

# Introduction to Data Management

## CSE 344

Lecture 24

Parallel Databases Wrap-up

# Announcement

- HW6:
  - Use “PARALLEL XXX”
  - Do turn in the result in Problem 3 ( $\approx 16\text{MB}$ )
  - Problem 4: extra credit but highly recommended!
- Wednesday: guest lecture Prof. Balazinska
- Thursday:
  - Reading assignment Sec.2.3.3-2.3.9 from [Mining of Massive Datasets](#), Rajaraman and Ullman
- Friday: Final Review with Paris Koutris


# Anatomy of a Query Execution

- Running problem #4
- 20 nodes = 1 master + 19 workers
- Using PARALLEL 50
- Let's see what happened


1h 16min

Only 19 reducers active, out of 50. Why?

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	33.17%	15816	10549	38	5229	0	0 / 0
reduce	4.17%	50	31	19	0	0	0 / 0



Map Completion Graph - [close](#)




3h 50min


Some errors start to occur. Watch this...

Completed. Sorting, and the rest of Reduce may proceed now

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	100.00%	15816	0	0	15816	0	0 / 18
reduce	32.42%	50	31	19	0	0	0 / 0




Map Completion Graph - [close](#)



3h 51min


Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	100.00%	15816	0	0	15816	0	0 / 18
reduce	37.72%	50	19	22	9	0	0 / 0

Completion Graph - [close](#)




3h 52min

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	100.00%	15816	0	0	15816	0	0 / 18
reduce	42.35%	50	11	20	19	0	0 / 0



Reduce Completion Graph - [close](#)


Some of the 19 reducers have finished...



...Next Batch of Reducers started

Reduce Completion Graph - [close](#)


Next Batch of 19 reducers



4h 18min

Several servers failed: "fetch error".  
Their map tasks need to be rerun. All reducers are waiting....

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	99.88%	15816	2638	30	13148	0	15 / 3337
reduce	48.42%	50	15	16	19	0	0 / 0




uce Completion Graph - [close](#)




7h 10min

Mappers finished,  
reducers resumed.

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	100.00%	15816	0	0	15816	0	26 / 5968
reduce	94.15%	50	0	6	44	0	0 / 8



uce Completion Graph - [close](#)




7h 20min

Success! 7hrs, 20mins.

Hadoop job\_201203041905\_0001 on ip-10-203-30-146

User: hadoop  
Job Name: PigLatin:DefaultJobName  
Job File:  
[hdfs://10.203.30.146:9000/mnt/var/lib/hadoop/tmp/mapred/staging/hadoop/.staging/job\\_201203041905\\_0001/job.xml](hdfs://10.203.30.146:9000/mnt/var/lib/hadoop/tmp/mapred/staging/hadoop/.staging/job_201203041905_0001/job.xml)  
Submit Host: ip-10-203-30-146.ec2.internal  
Submit Host Address: 10.203.30.146  
Job-ACLs: All users are allowed  
Job Setup: [Successful](#)  
Status: Succeeded  
Started at: Sun Mar 04 19:08:29 UTC 2012  
Finished at: Mon Mar 05 02:28:39 UTC 2012  
Finished in: 7hrs, 20mins, 10sec  
Job Cleanup: [Successful](#)  
Black-listed TaskTrackers: 3

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
<a href="#">map</a>	100.00% <div><div></div></div>	15816	0	0	<a href="#">15816</a>	0	<a href="#">26 / 5968</a>
<a href="#">reduce</a>	100.00% <div><div></div></div>	50	0	0	<a href="#">50</a>	0	0 / <a href="#">14</a>



# Hash Join

```
Users = load `users` as (name, age);  
Pages = load `pages` as (user, url);  
Jnd = join Users by name, Pages by user;
```



Pages

The diagram illustrates the Hash Join process. It consists of two vertical blue rectangles. The left rectangle is labeled 'Pages' and the right rectangle is labeled 'Users'. Both rectangles have a blue gradient, with the top being a lighter shade and the bottom being a darker shade. The labels 'Pages' and 'Users' are centered near the top of their respective rectangles.

