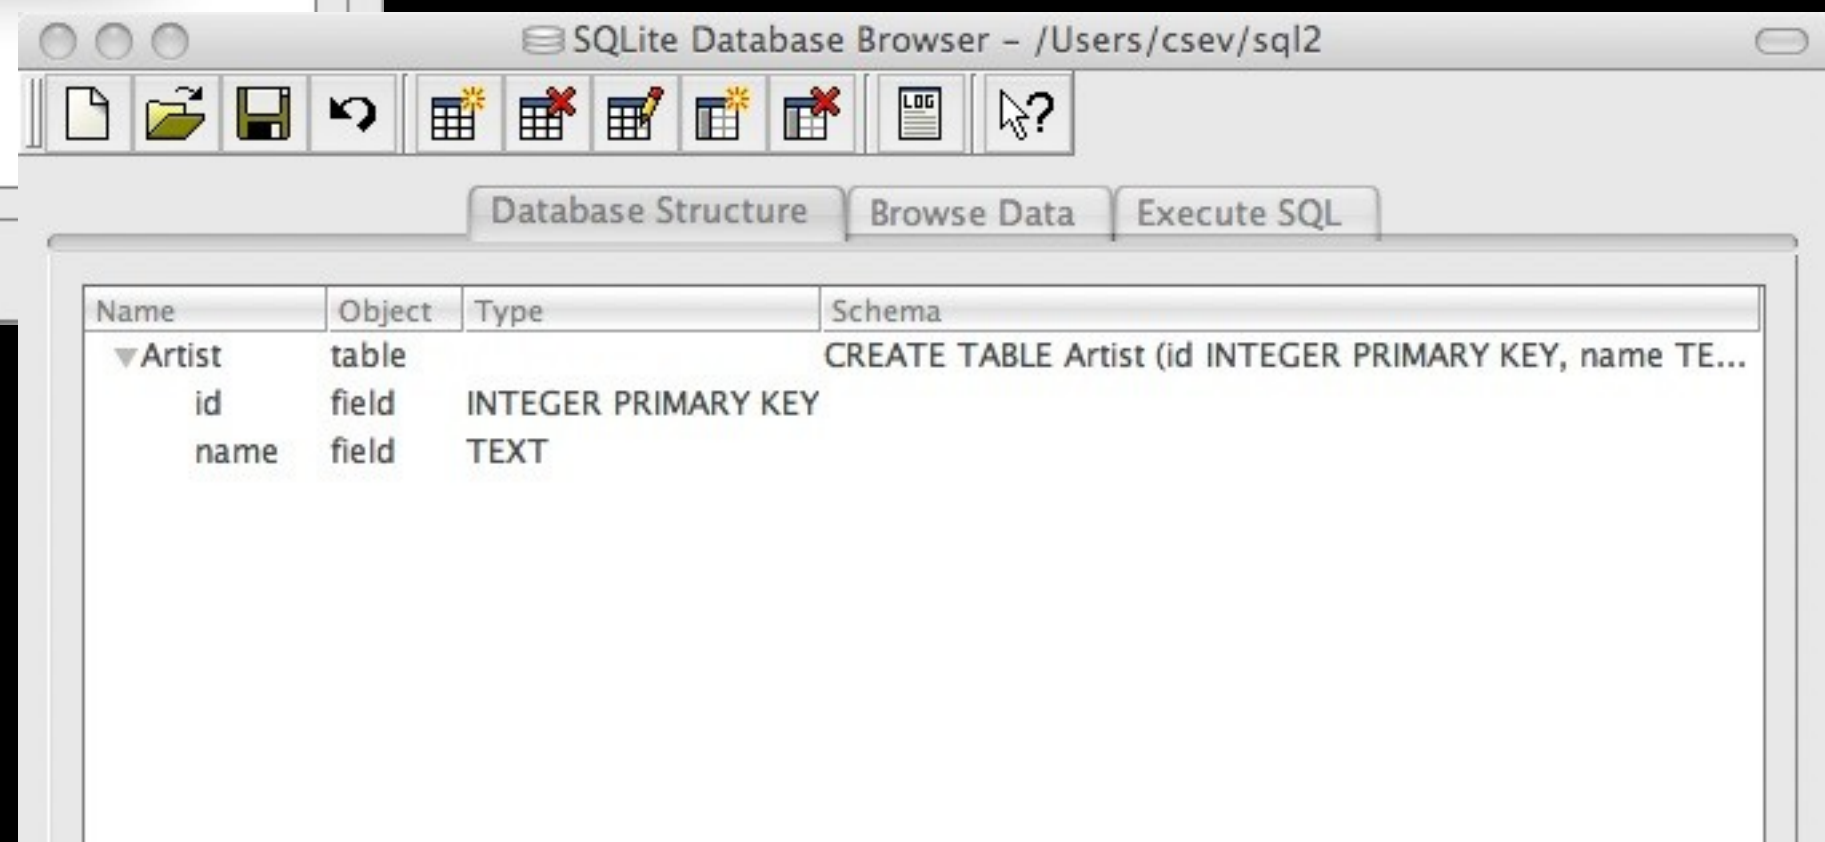
