Case Study 4: Collaborative Filtering

Graph-Parallel Problems

Synchronous v. Asynchronous Computation

Machine Learning/Statistics for Big Data
CSE599C1/STAT592, University of Washington

Carlos Guestrin

March 12th, 2013

Needless to Say, We Need Machine Learning for Big Data

flickr
6 Billion
Flickr Photos

28 Million
Wikipedia Pages

facebook
1 Billion
Facebook Users

YouTube
72 Hours a Minute
YouTube

The New York Times
SundayReview

WORLD U.S. N.Y. / REGION BUSINESS TEC

NEWS ANALYSIS
The Age of Big Data
By STEVE LOHR
Published: February 11, 2012

"… data a new class of economic asset, like currency or gold."

CPUs Stopped Getting Faster…

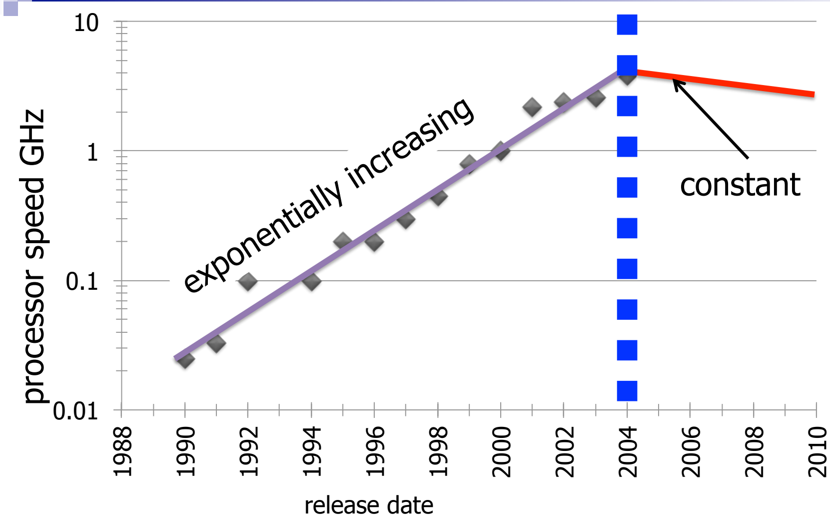

ML in the Context of Parallel Architectures

GPUs

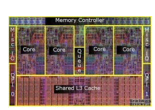

Multicore

Clusters

amazon web services™

Clouds

Supercomputers

- But scalable ML in these systems is hard, especially in terms of:
  1. Programmability
  2. Data distribution
  3. Failures

Move Towards Higher-Level Abstraction

- Distributed computing challenges are hard and annoying!
  1. Programmability
  2. Data distribution
  3. Failures
- High-level abstractions try to simplify distributed programming by hiding challenges:
  - Provide different levels of robustness to failures, optimizing data movement and communication, protect against race conditions…
  - Generally, you are still on your own WRT designing parallel algorithms
- Some common parallel abstractions:
  - Lower-level:
    - Pthreads: abstraction for distributed threads on single machine
    - MPI: abstraction for distributed communication in a cluster of computers
  - Higher-level:
    - Map-Reduce (Hadoop: open-source version): mostly data-parallel problems
    - GraphLab: for graph-structured distributed problems

Simplest Type of Parallelism: Data Parallel Problems

- You have already learned a classifier w^*
  - What's the test error? $err = \frac{1}{N_{test}} \sum_i \frac{1}{2} \left| y^{(i)} - \text{sign}(w^* \cdot x^{(i)}) \right|$
- You have 10B labeled documents and 1000 machines

locally compute error on subset of data

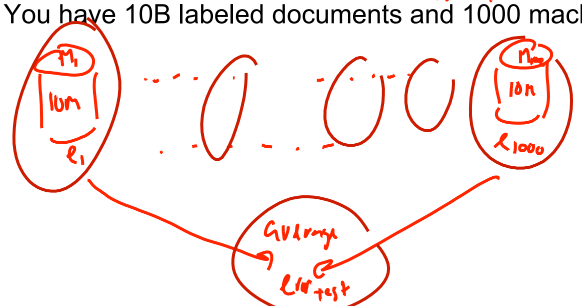

- Problems that can be broken into independent subproblems are called data-parallel (or embarrassingly parallel)
- ~~Map-Reduce~~ is a great tool for this…
  - Focus of today's lecture
  - but first a simple example