Users



# Hash Join

```
Users = load 'users' as (name, age);  
Pages = load 'pages' as (user, url);  
Jnd = join Users by name, Pages by user;
```

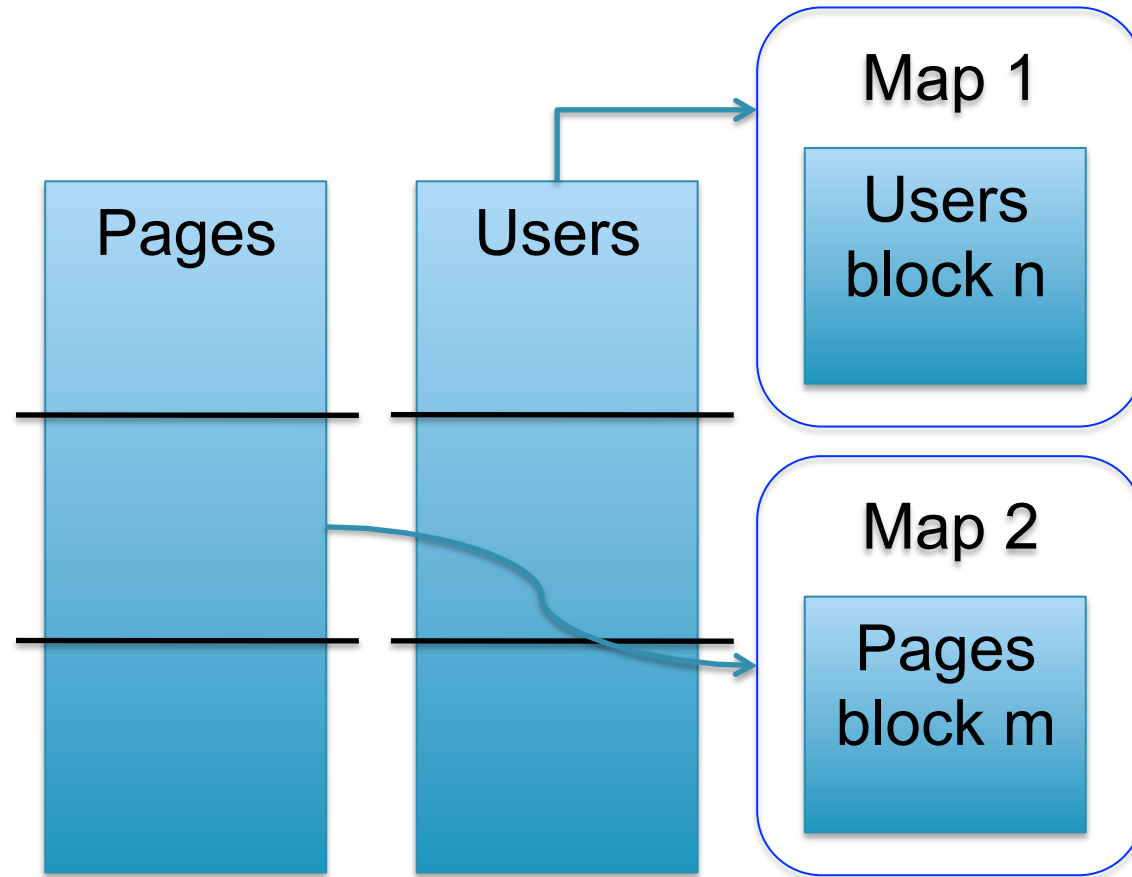