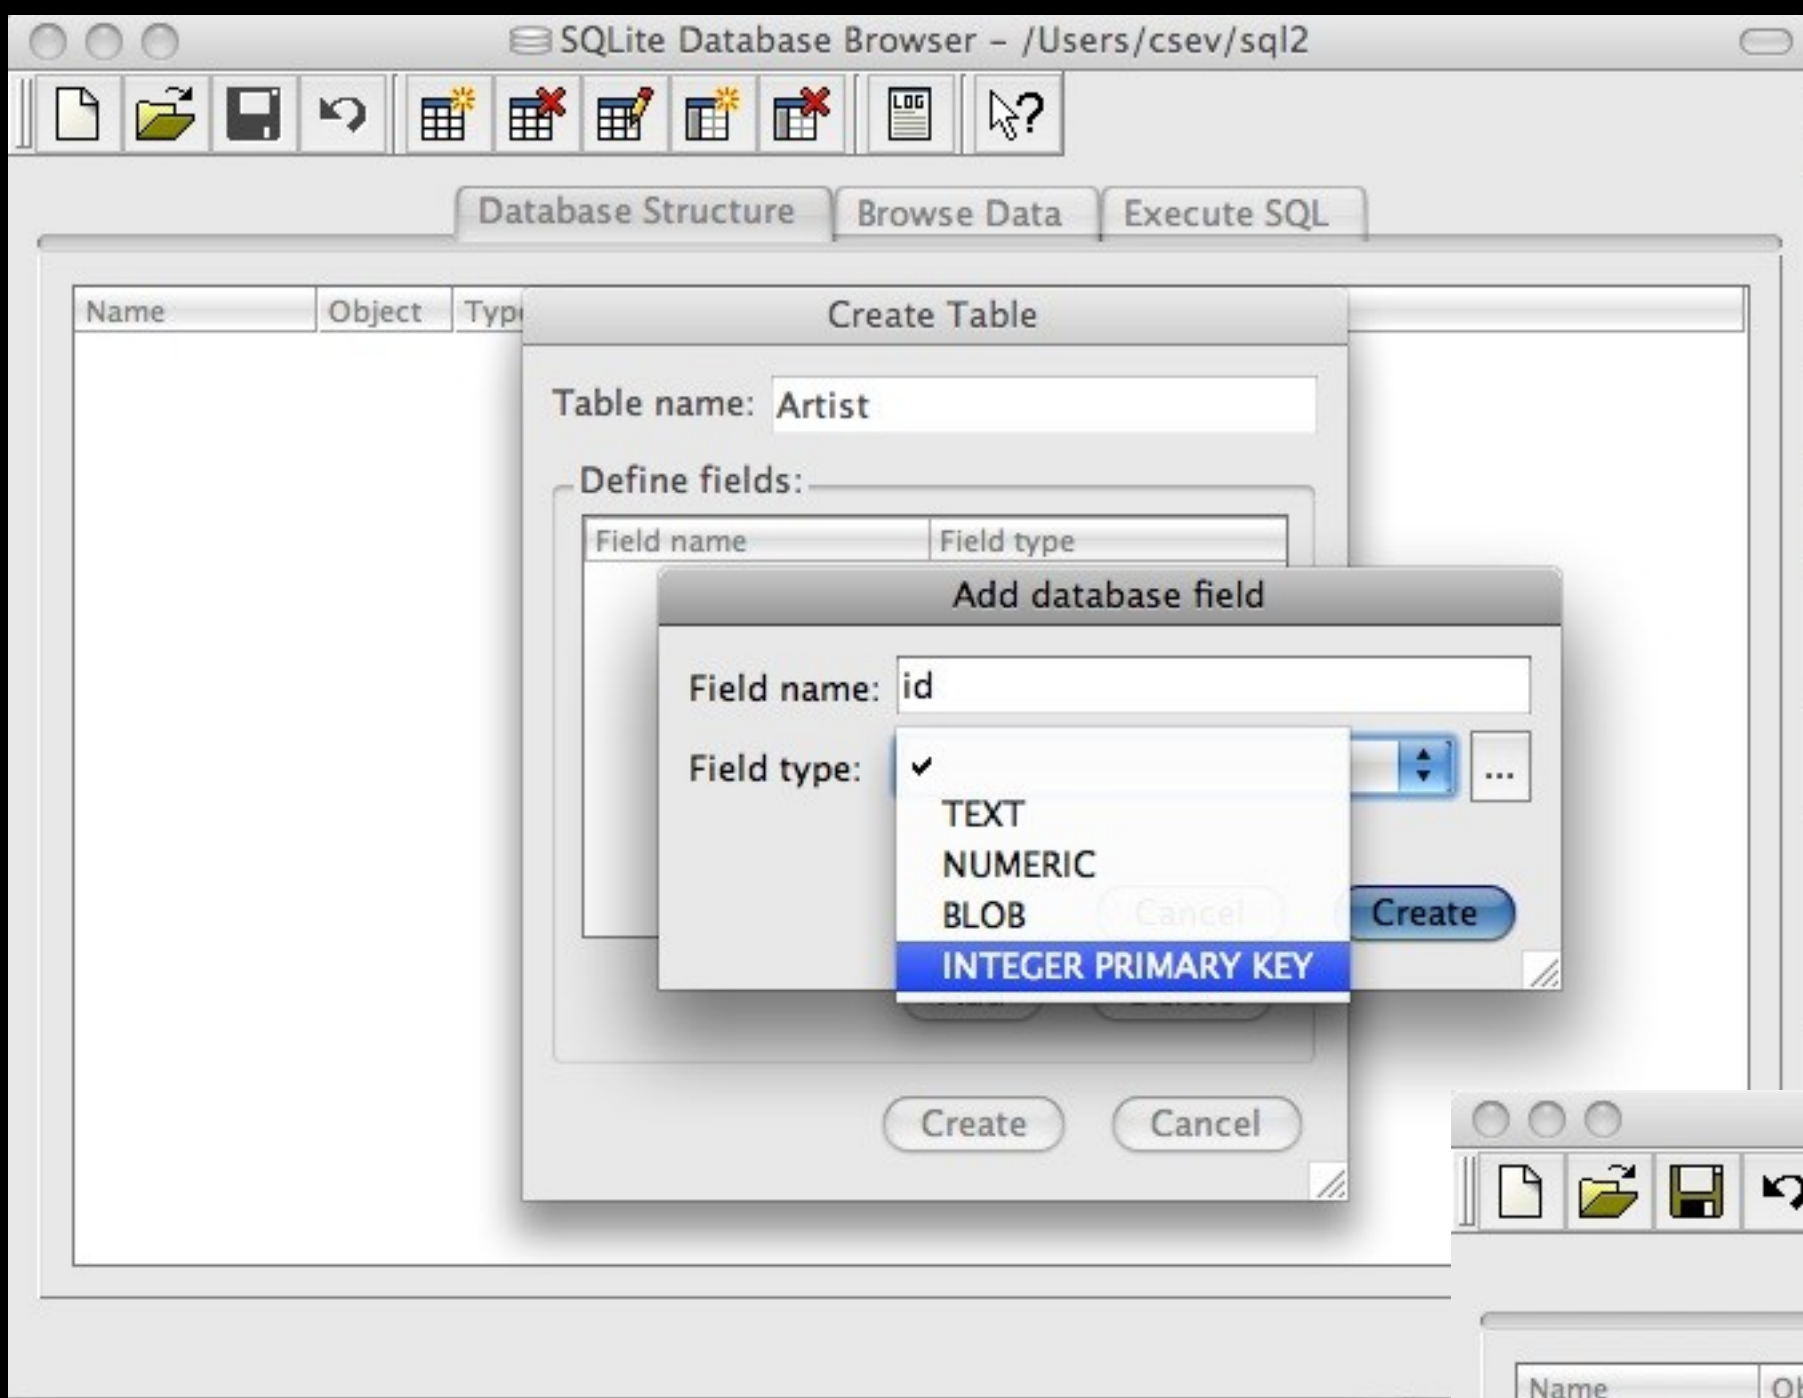


