History and Types of Databases

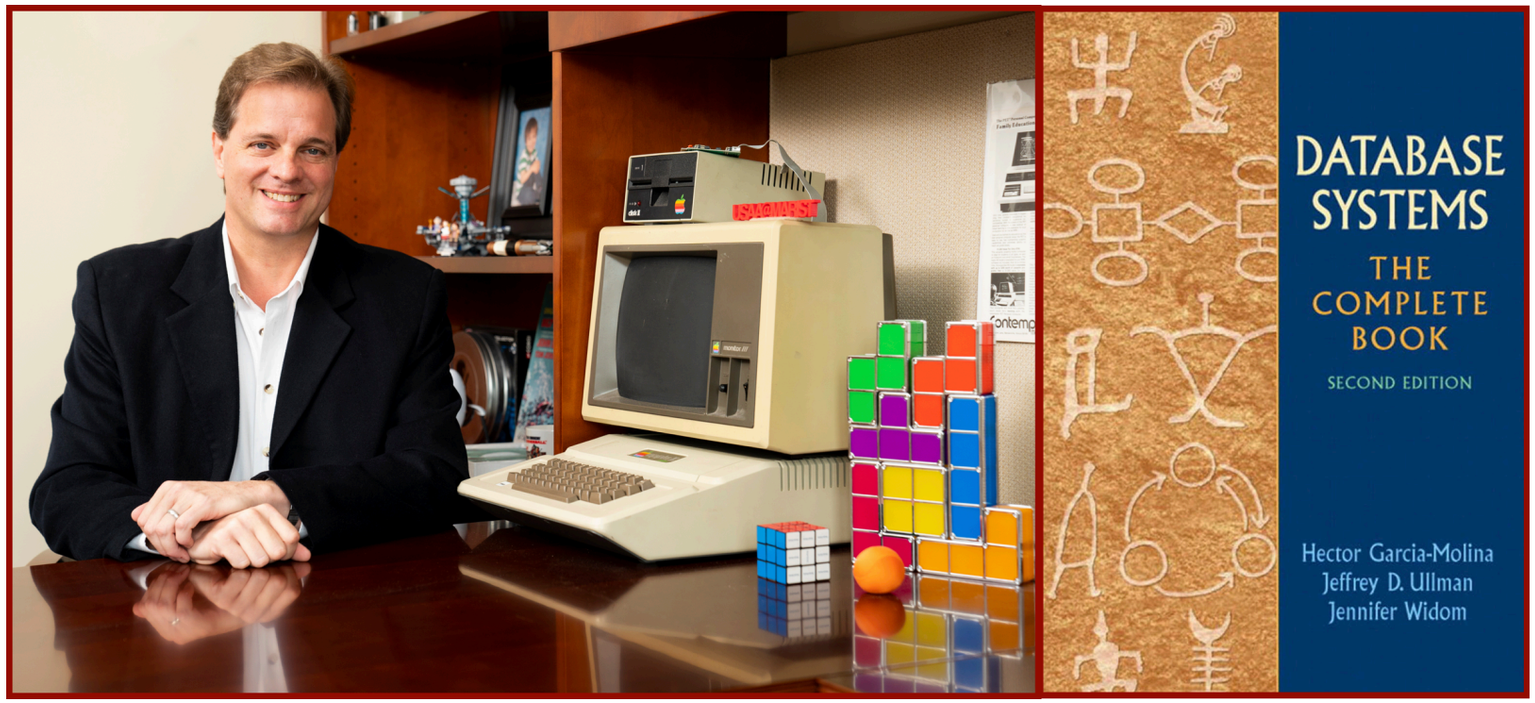

Alan G. Labouseur, Ph.D.
Alan.Labouseur@Marist.edu

Data?

What is data?

?

Data?

What is data?

```
1  007  42  21  12  90  125  86 75 30 9
```

Data?

What is data?

What does it mean?

```
1  007  42  21  12  90  125  86 75 30 9
```

Data?

What is data?

What does it mean?

```
1  007  42  21  12  90  125  86 75 30 9
```

We need context.

Data?

What is data?

What does it mean?

```
1  007  42  21  12  90  125  86 75 30 9
```

With context we can draw conclusions from data.

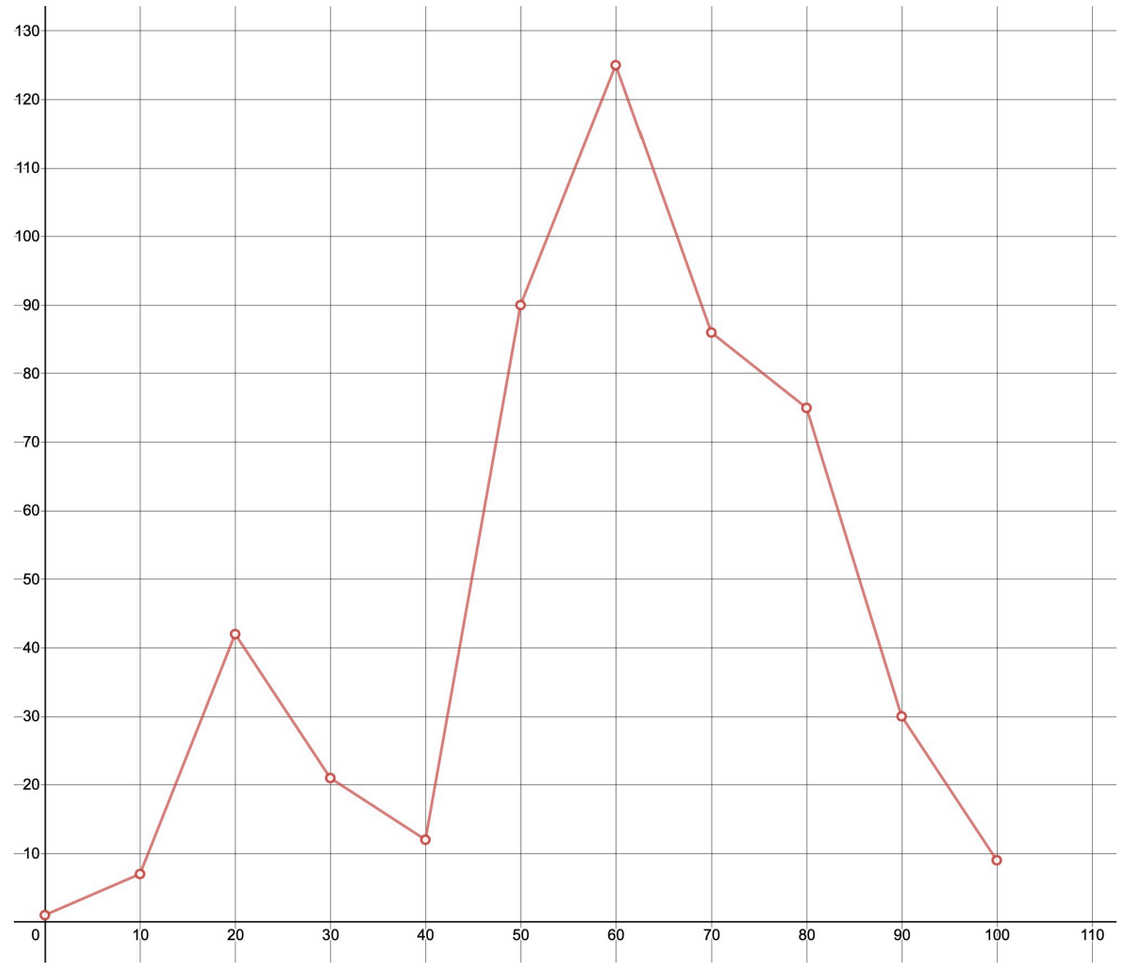

With context
we have **information**.

Data is Dangerous

What is data?

What does it mean?

```
1  007  42  21  12  90  125  86 75 30 9
```

What if we're wrong?

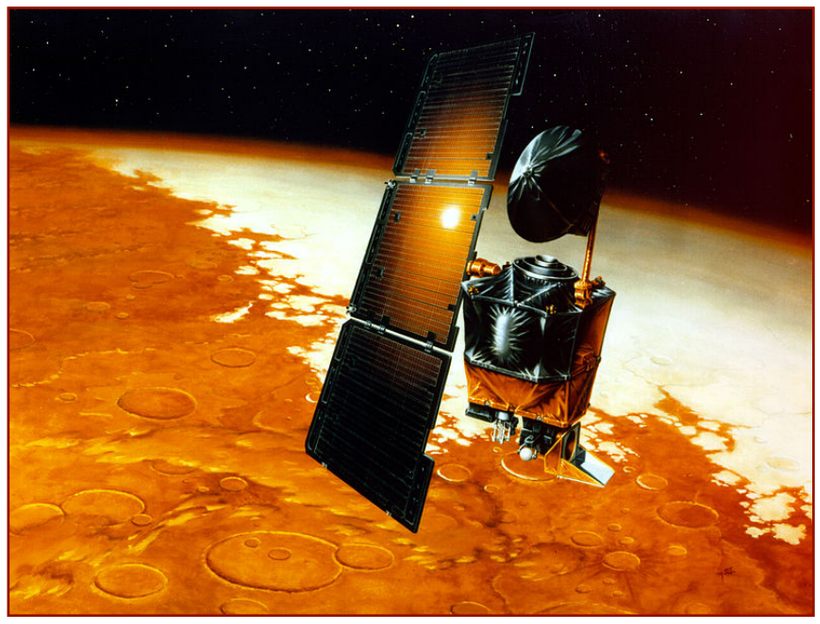

Data is Dangerous, Information is Valuable

Data + Context = Information

Information is valuable.
Information is difficult to obtain.
Information is what we want.