Data Parallelism (MapReduce)

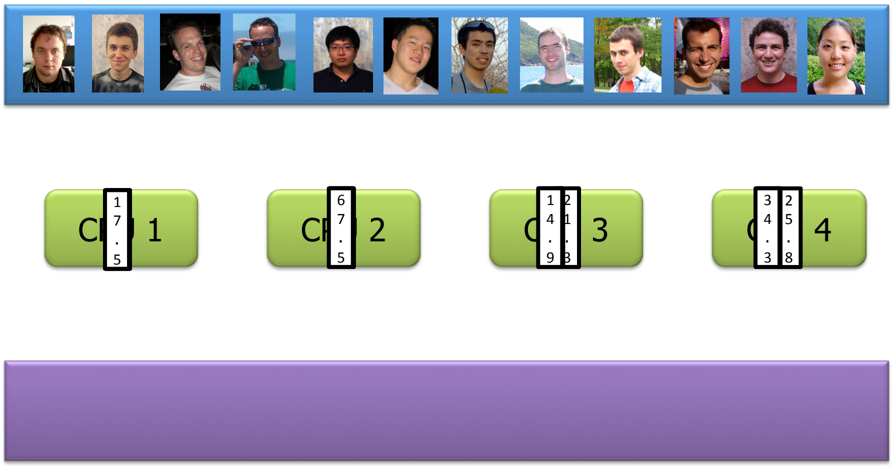

*Solve a huge number of **independent** subproblems, e.g., extract features in images*

Map-Reduce Abstraction

- Map: Transforms a data element
  - Data-parallel over elements, e.g., documents
  - Generate (key,value) pairs
    - "value" can be any data type

('UW', 17)

in this example: ('Mary',1) ('UW',1) ('Mary',1)

- Reduce: Take all values associated with a key and aggregate
  - Aggregate values for each key
  - Must be commutative-associate operation
  - Data-parallel over keys
  - Generate (key,value) pairs

reduce ('UW', [1, 17, 0, 0, 12])
emit ('UW', 30)

word count

```
map (document)
  for word in doc
    emit (word, 1)

Reduce (word, count: list(int))
  c = 0
  for i in count
    c += count[i]
  emit (word, c)
```

- Map-Reduce has long history in functional programming
  - But popularized by Google, and subsequently by open-source Hadoop implementation from Yahoo!

Map-Reduce – Execution Overview

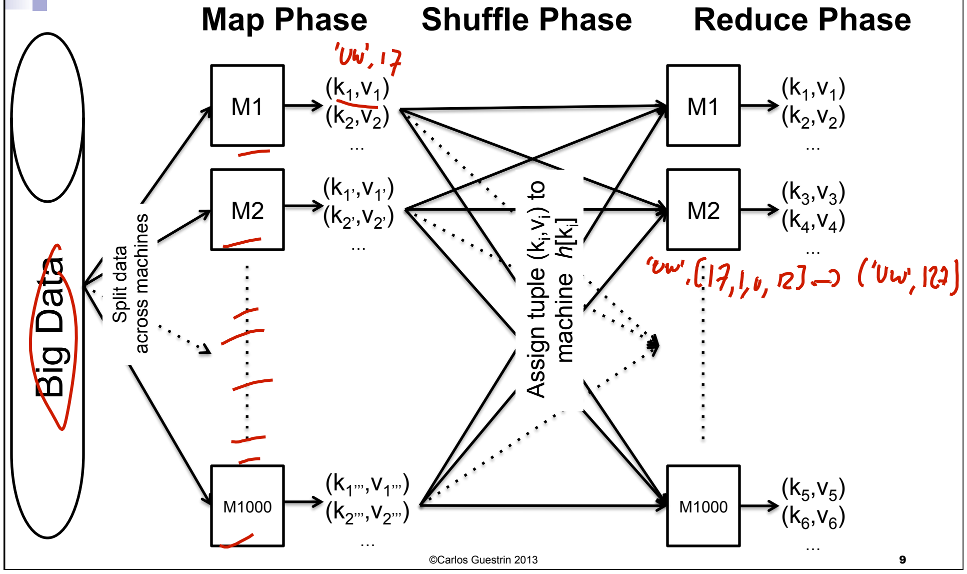

Issues with Map-Reduce Abstraction

- Often all data gets moved around cluster
  - Very bad for iterative settings
- Definition of Map & Reduce functions can be unintuitive in many apps
  - Graphs are challenging
- Computation is synchronous