Table
Primary key
Logical key
Foreign key

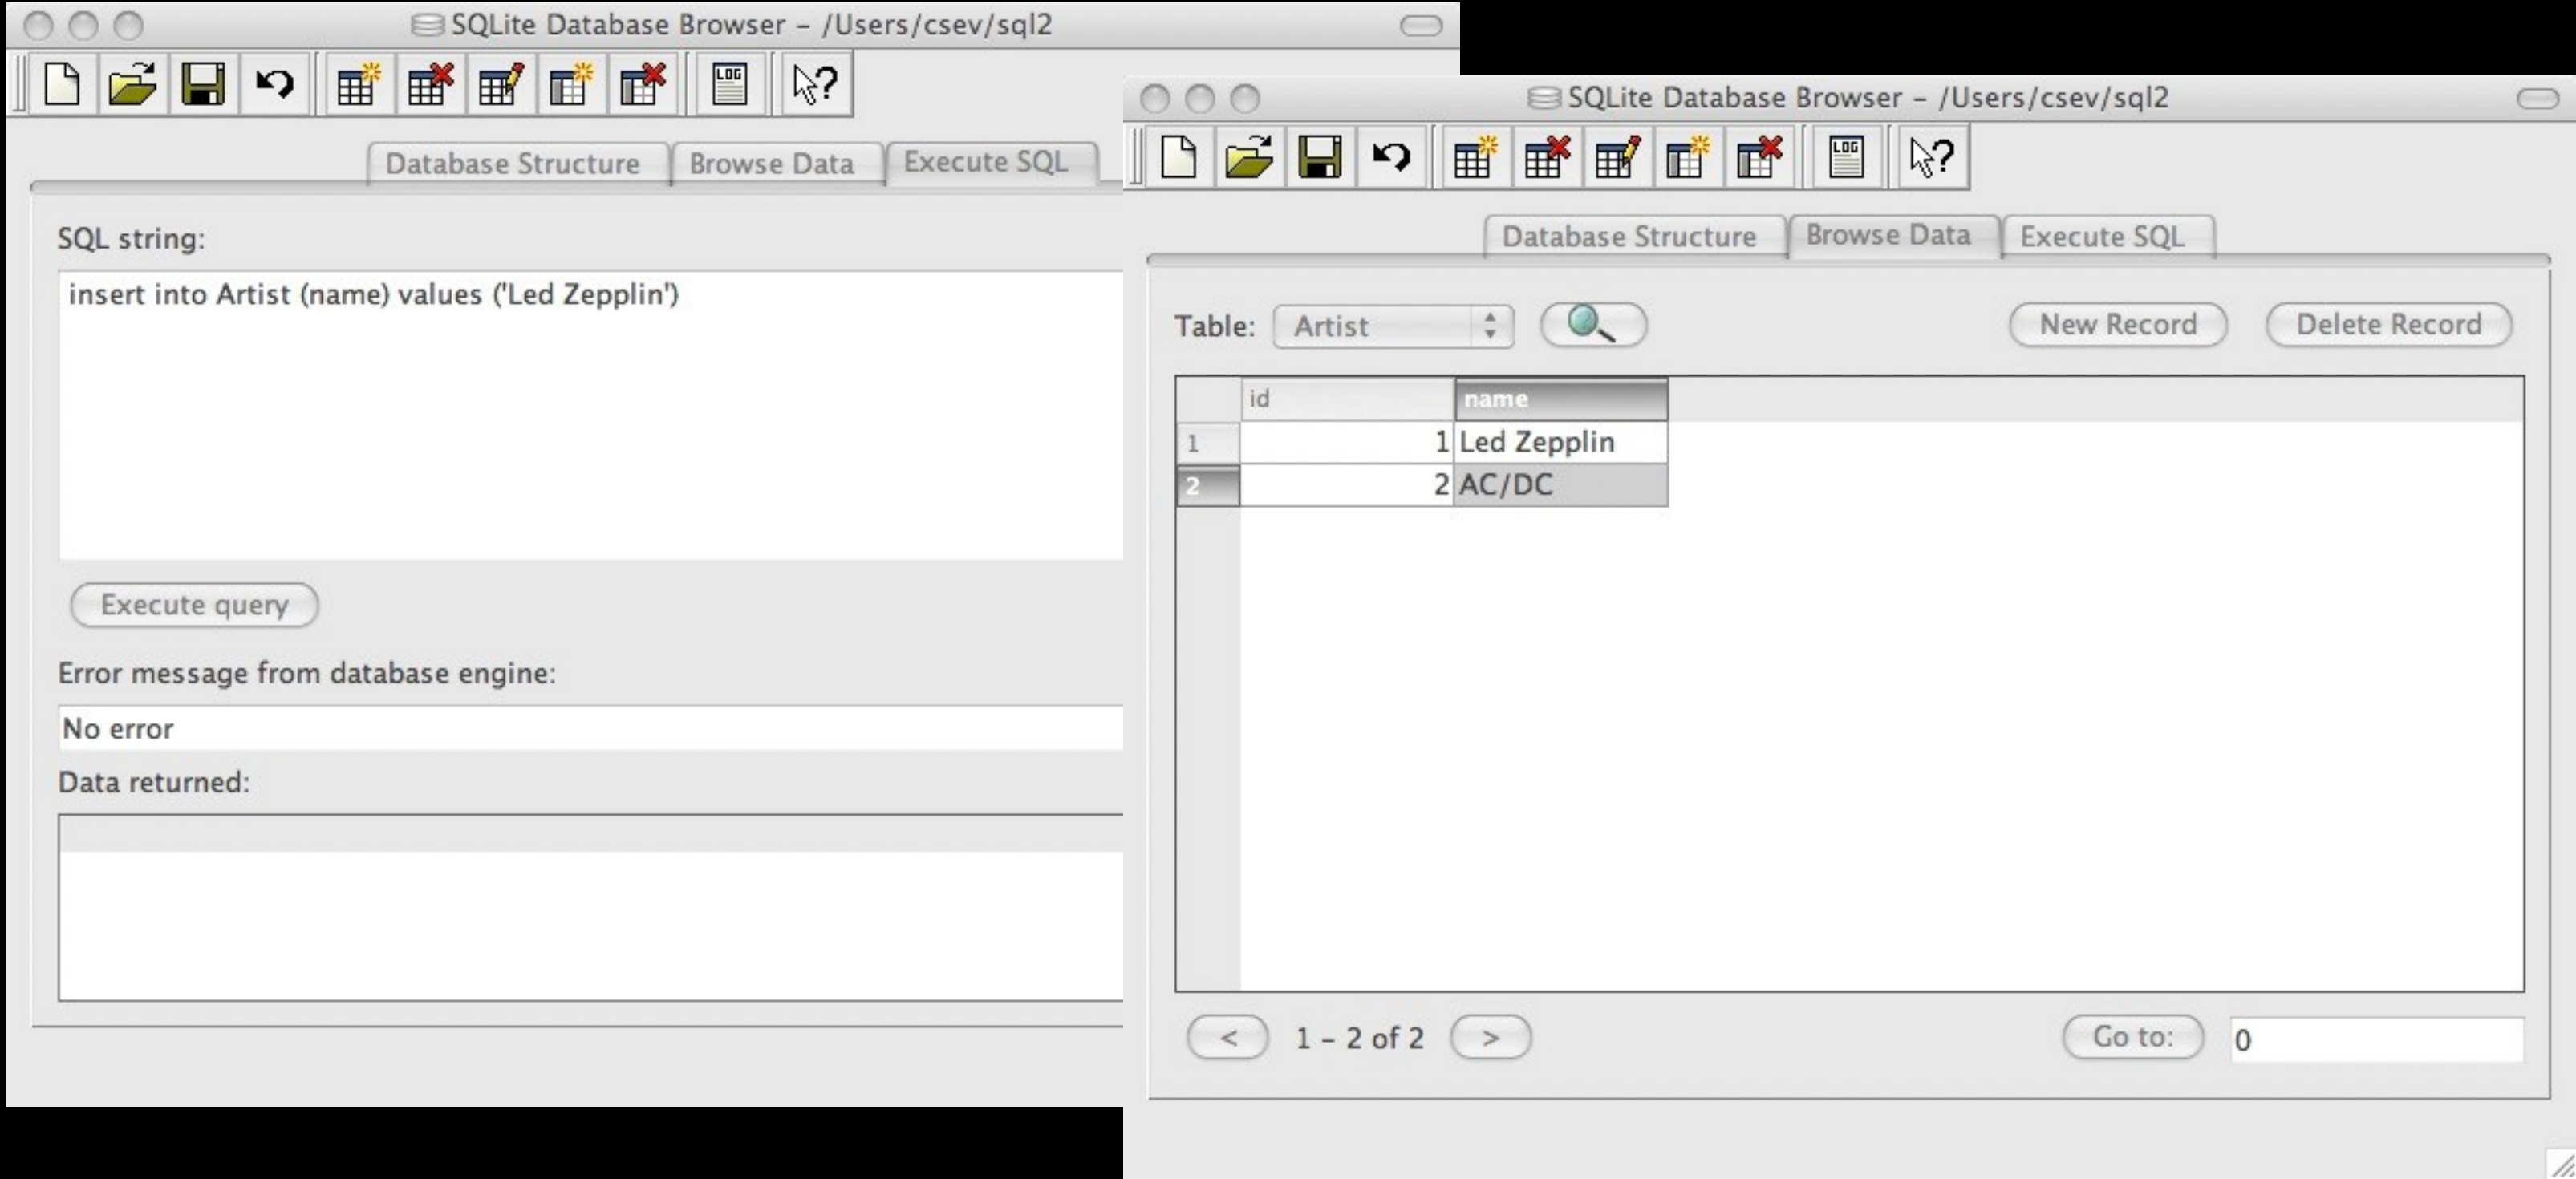
Naming FK `artist_id` is a convention.