Pages



Users

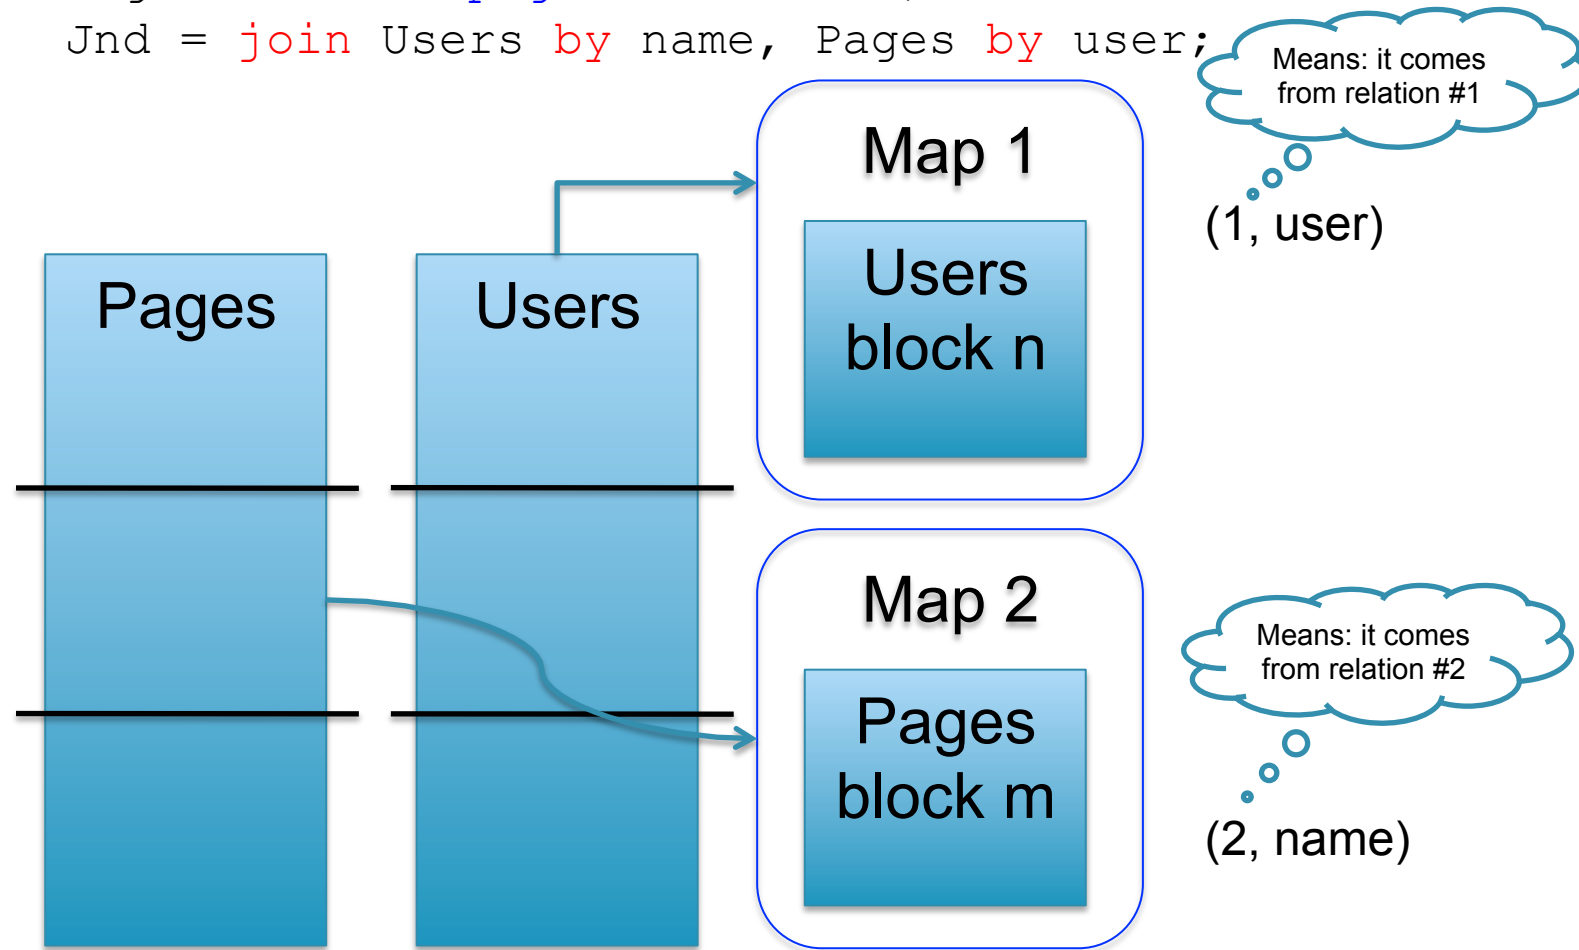
# Hash Join

```
Users = load 'users' as (name, age);  
Pages = load 'pages' as (user, url);  
Jnd = join Users by name, Pages by user;
```