SGD for Matrix Factorization in Map-Reduce?

$$\epsilon_t = L_u^{(t)} \cdot R_v^{(t)} - r_{uv}$$

$$\begin{bmatrix} L_u^{(t+1)} \\ R_v^{(t+1)} \end{bmatrix} \leftarrow \begin{bmatrix} (1 - \eta_t \lambda_u) L_u^{(t)} - \eta_t \epsilon_t R_v^{(t)} \\ (1 - \eta_t \lambda_v) R_v^{(t)} - \eta_t \epsilon_t L_u^{(t)} \end{bmatrix}$$

- Map and Reduce functions???
- Map-Reduce:
  - Data-parallel over all mappers
  - Data-parallel over reducers with same key
- Here, one update at a time!

Matrix Factorization as a Graph

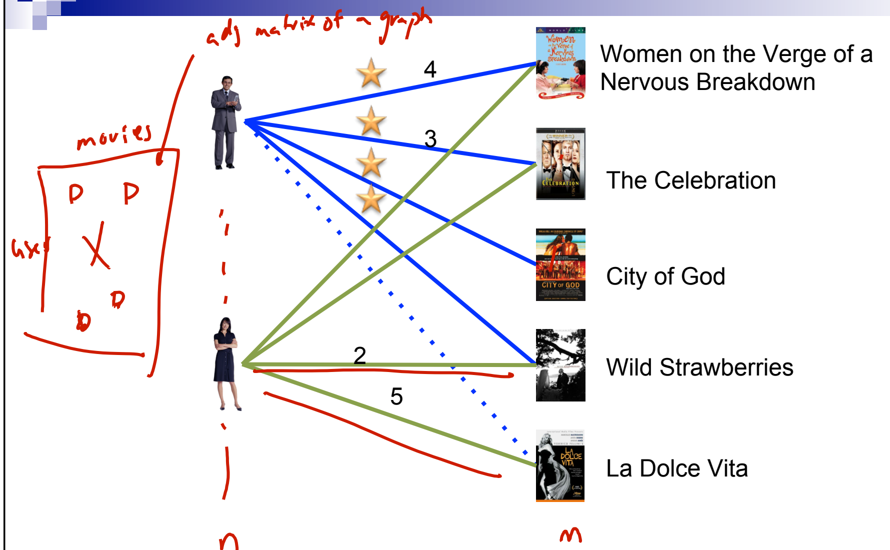

Flashback to 1998

Why?

First Google advantage:
a **Graph Algorithm** & a **System to Support** it!

Social Media **Science** **Advertising** **Web**

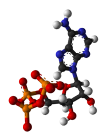

- **Graphs** encode the **relationships** between:

People Products Ideas
Facts Interests

- **Big: 100 billions** of **vertices** and **edges** and rich metadata
  - Facebook (10/2012): 1B users, 144B friendships
  - Twitter (2011): 15B follower edges