And to get it, we need to impose structure for context.

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

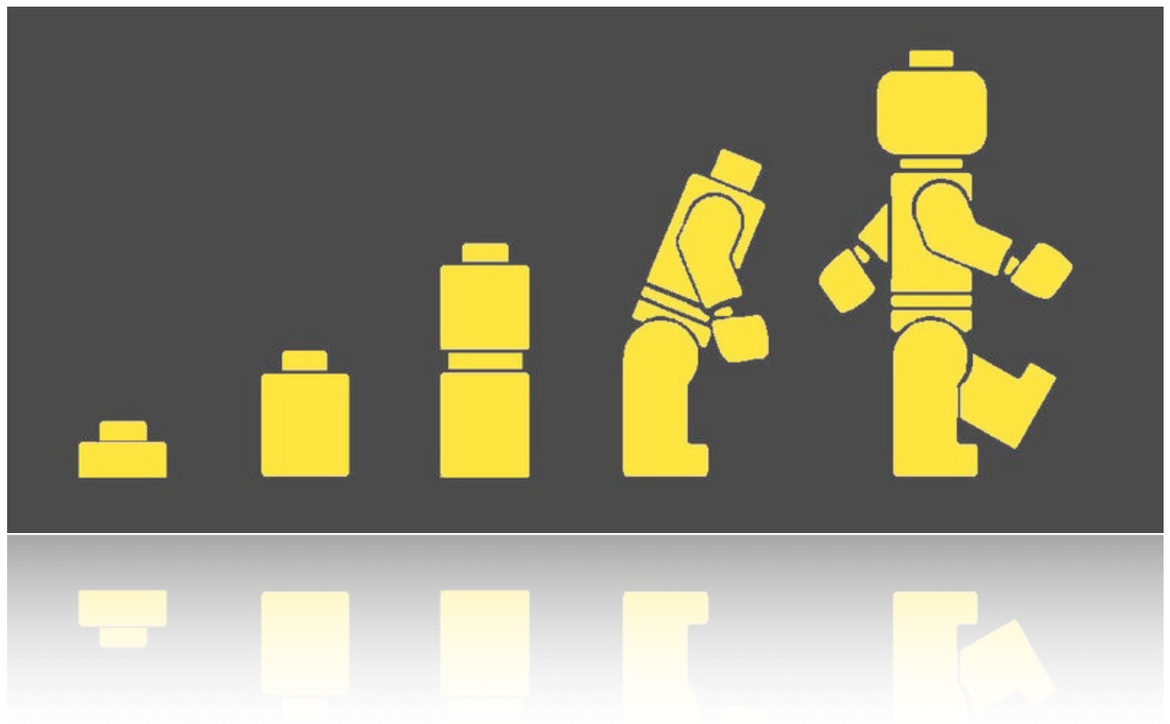

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

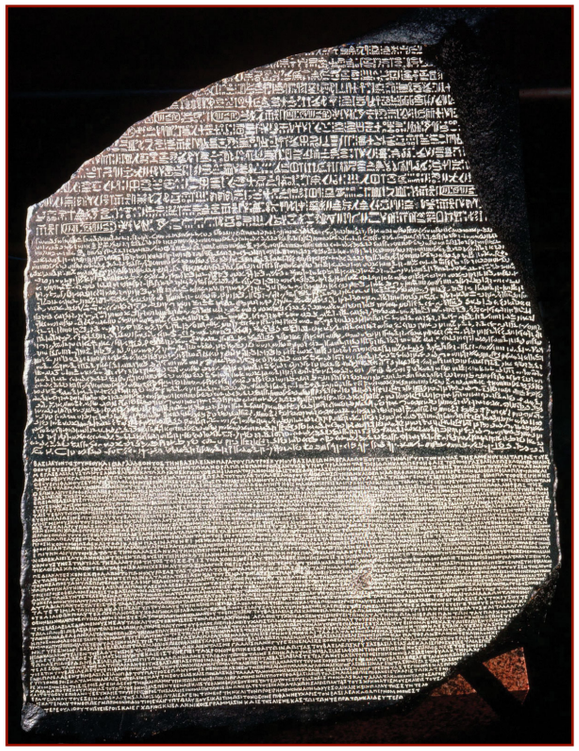

Heavy data

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

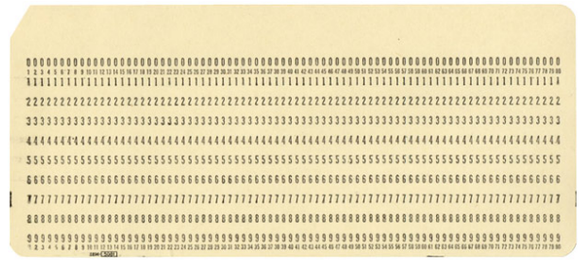

1890 Census

Big data

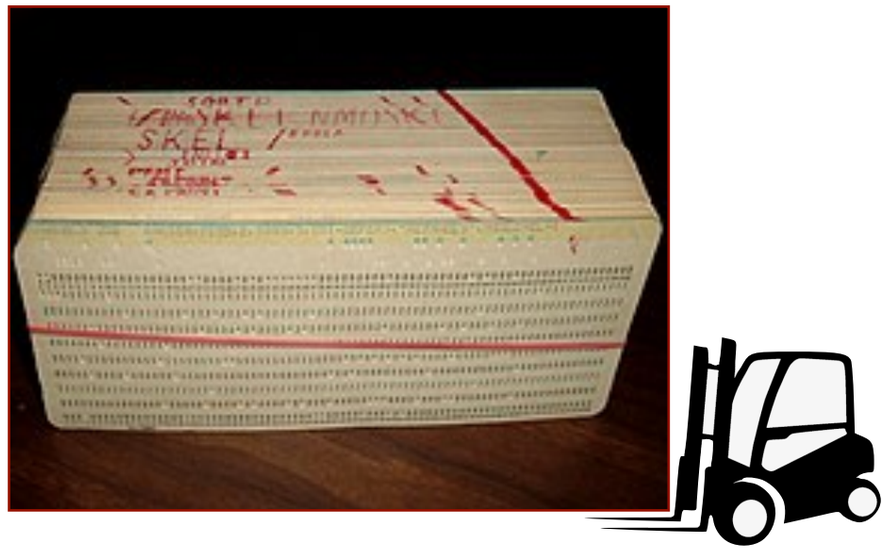

(Still heavy.)

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

- Files of Records of Fields for a D&D-type game

Players File

Player 1 Record

Player 1 Fields

```
id    : 1
name  : James
rank  : Captain
items: wand,
       gem
```

Player 2 Record

Player 2 Fields

```
id    : 2
name  : Leonard
rank  : Admiral
items: gem,
       mace
```