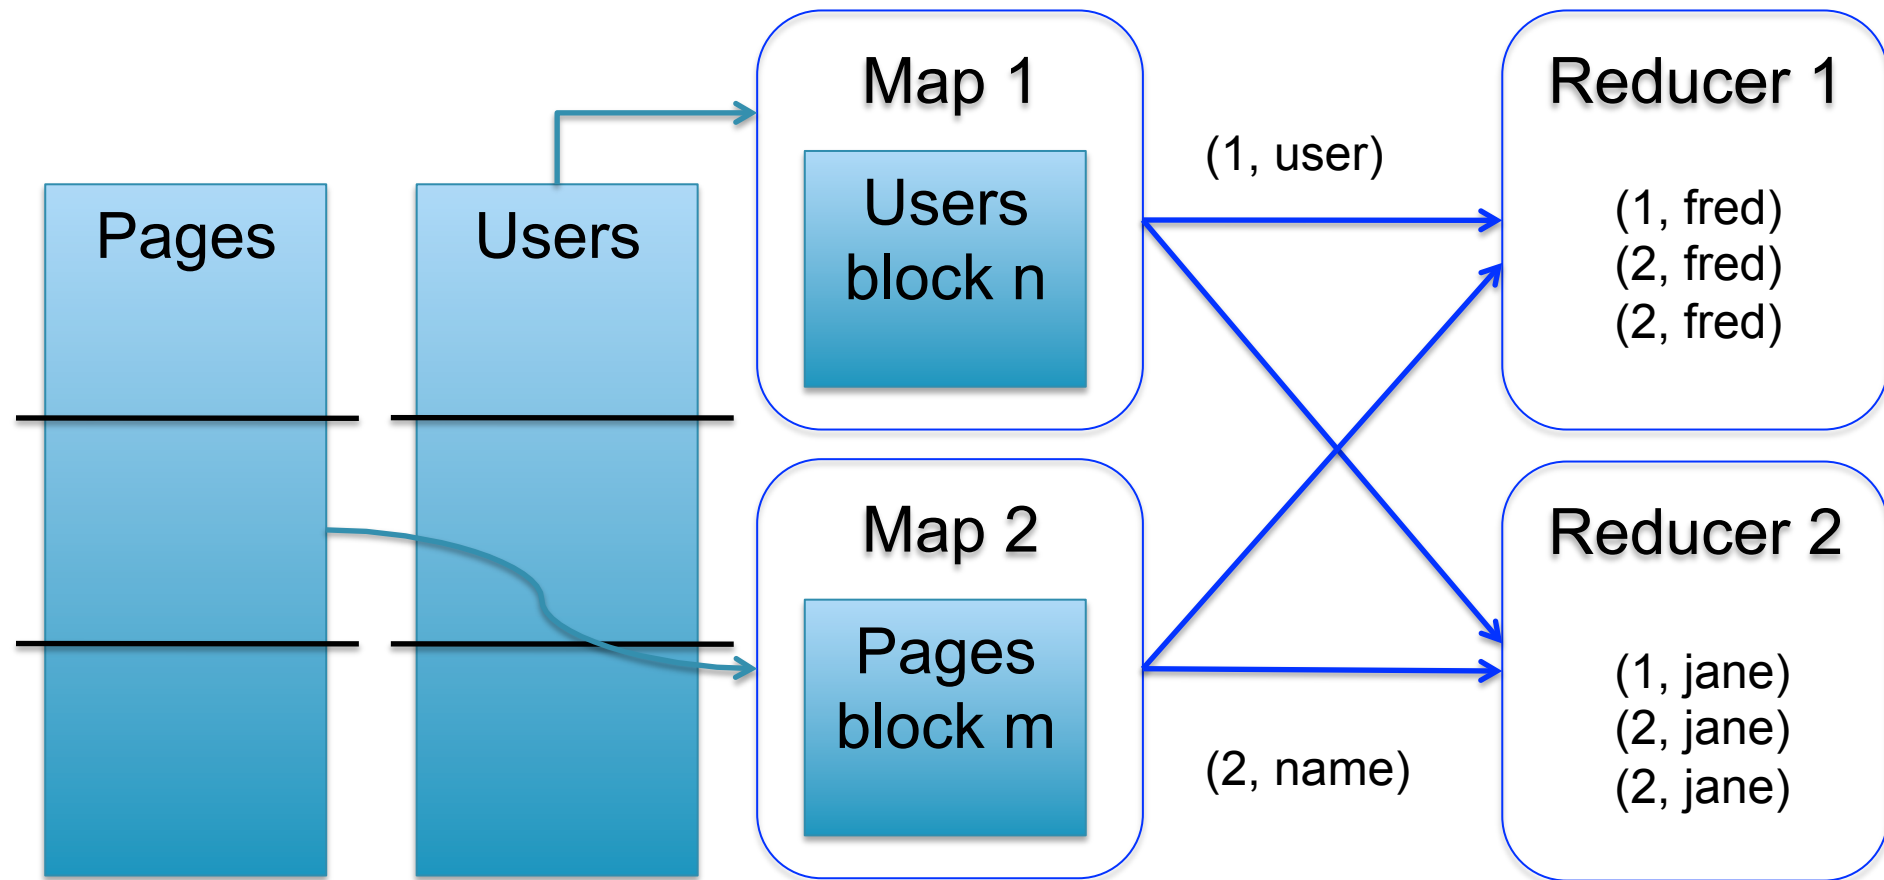
# Hash Join

```
Users = load 'users' as (name, age);  
Pages = load 'pages' as (user, url);  
Jnd = join Users by name, Pages by user;
```