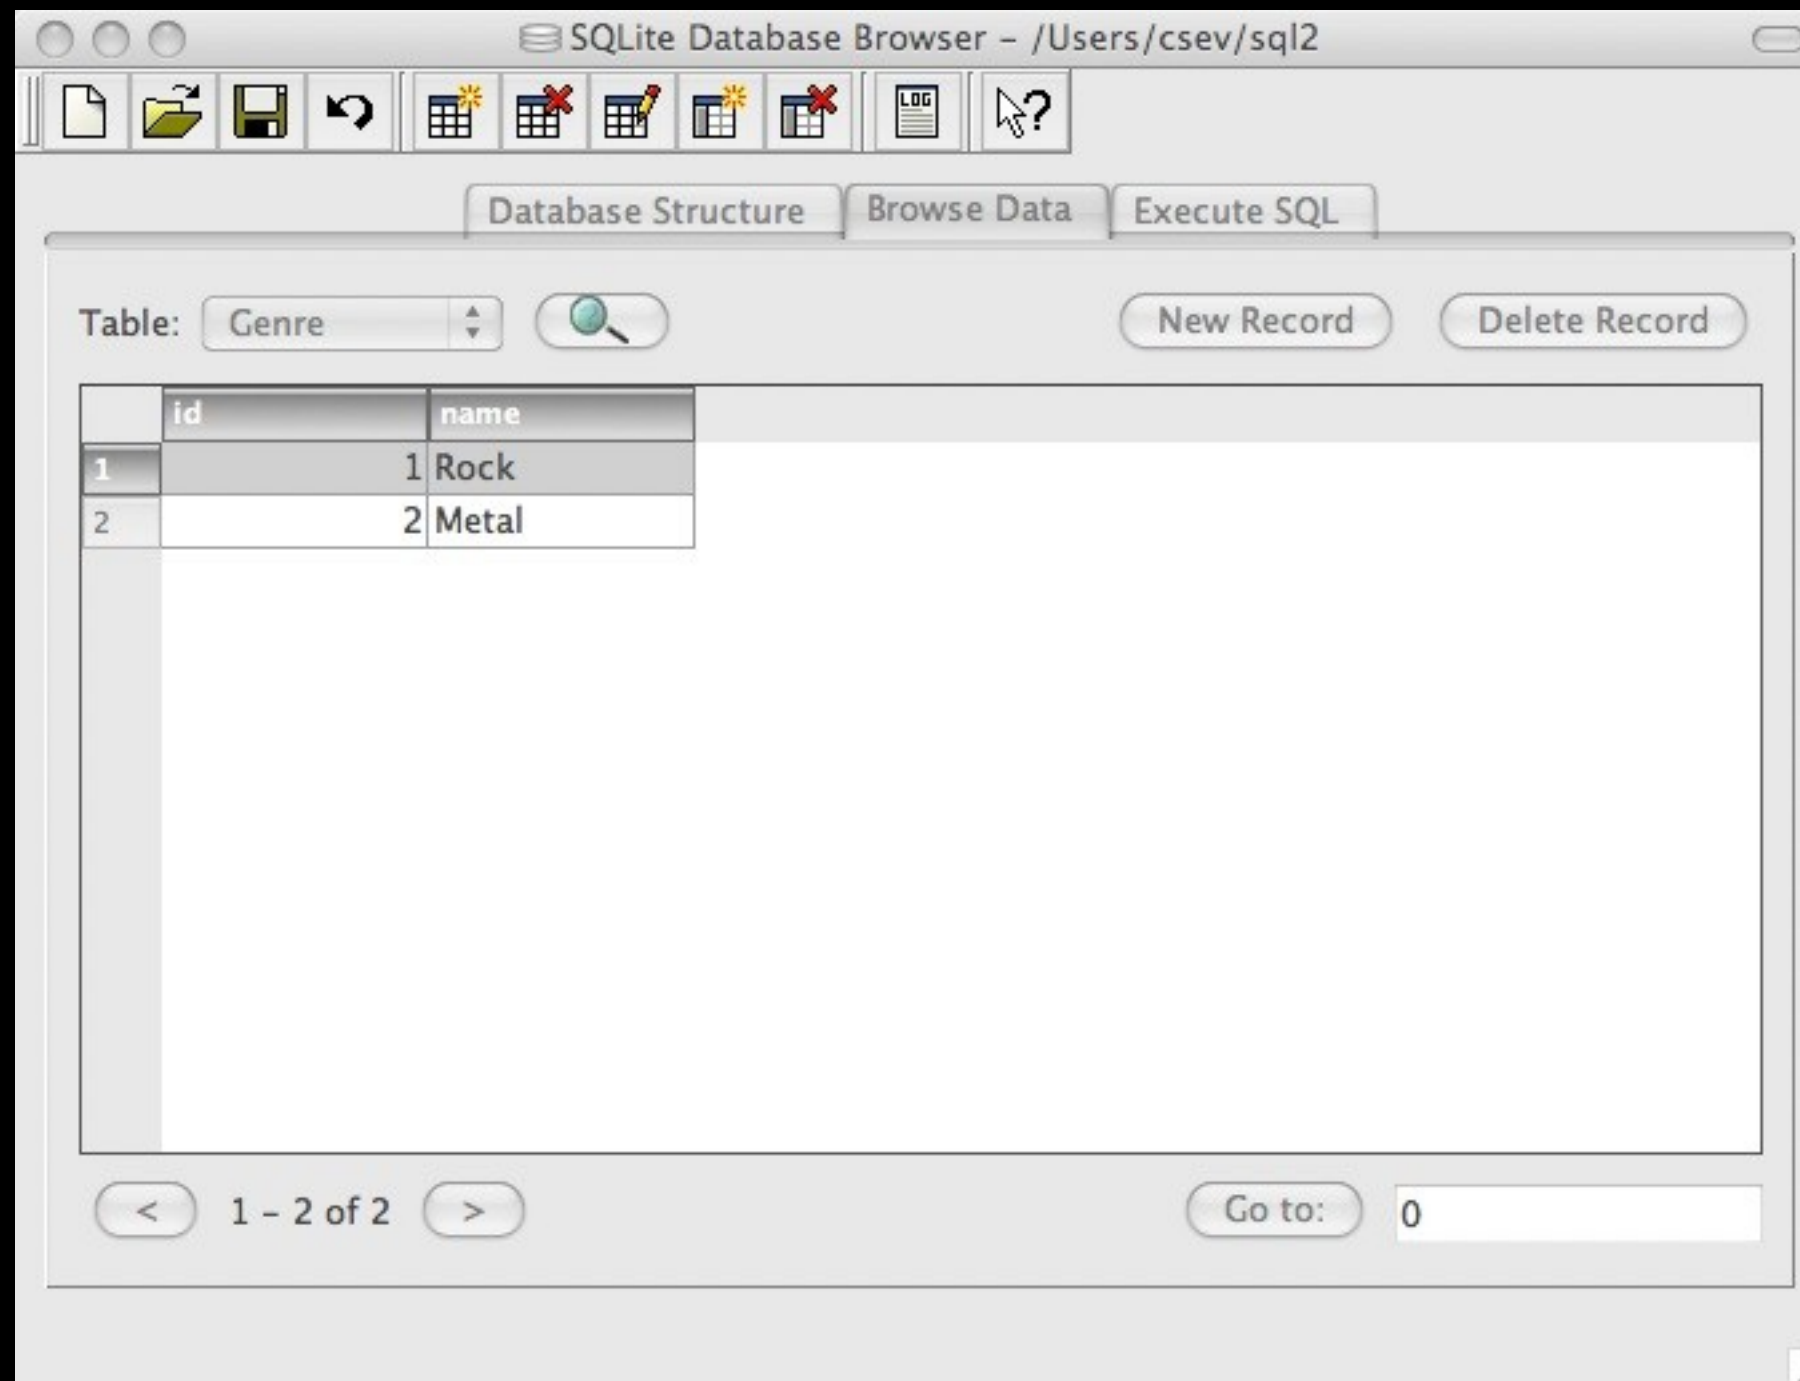
SQLite Database Browser - /Users/csev/sql2