Items File

Item 1 Record

Item 1 Fields

```
id    : A
name  : wand
desc  : ...
```

Item 2 Record

Item 2 Fields

```
id    : B
name  : gem
desc  : ...
```

Item 3 Record

Item 3 Fields

```
id    : C
name  : mace
desc  : ...
```

Item 4 Record

Item 4 Fields

```
id    : D
name  : sword
desc  : ...
```

Evolution

Consider the evolution of Data Management†

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

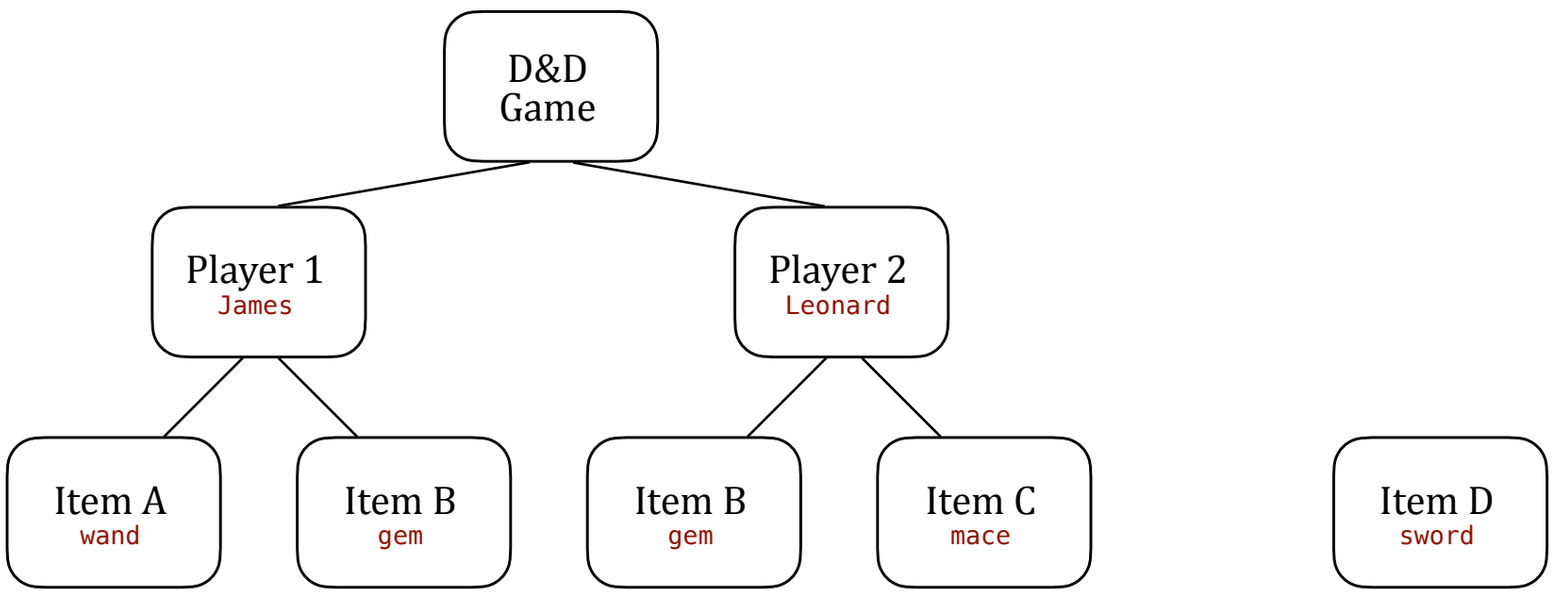

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

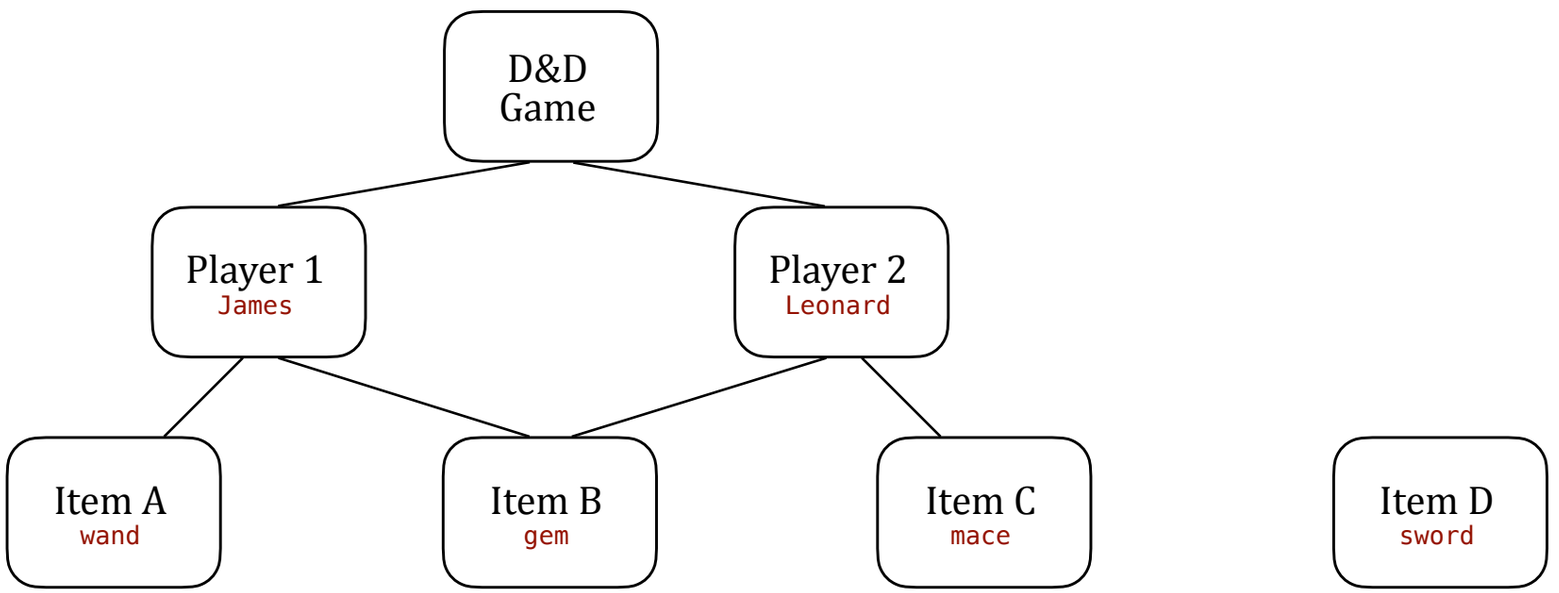

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

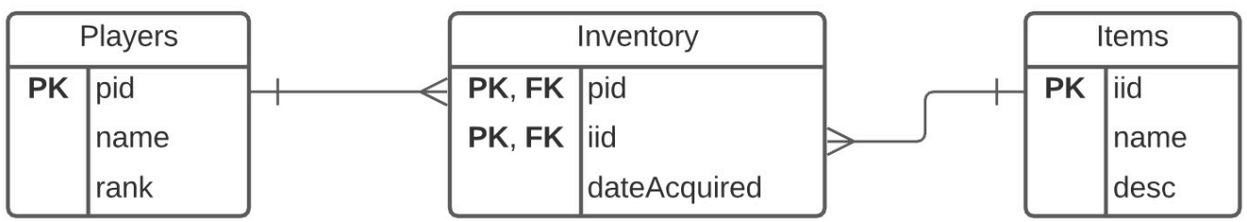


```
CAP=# select *
CAP-# from Players;
 pid |  name   |  rank
-----+---------+---------
   1 | James   | Captain
   2 | Leonard | Admiral
(2 rows)

CAP=# select *
CAP-# from Items;
 iid | name  | descr
-----+-------+-------
 A   | wand  | ...
 B   | gem   | ...
 C   | mace  | ...
 D   | sword | ...
(4 rows)

CAP=# select *
CAP-# from Inventory;
 pid | iid | dateacquired
-----+-----+--------------
   1 | A   | 2020-01-23
   1 | B   | 2020-01-23
   2 | B   | 2020-01-23
   2 | C   | 2020-01-23
(4 rows)
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

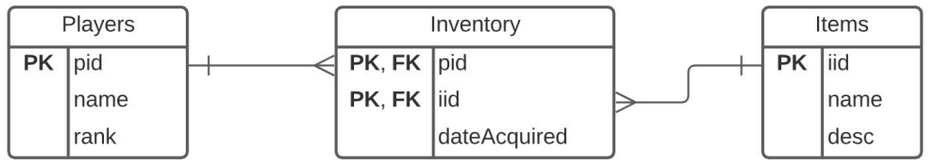


```
DB=# -- Players and their Items
DB=# select Players.name, Items.name
DB=# from Players inner join Inventory on Players.pid = Inventory.pid
DB=#              inner join Items on Inventory.iid = Items.iid
DB=# ;

  name   | name
---------+------
 James   | wand
 James   | gem
 Leonard | gem
 Leonard | mace
(4 rows)
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

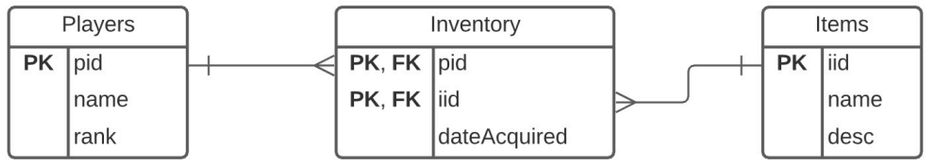


```
DB=# -- Unused Items
DB=# select *
DB=# from Items
DB=# where iid not in (select iid
DB=#                   from Inventory);

 iid | name  | descr
-----+-------+-------
 D   | sword | ...
(1 row)
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

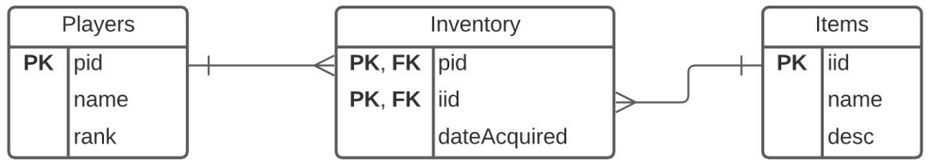


```
DB=# -- Item use count v1
DB=# select iid, count(iid)
DB=# from Inventory
DB=# group by iid
DB=# order by count(iid) DESC;

 iid | count
-----+-------
 B   |     2
 C   |     1
 A   |     1
(3 rows)
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

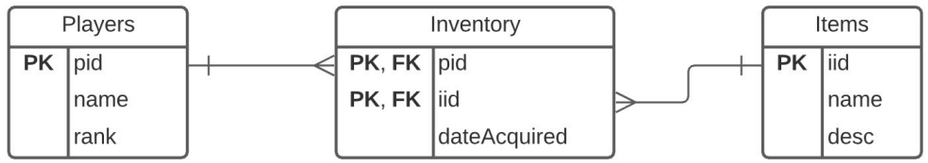


```
DB=# -- Item use count v2, now with item names!
DB=# select Inventory.iid, Items.name, count(Inventory.iid)
DB=# from Inventory inner join Items on Inventory.iid = Items.iid
DB=# group by Inventory.iid, Items.name
DB=# order by count(Inventory.iid) DESC
DB=# ;

 iid | name | count
-----+------+-------
 B   | gem  |     2
 A   | wand |     1
 C   | mace |     1
(3 rows)
```