# Hash Join

```
Users = load 'users' as (name, age);  
Pages = load 'pages' as (user, url);  
Jnd = join Users by name, Pages by user;
```



# Hash Join

```
Users = load 'users' as (name, age);  
Pages = load 'pages' as (user, url);  
Jnd = join Users by name, Pages by user;
```

```
map(String usr, String value):  
    // usr: either Users.name or Pages.user  
    // value.relation is either 'Users' or 'Pages'  
    if value.relation='Users':  
        EmitIntermediate(usr, (1, value));  
    else  
        EmitIntermediate(usr, (2, value));
```

```
reduce(String usr, Iterator values):  
    Users = empty; Pages = empty;  
    for each v in values:  
        if v.type = 1: Users.insert(v)  
        else Pages.insert(v);  
    for v1 in Users, for v2 in Pages  
        Emit(usr, v1,v2);
```

# Broadcast Join

```
Users = load `users` as (name, age);  
Pages = load `pages` as (user, url);  
Jnd = join Pages by user, Users by name using "replicated";
```



Pages

Users

# Broadcast Join

```
Users = load `users` as (name, age);  
Pages = load `pages` as (user, url);  
Jnd = join Pages by user, Users by name using "replicated";
```



Pages

Users

# Broadcast Join

```
Users = load `users` as (name, age);  
Pages = load `pages` as (user, url);  
Jnd = join Pages by user, Users by name using "replicated";
```