Database Structure | Browse Data | Execute SQL

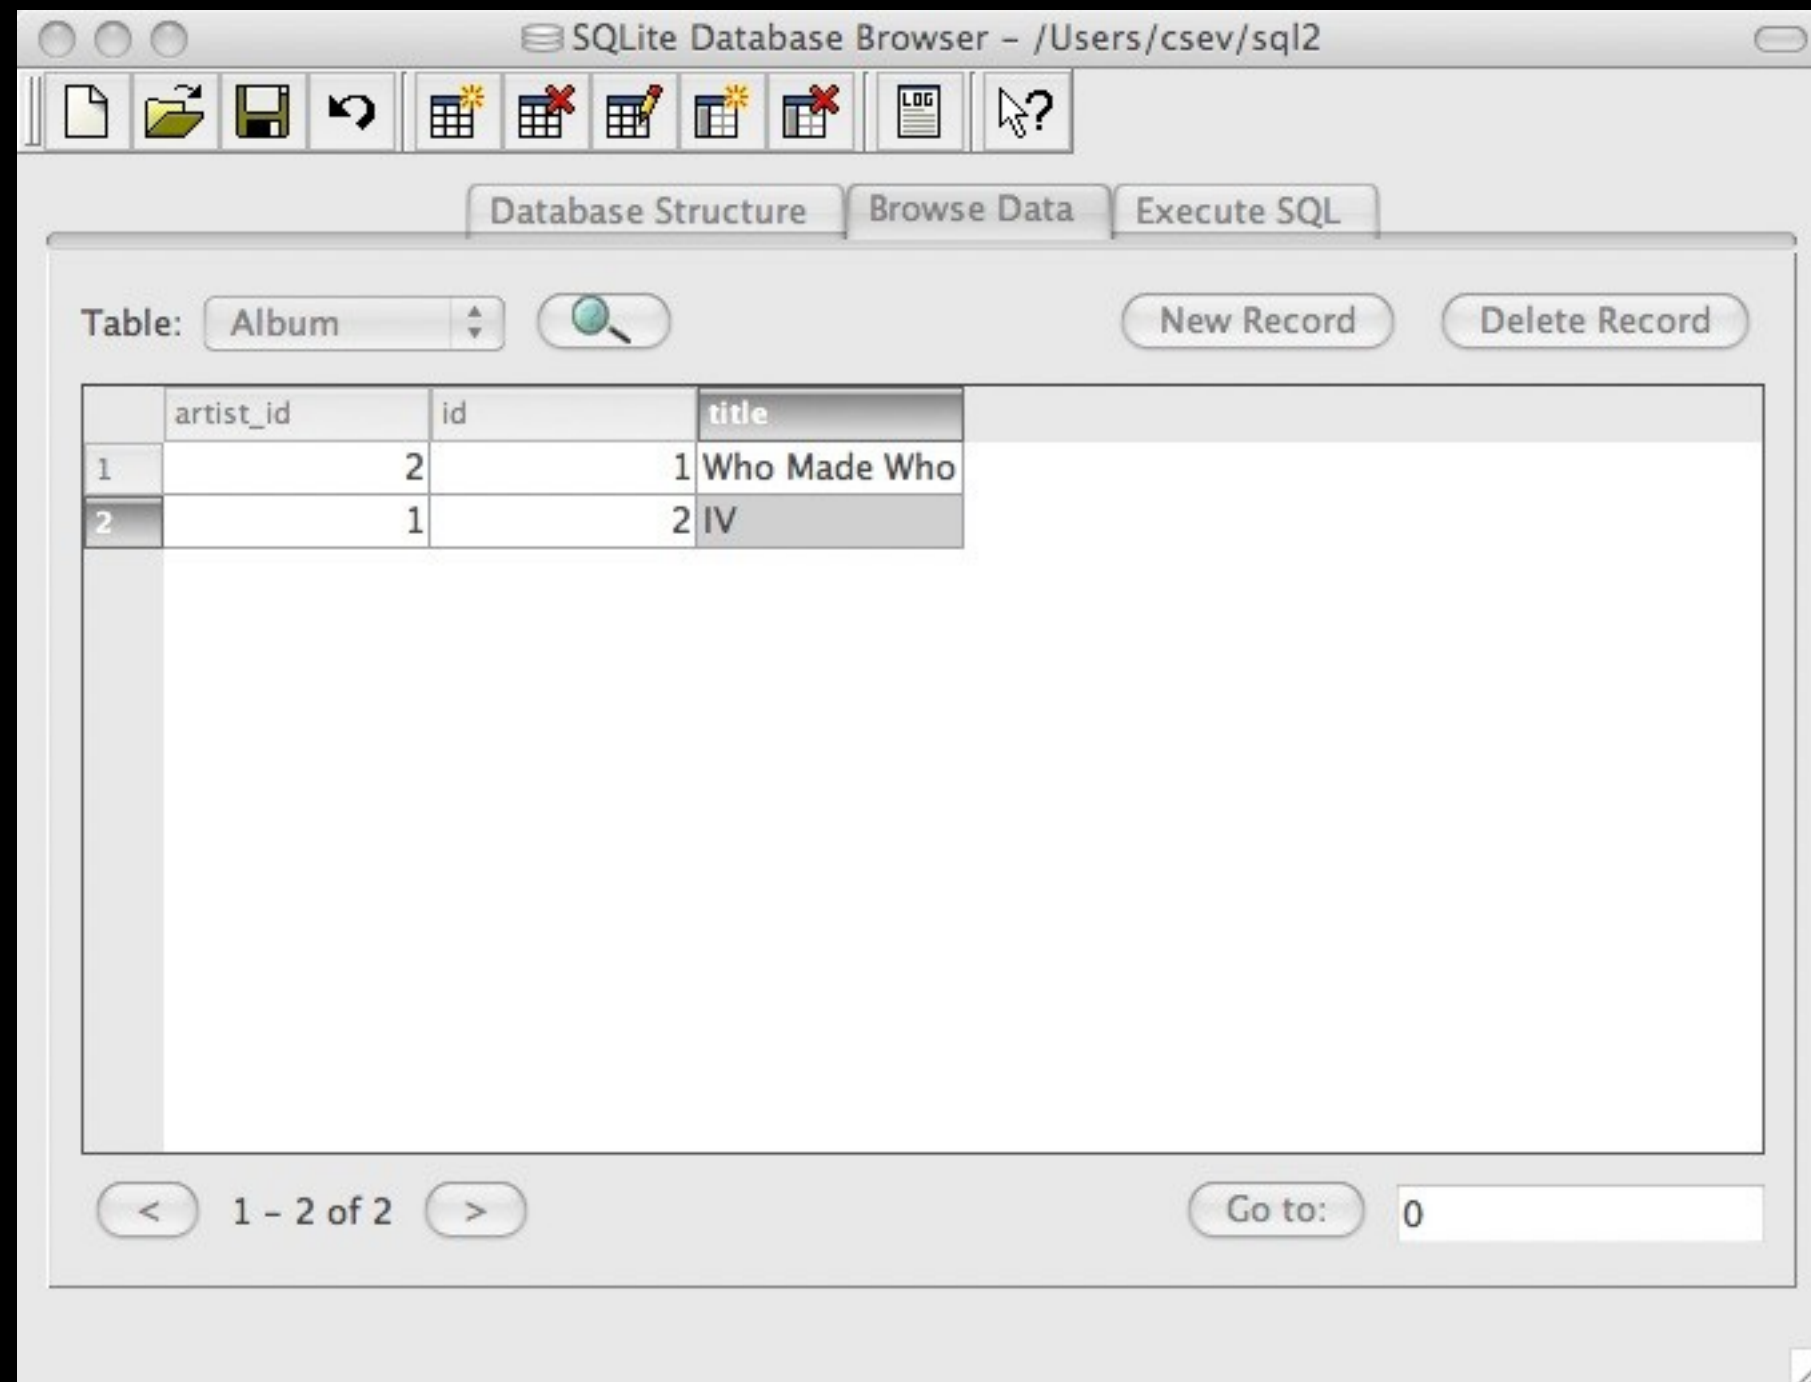
Name	Object	Type	Schema
▼ Album	table		CREATE TABLE Album (artist_id NUMERIC, id INTEGER PRI...
artist_id	field	NUMERIC	
id	field	INTEGER PRIMARY KEY	
title	field	TEXT	
▼ Artist	table		CREATE TABLE Artist (id INTEGER PRIMARY KEY, name TE...
id	field	INTEGER PRIMARY KEY	
name	field	TEXT	
▼ Genre	table		CREATE TABLE Genre (id INTEGER PRIMARY KEY, name T...
id	field	INTEGER PRIMARY KEY	
name	field	TEXT	
▼ Track	table		CREATE TABLE Track (album_id NUMERIC, count NUMERI...
album_id	field	NUMERIC	
count	field	NUMERIC	
genre_id	field	NUMERIC	
id	field	INTEGER PRIMARY KEY	
len	field	NUMERIC	
rating	field	NUMERIC	
title	field	TEXT	