Evolution

SQL Script for Player, Items, and Inventory tables and a few queries

```
create table Players (
   pid   int not null,
   name  text,
   rank text,
 primary key (pid)
);

insert into Players(pid, name, rank)
values (1, 'James',   'Captain'),
       (2, 'Leonard', 'Admiral');

select *
from Players;


create table Items (
   iid   char(1) not null,
   name  text,
   descr text,
 primary key (iid)
);

insert into Items (iid, name, descr)
values ('A', 'wand',  '...'),
       ('B', 'gem',   '...'),
       ('C', 'mace',  '...'),
       ('D', 'sword', '...');

select *
from Items;
```

```
create table Inventory (
   pid          int     not null references Players(pid),
   iid          char(1) not null references Items(iid),
   dateAcquired date,
  primary key(pid, iid)
);

insert into Inventory (pid, iid, dateAcquired)
values (1, 'A', '2020-01-23'),
       (1, 'B', '2020-01-23'),
       (2, 'B', '2020-01-23'),
       (2, 'C', '2020-01-23');

select *
from Inventory;

-- Players and their Items
select Players.name, Items.name
from Players inner join Inventory on Players.pid = Inventory.pid
             inner join Items on Inventory.iid = Items.iid;

-- Unused Items
select *
from Items
where iid not in (select iid
                  from Inventory);

-- Item use count v1
select iid, count(iid)
from Inventory
group by iid
order by count(iid) DESC;

-- Item use count v2, now with item names!
select Inventory.iid, Items.name, count(Inventory.iid)
from Inventory inner join Items on Inventory.iid = Items.iid
group by Inventory.iid, Items.name
order by count(Inventory.iid) DESC;
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

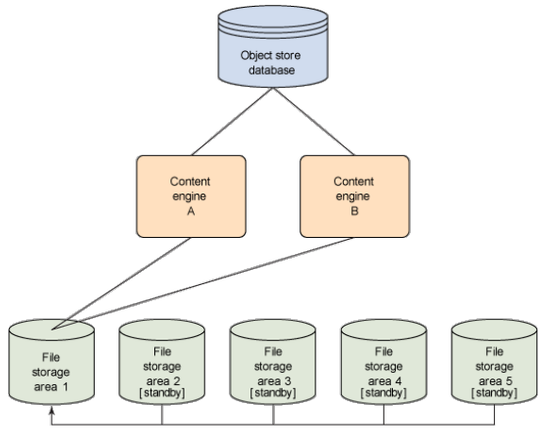


```
Players:
{
  {1, {James, Captain,{A,B}}}
  {2, {Leonard, Admiral,{B,C}}}
}

Items:
{
  {A, {wand, ..., {1}}}
  {B, {gem, ..., {1,2}}}
  {C, {mace, ..., {3}}}
  {D, {sword, ..., {}}}
}
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

name	*address*		*birthdate*	*movies*		
Fisher	*street*	*city*	9/9/99	*title*	*year*	*length*
	Maple	H'wood		Star Wars	1977	124
	Locust	Malibu		Empire	1980	127
				Return	1983	133
Hamill	*street*	*city*	8/8/88	*title*	*year*	*length*
	Oak	B'wood		Star Wars	1977	124
				Empire	1980	127
				Return	1983	133

Figure 10.15: A nested relation for stars and their movies

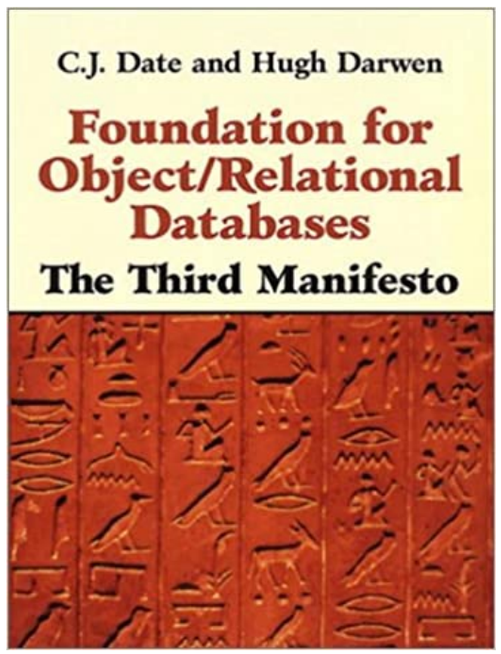

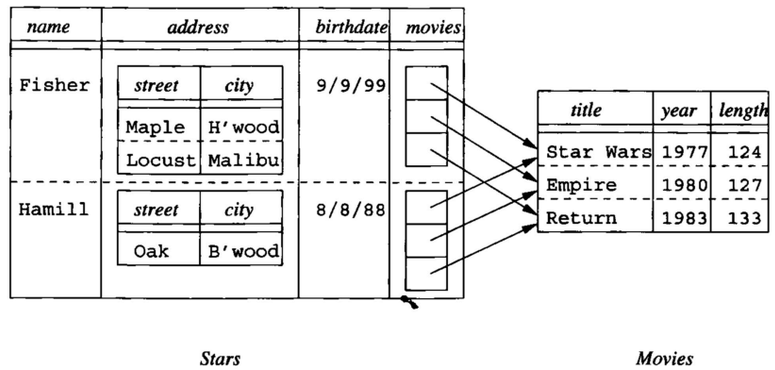

Figure 10.16: Sets of references as the value of an attribute

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

A **graph** is like a **network** in most ways.

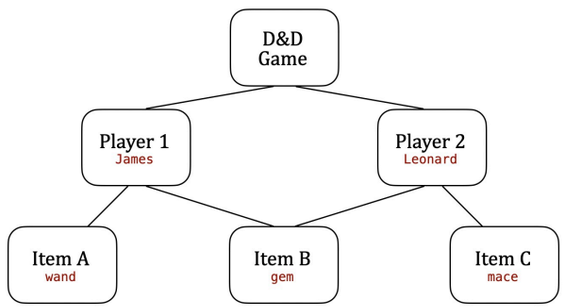

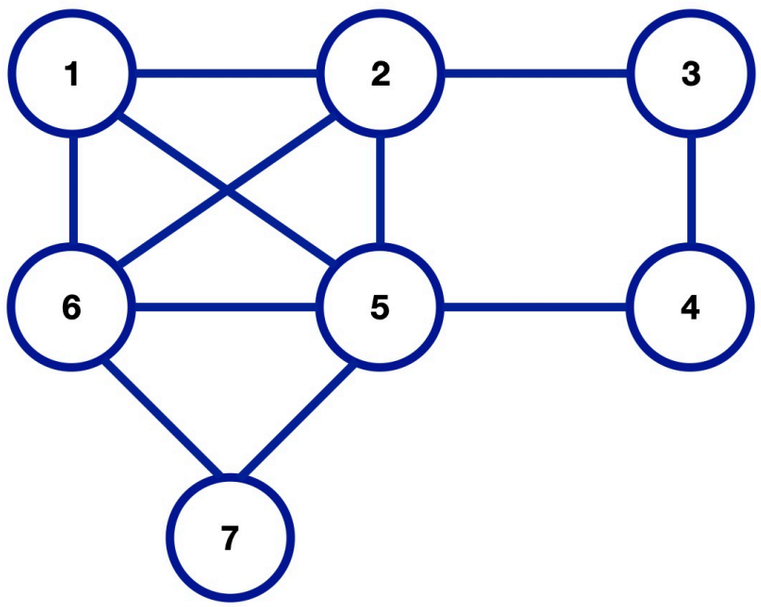

But graph databases are modern tools for managing them and gaining insight from the data pile.

Evolution

Consider the evolution of Data Management

Graph . . .

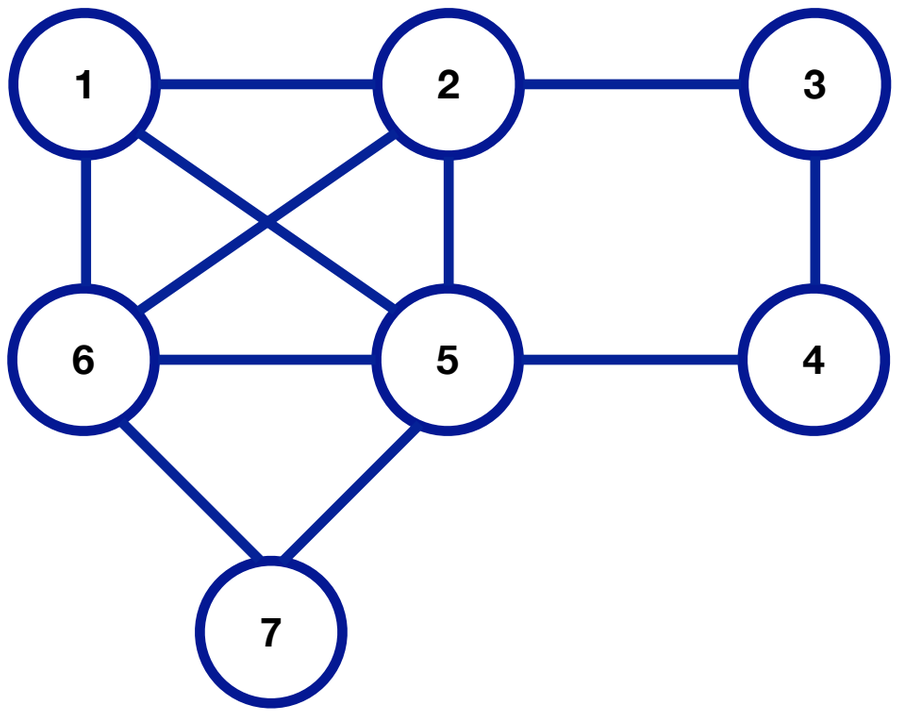

as Matrix

	1	2	3	4	5	6	7
1	.	1	.	.	1	1	.
2	1	.	1	.	1	1	.
3	.	1	.	1	.	.	.
4	.	.	1	.	1	.	.
5	1	1	.	1	.	1	1
6	1	1	.	.	1	.	1
7	.	.	.	.	1	1	.