Pages



Users

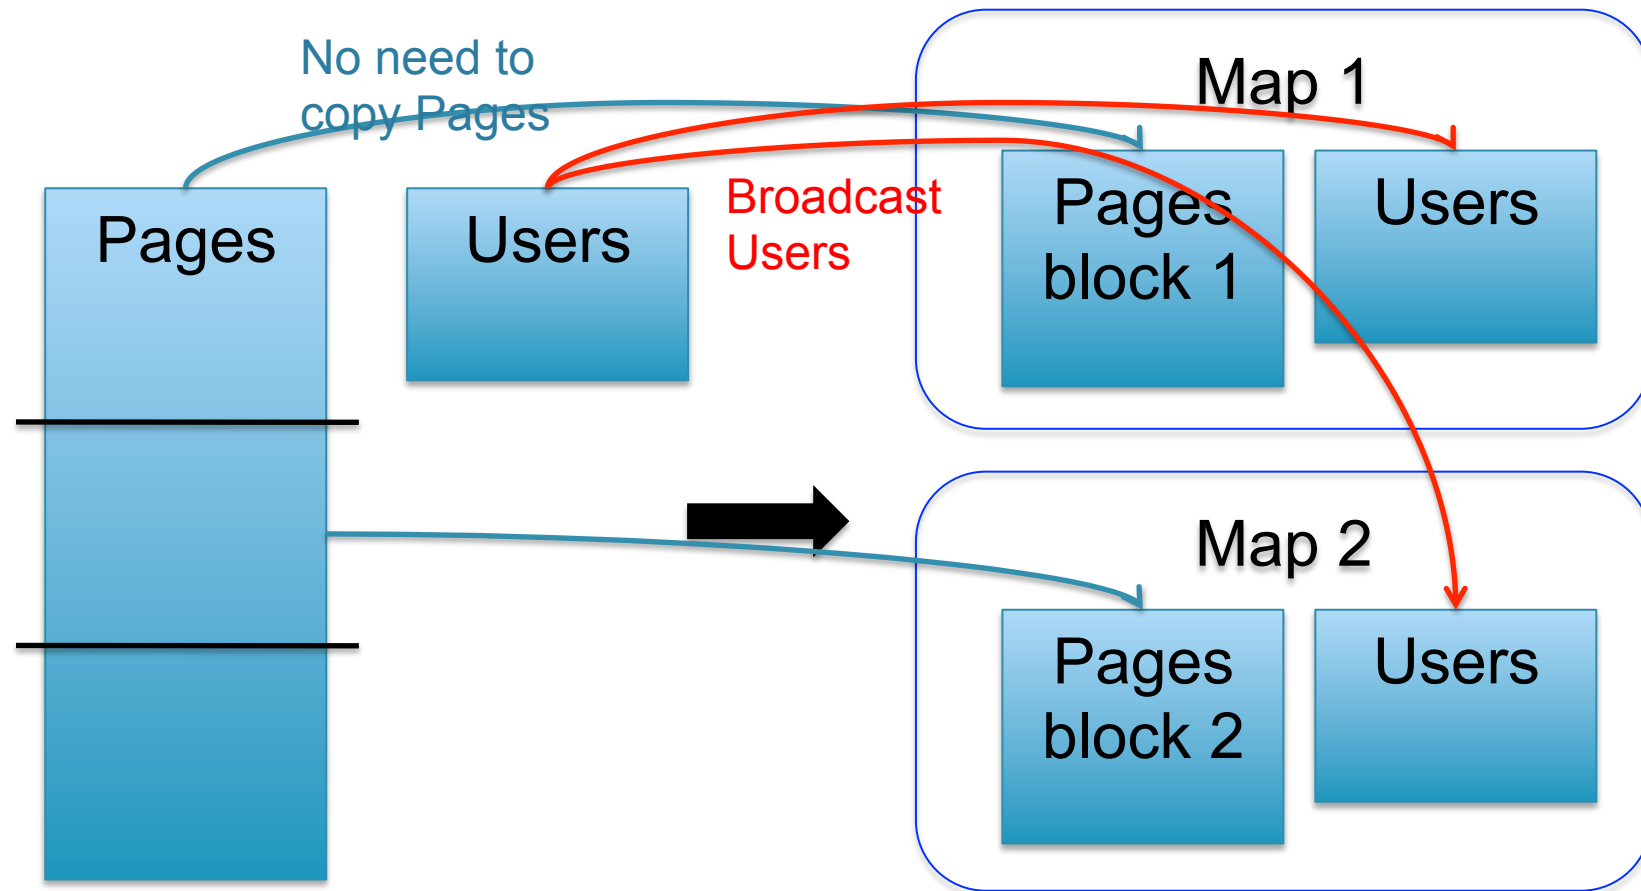
Map 1

Map 2



# Broadcast Join

```
Users = load `users` as (name, age);  
Pages = load `pages` as (user, url);  
Jnd = join Pages by user, Users by name using "replicated";
```



# Parallel DBs v.s. Map-Reduce

## **Parallel DB**

- Plusses
- Minuses

## **Map-Reduce**

- Minuses
- Plusses

# Parallel DBs v.s. Map-Reduce

## Parallel DB

- **Plusses**
  - Efficient binary format
  - Indexes, physical tuning
  - Cost-based optimization
- **Minuses**
  - Difficult to import data
  - Lots of baggage: logging, transactions

## Map-Reduce

- **Minuses**
  - Lots of time spent parsing!
  - Text files
  - “Optimizers is between your eyes and your keyboard”
- **Plusses**
  - Any data
  - Lightweight, easy to speedup