insert into Artist (name) values ('Led Zeppelin')
insert into Artist (name) values ('AC/DC')

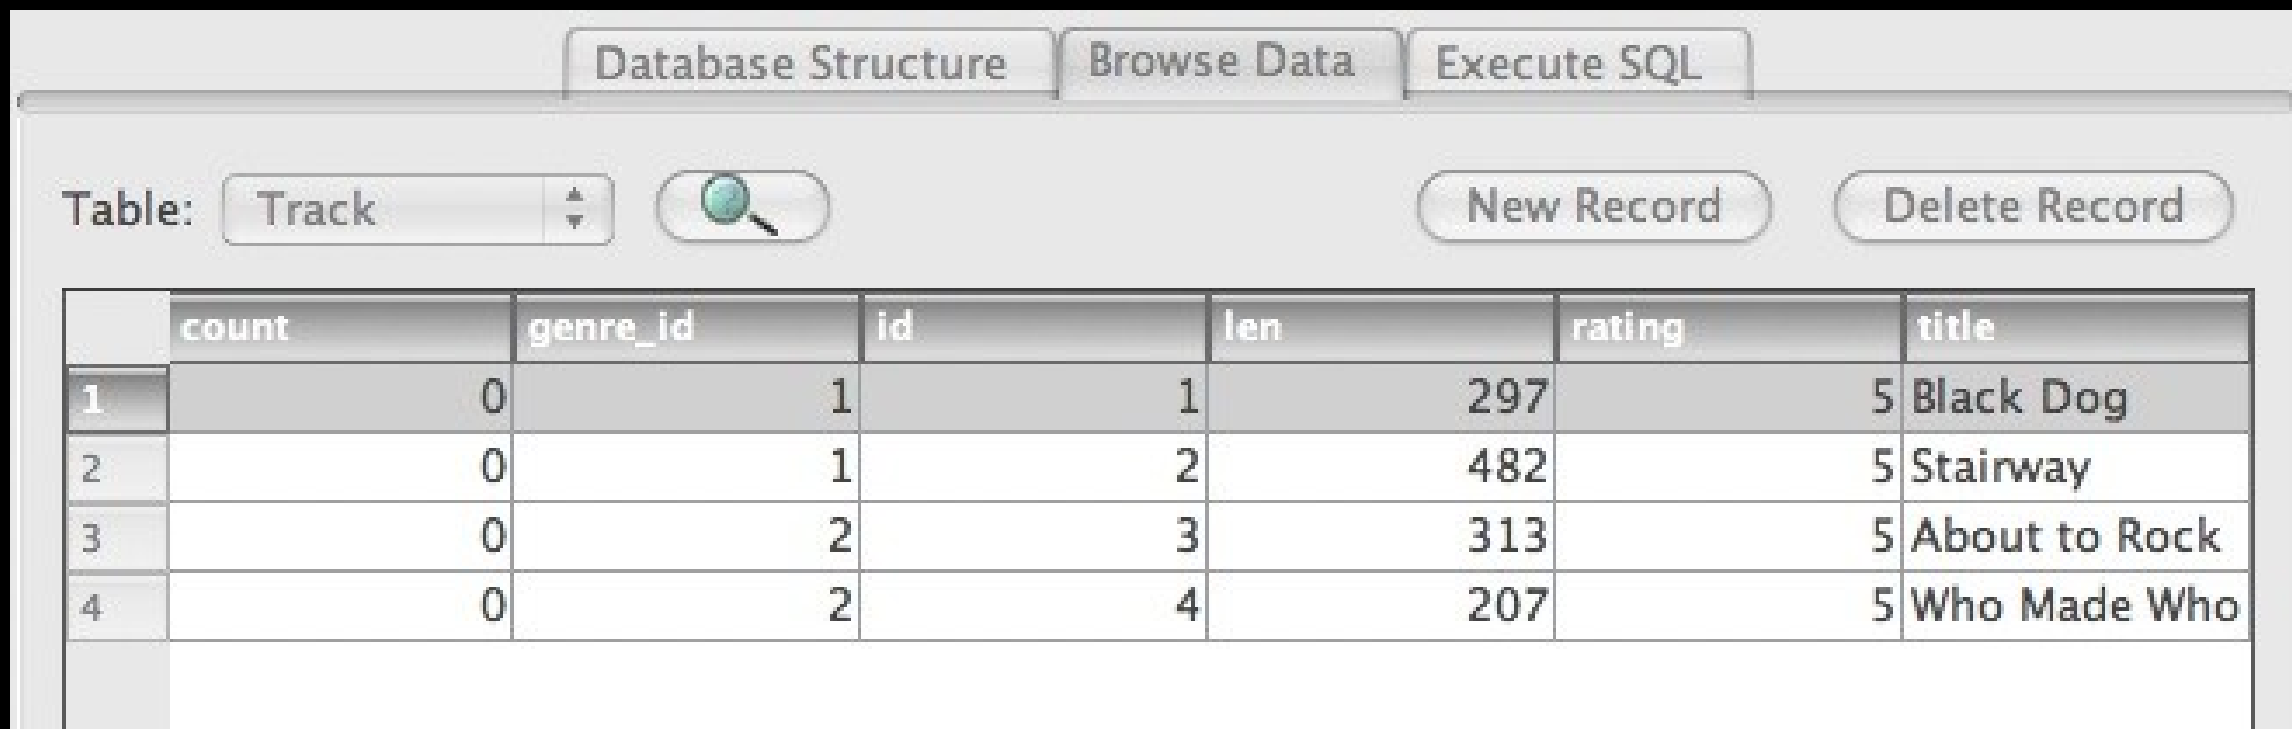


insert into Genre (name) values ('Rock')
insert into Genre (name) values ('Metal')



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



The screenshot shows a database management interface with three tabs: 'Database Structure', 'Browse Data', and 'Execute SQL'. The 'Track' table is selected, and the interface includes buttons for 'New Record' and 'Delete Record'. The table contains the following data:

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

We have relationships!

Table: Track

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

Table: Album

	artist_id	id	title
1	2	1	Who Made Who
2	1	2	IV

Table: Genre

	id	name
1	1	Rock
2	2	Metal

Table: Artist

	id	name
1	1	Led Zeppelin
2	2	AC/DC

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 the keys which make the connection between the tables using an ON clause

Table: Track

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

Data returned:

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

Table: Genre

	id	name
1	1	Rock
2	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
```

Data returned:

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

What we want to see

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						25
<input checked="" type="checkbox"/>	Paranoid						22
<input checked="" type="checkbox"/>	Planet Caravan						25
<input checked="" type="checkbox"/>	Iron Man						26
<input checked="" type="checkbox"/>	Electric Funeral						22
<input checked="" type="checkbox"/>	Hand of Doom						23
<input checked="" type="checkbox"/>	Rat Salad						31
<input checked="" type="checkbox"/>	Jack the Stripper/Fairies						24
<input checked="" type="checkbox"/>	Bomb Squad (TECH)						1
<input checked="" type="checkbox"/>	clay techno						2
<input checked="" type="checkbox"/>	Heavy						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

Data returned:

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

Python and SQLite3

<http://www.python.org/doc/2.5.2/lib/module-sqlite3.html>

SQLite3 is built into Python

- Since SQLite is simple and small and designed to be “embedded” - Python decided to embed SQLite into Python
- You simply “import sqlite3” and open a connection to the database and start doing SQL commands

<http://www.python.org/doc/2.5.2/lib/module-sqlite3.html>

SQLite3 is built into Python

```
import sqlite3

# Open up the database file and get a cursor
conn = sqlite3.connect('music.db')
c = conn.cursor()

print "Genre Rows"
c.execute('select * from Genre')
for row in c :
    print row
```

SQLite stores all tables and data in a single file.

```
$ python sql1.py
Genre Rows
(1, u'Rock')
(2, u'Metal')
$ ls
music.db  sql1.py  sql2.py
```

```
import sqlite3

# Open up the database file and get a cursor
conn = sqlite3.connect('music.db')
c = conn.cursor()

print "Inserting Country"
c.execute('insert into Genre (name) values ( ? )', ( 'Country', ) )

print "Genre Rows"
c.execute('select * from Genre')
for row in c :
    print row

print "Deleting Country"
c.execute("delete from Genre where name='Country'")

print "Genre Rows"
c.execute('select * from Genre')
for row in c :
    print row
```

```
$ python sql2.py
Inserting Country
Genre Rows
(1, u'Rock')
(2, u'Metal')
(3, u'Country')
Deleting Country
Genre Rows
(1, u'Rock')
(2, u'Metal')
```

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 SI572 - Database Design

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