Evolution

Consider the evolution of Data Management

Graph . . .

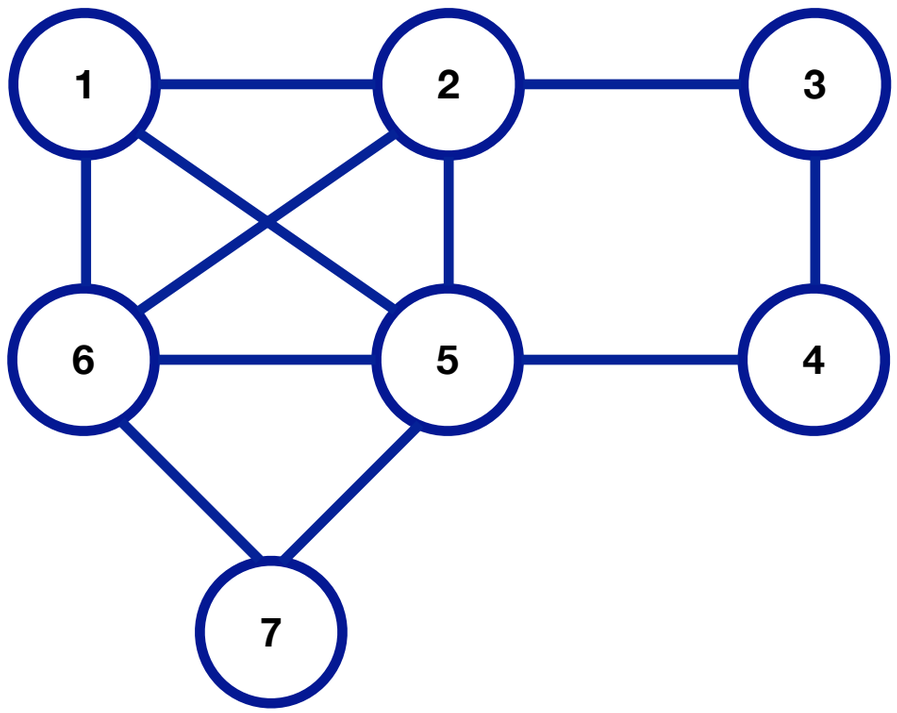

as Adjacency List

```
[1]  2  5  6
[2]  1  3  5  6
[3]  2  4
[4]  3  5
[5]  1  2  4  6  7
[6]  1  2  5  7
[7]  5  6
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

G* The Dynamic Graph Database

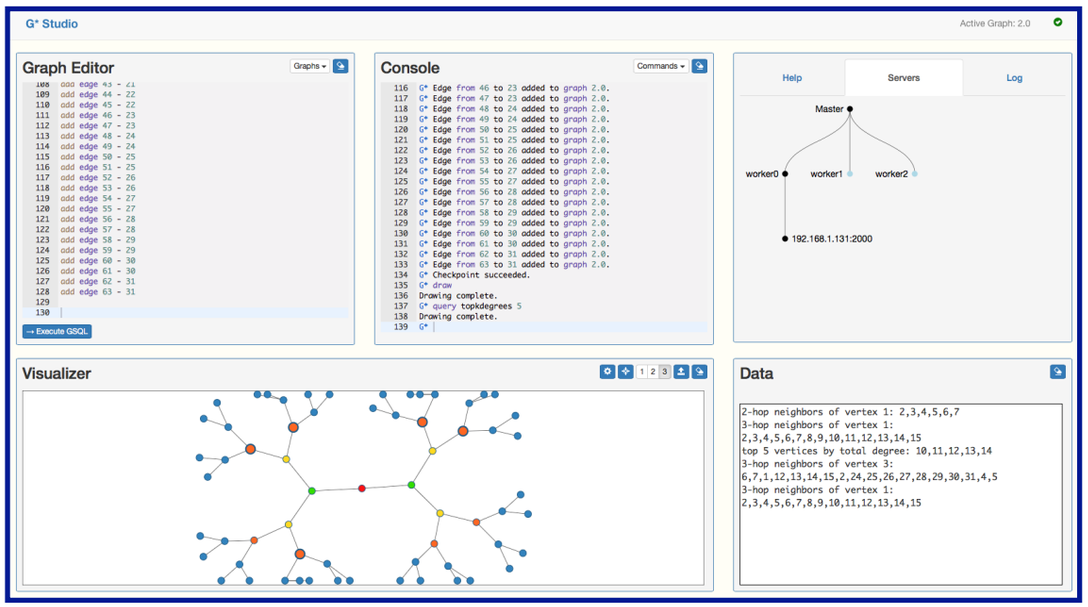

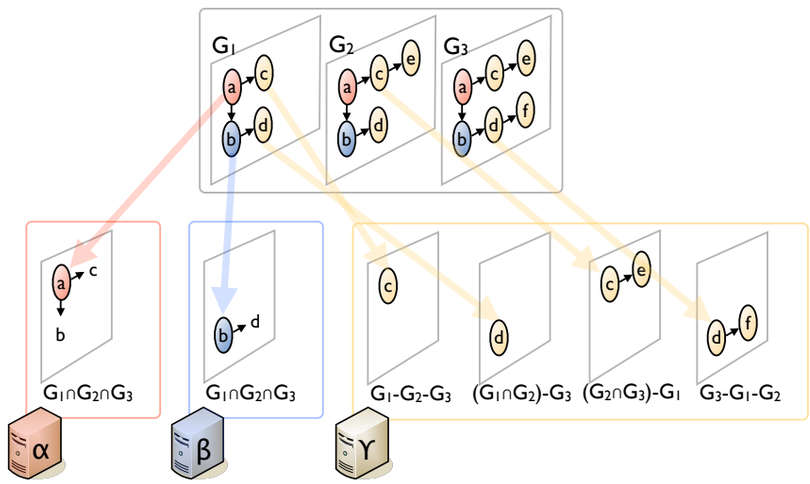

Evolution

G* The Dynamic Graph Database

Browser Application

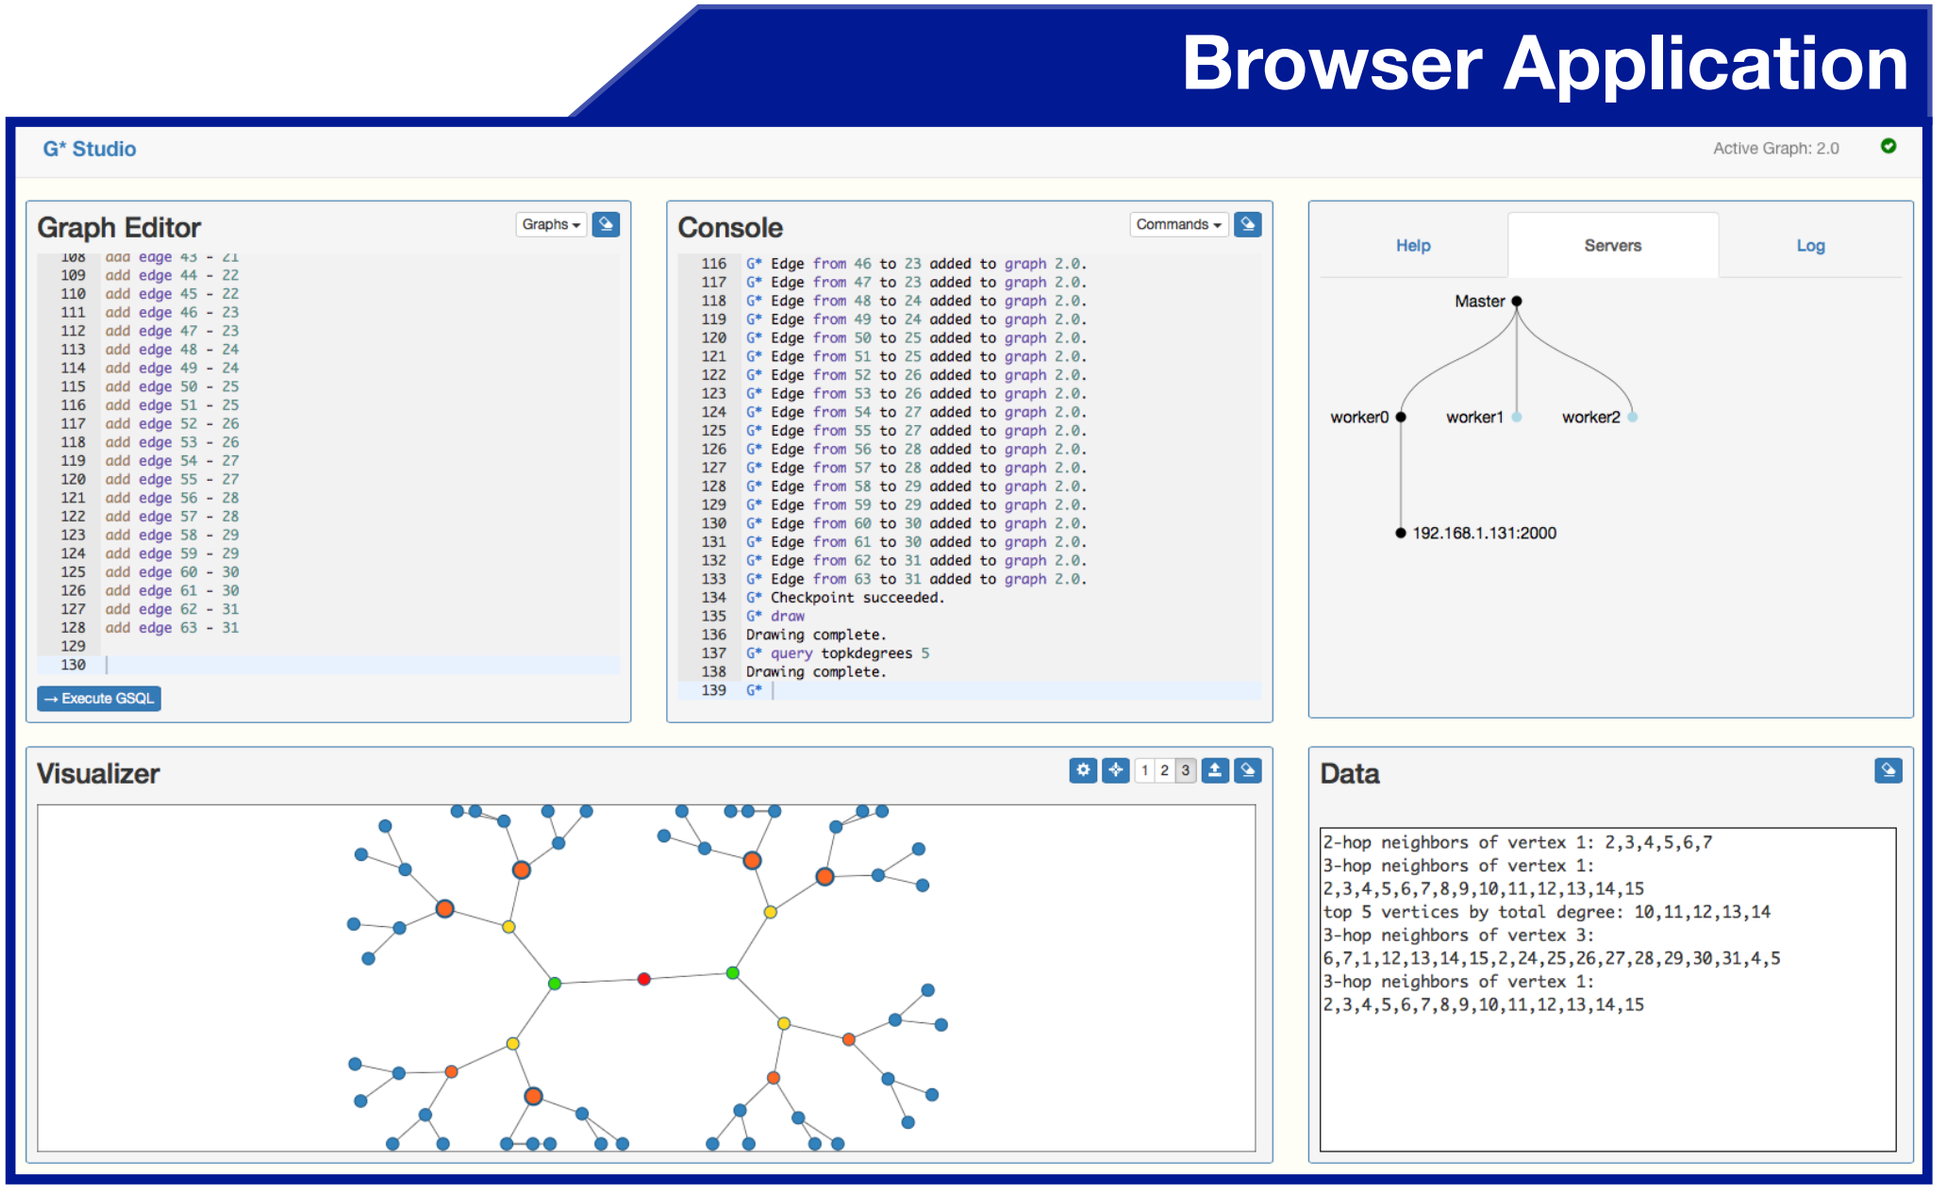

Evolution

G* The Dynamic Graph Database

Graph Editor

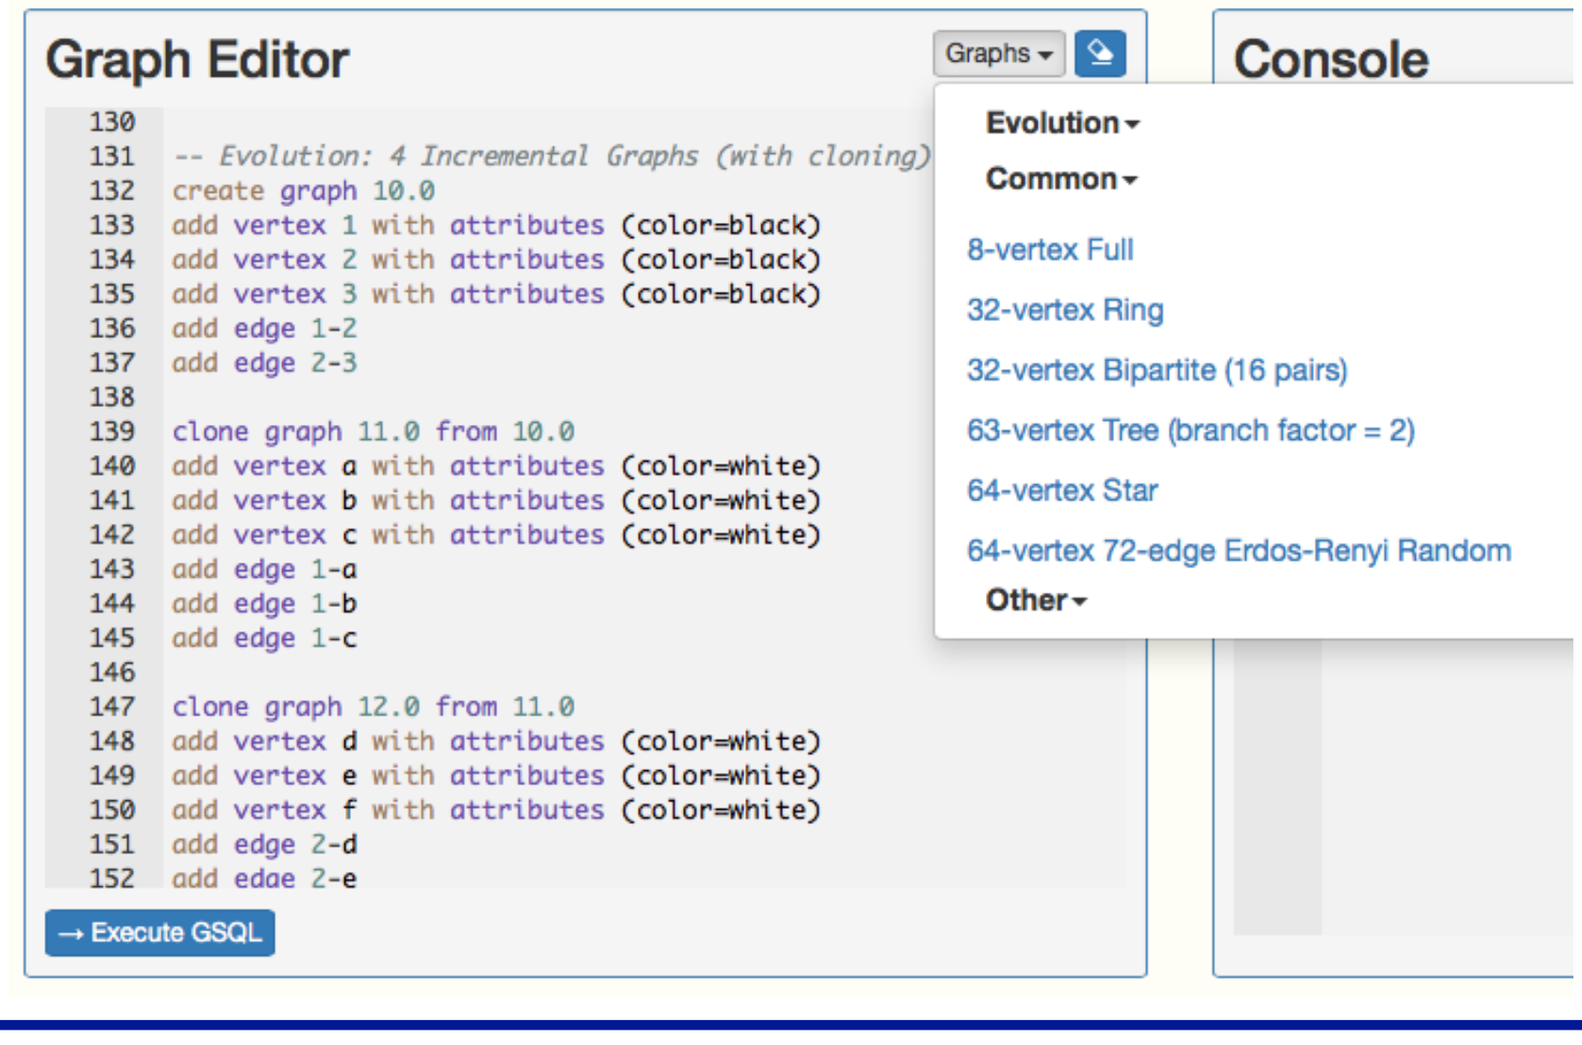

Evolution

G* The Dynamic Graph Database

Interactive Console

Console

```
  5  Graph 2.0 :
  6    Vertices: 63
  7    Edges   : 62
  8  Graph 1.0 :
  9    Vertices: 4
 10    Edges   : 2
 11  Graph 0.0 :
 12    Vertices: 2
 13    Edges   : 1
 14  G* create graph 4
 15  New graph 4.0 was created.
 16  G* add vertex Kirk
 17  Vertex Kirk added to graph 4.0.
 18  G* add vertex Spock
 19  Vertex Spock added to graph 4.0.
 20  G* add vertex McCoy
 21  Vertex McCoy added to graph 4.0.
 22  G* add edge Kirk-Spock
 23  Edge from Kirk to Spock added to graph 4.0.
 24  G* add edge Kirk-McCoy
 25  Edge from Kirk to McCoy added to graph 4.0.
 26  G* draw
 27  Drawing complete.
 28  G*
```

Commands

Information

Queries

Degree Distribution

Top-k vertices by degree

Top-k vertices with the largest change in degree over consecutive graph snapshot pairs

Evolution

G* The Dynamic Graph Database

Visualizer and Data

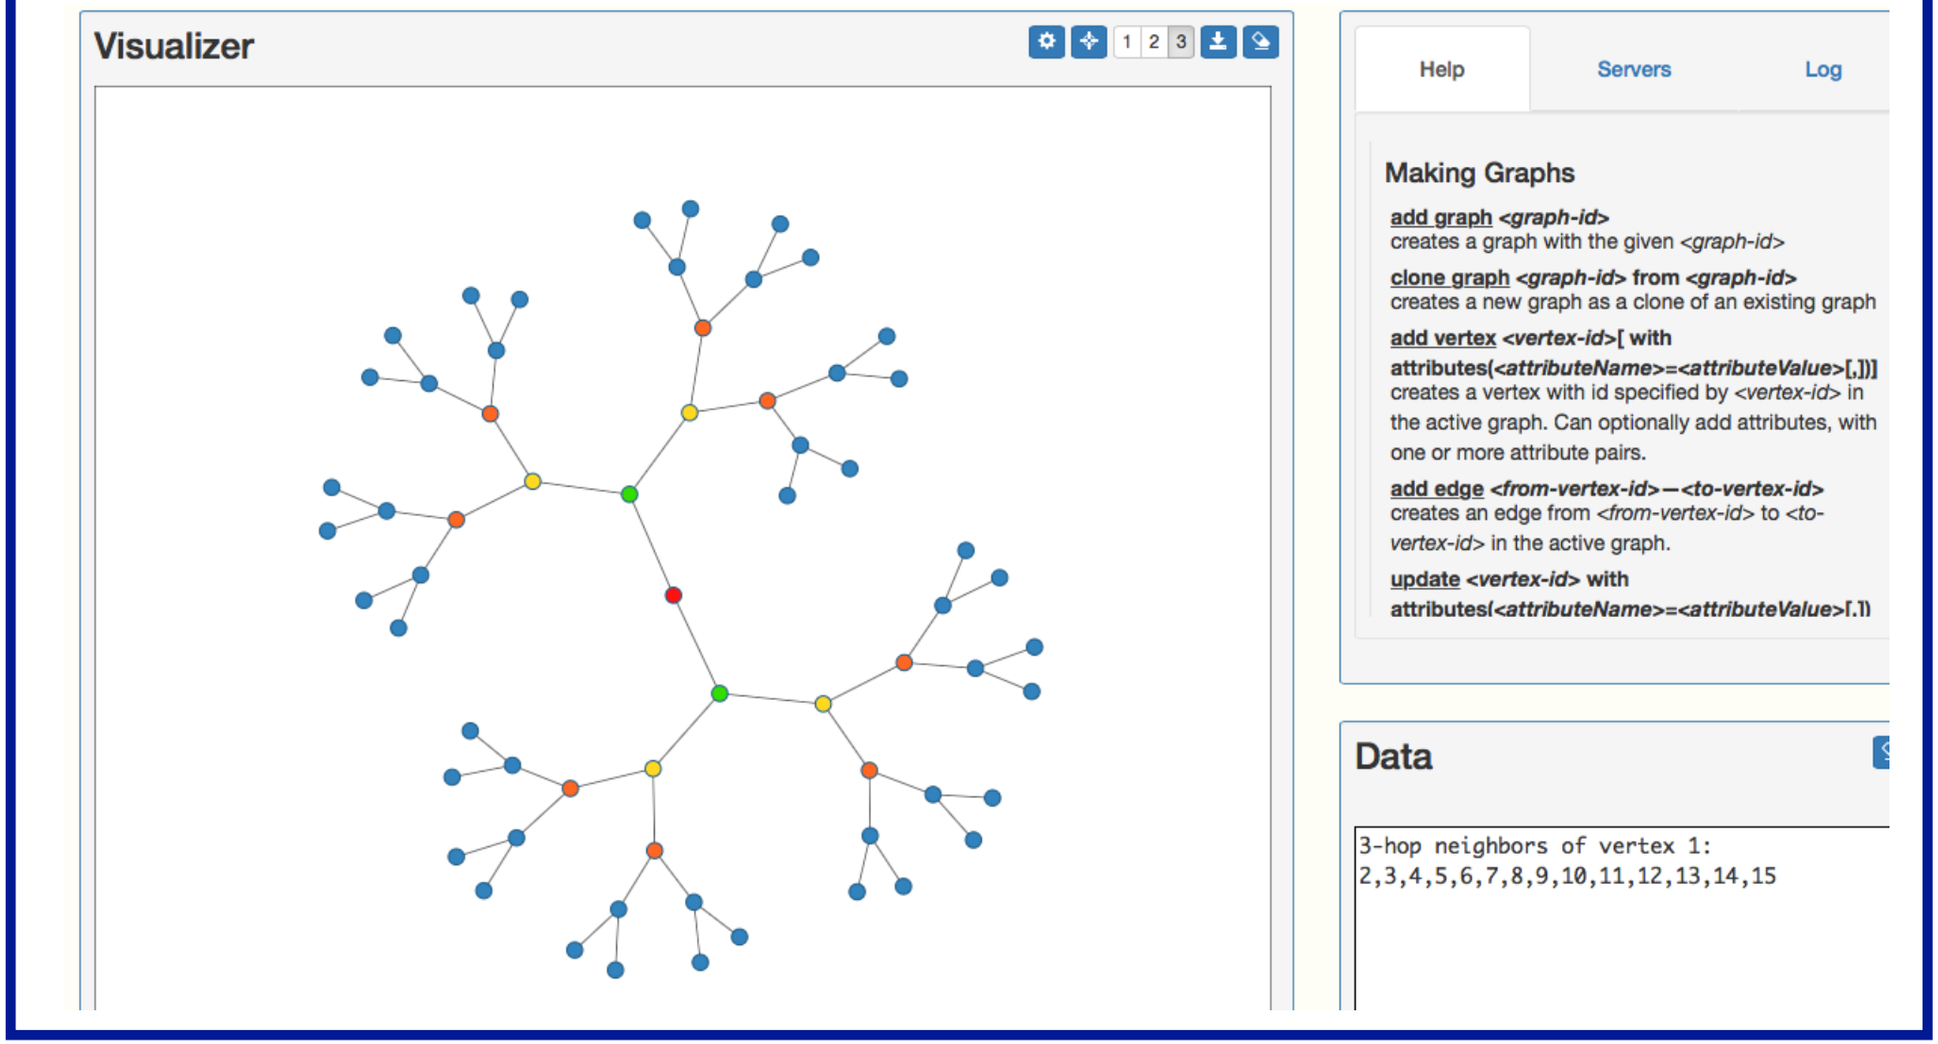

Evolution

G* The Dynamic Graph Database

Top-k Query

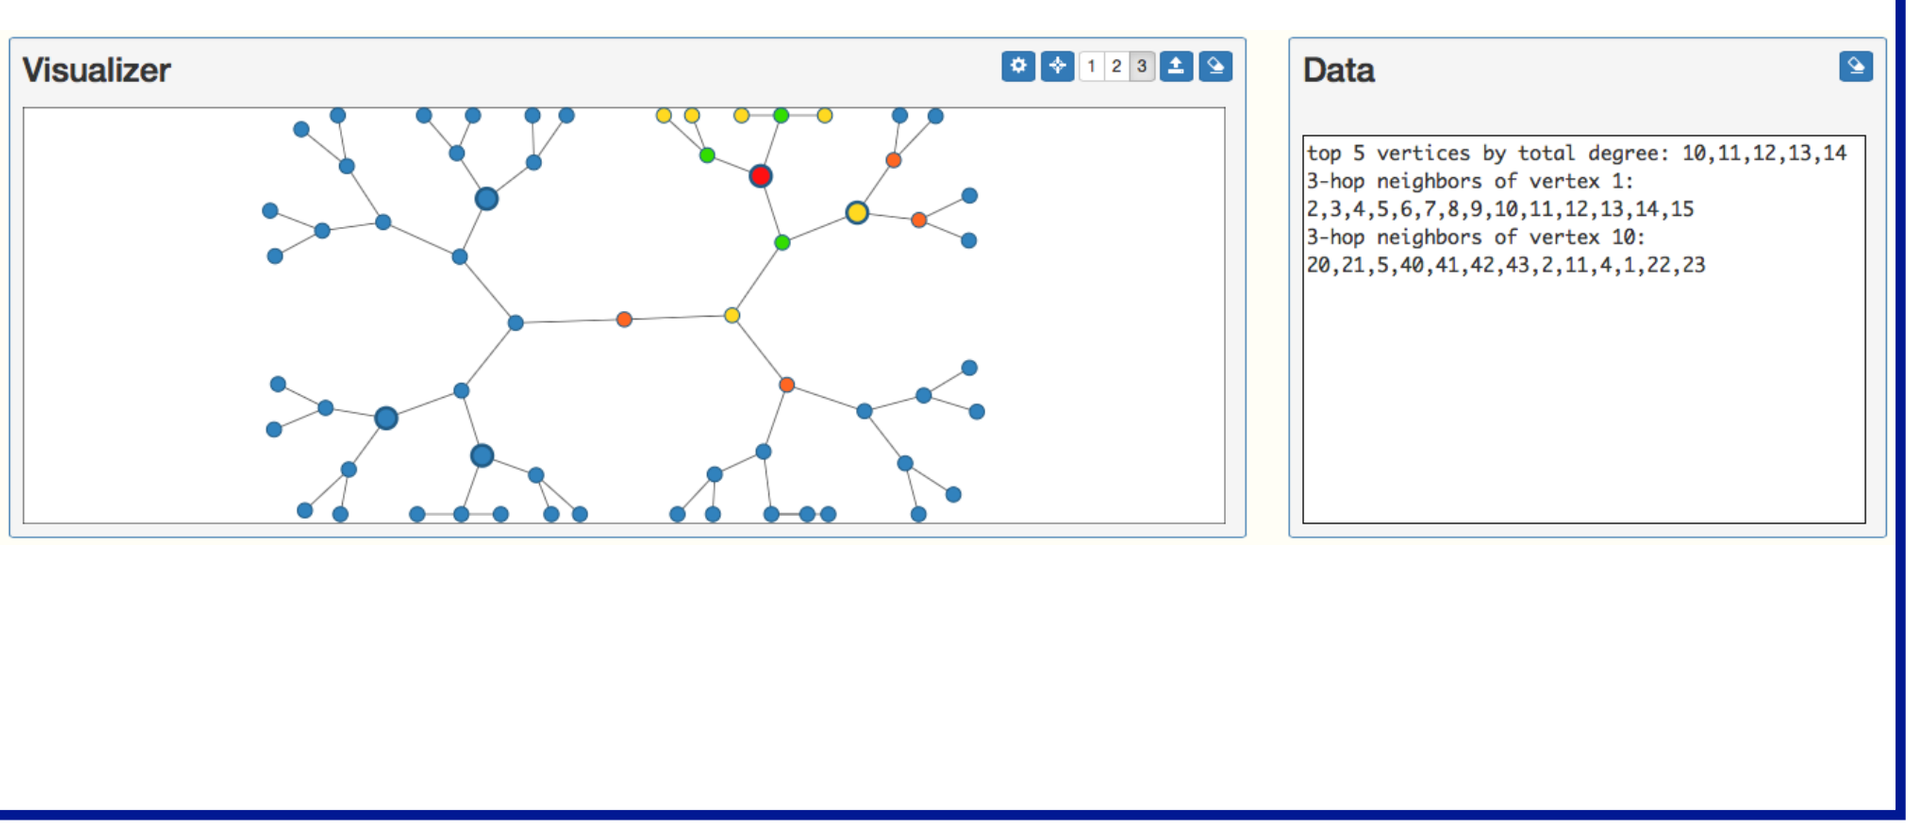

Evolution

G* The Dynamic Graph Database

Degree Distribution Query

Erdős-Rényi random graph

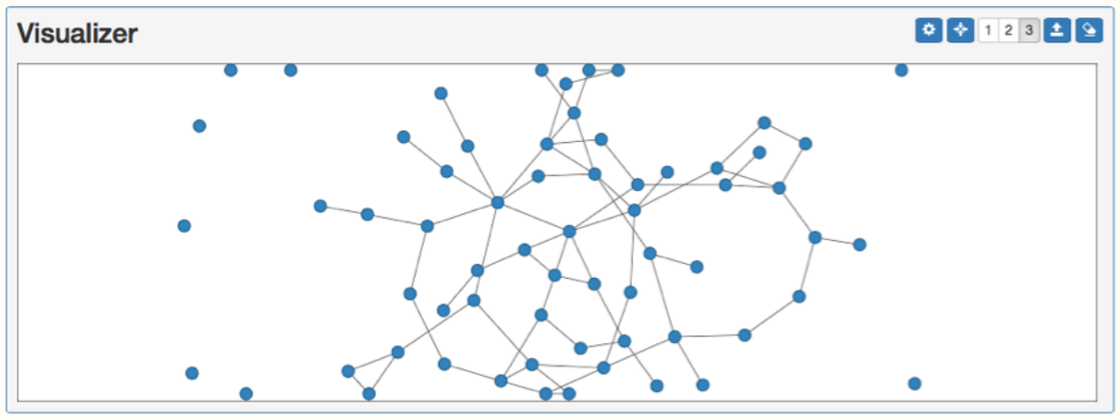

Degree Distribution

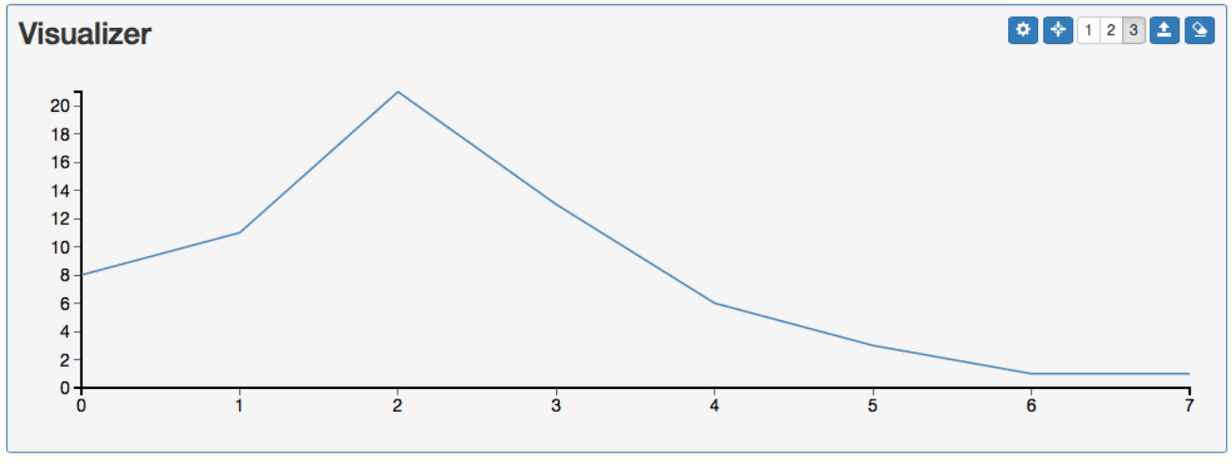

Data

```
degree distribution:
total_degree , count
0 , 8
1 , 11
2 , 21
3 , 13
4 , 6
5 , 3
6 , 1
7 , 1
```

Evolution

G* The Dynamic Graph Database

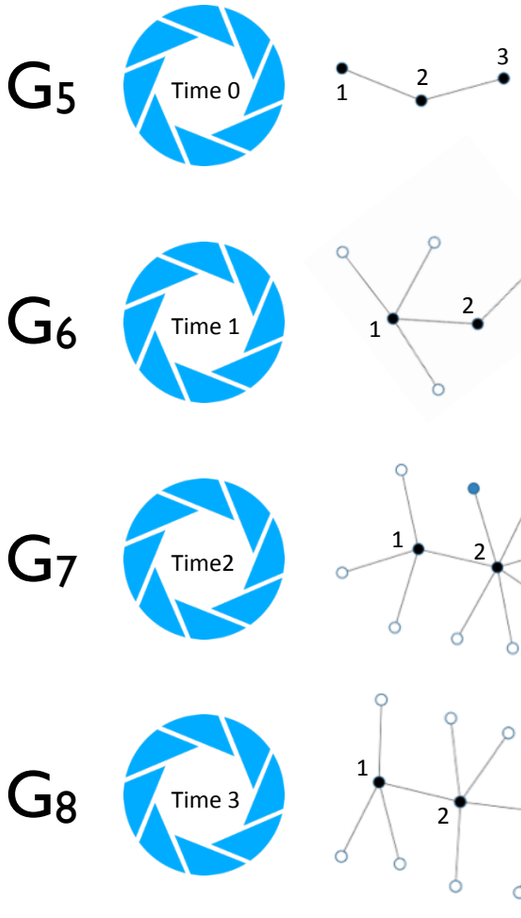

Data

```
top 20 vertices with the largest change in
degree over consecutive graph snapshot pairs
from 6 to 8:
shapshotPairs , vertexID , change
5->6 ,              1 ,       +3
6->7 ,              2 ,       +5
7->8 ,              3 ,       +3
5->6 ,              2 ,        0
5->6 ,              3 ,        0
6->7 ,              1 ,        0
6->7 ,              3 ,        0
6->7 ,              a ,        0
 .
 .
 .
7->8 ,              2 ,       -2
```

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- relational databases
- object stores
- object-relational databases (Third Manifesto?)
- graph databases

What should we concentrate on?
Where should we spend our time?

Evolution

Consider the evolution of Data Management

- stone tablets
- punched cards
- flat files on tape
- hierarchical databases on DASD
- network databases on disk
- **relational databases**
- object stores
- object-relational databases (Third Manifesto?)
- **graph databases**

We will spend most of our time on the Relational model and relational databases. And a little time on graphs.