Facebook Graph

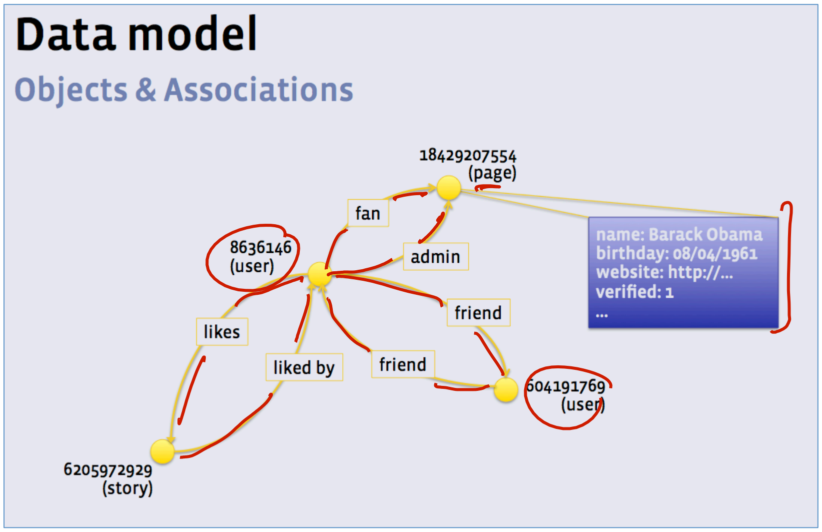

Slide from Facebook Engineering presentation

Label a Face and Propagate

Pairwise similarity not enough…

©Carlos Guestrin 2013

17

Propagate Similarities & Co-occurrences for Accurate Predictions

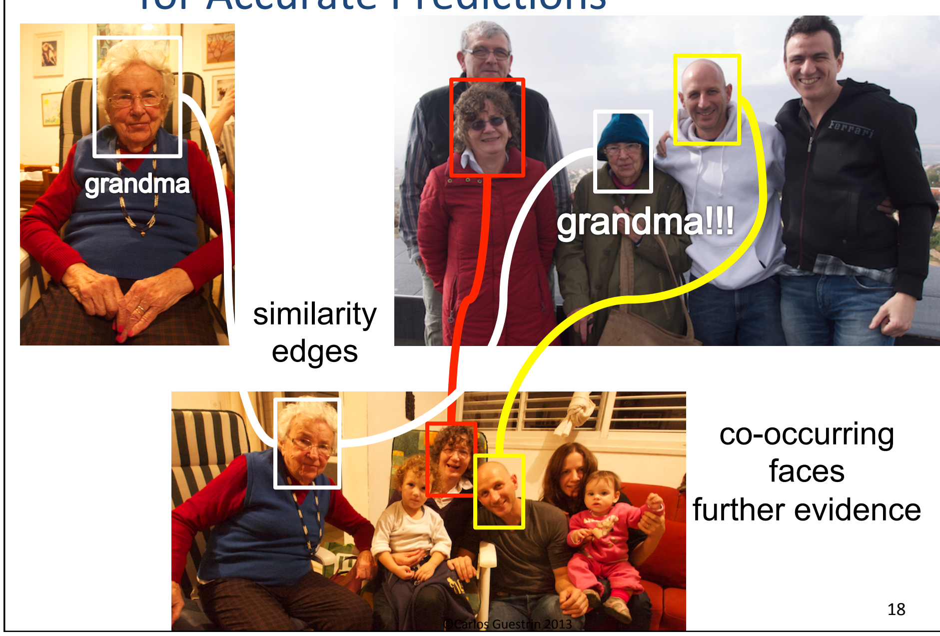

co-occurring faces further evidence

18

©Carlos Guestrin 2013

Example: *Estimate Political Bias*

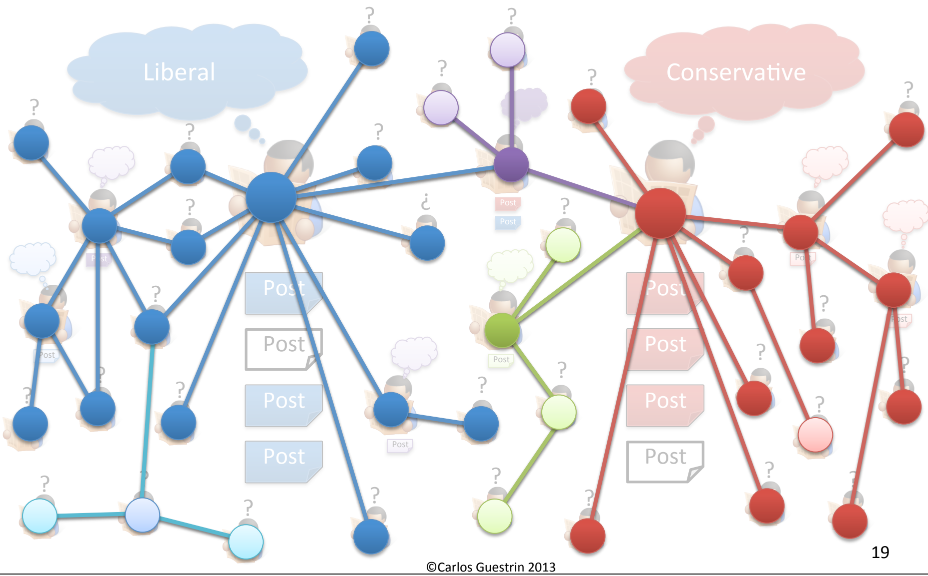

Latent Topic Modeling (LDA)

✓ docs

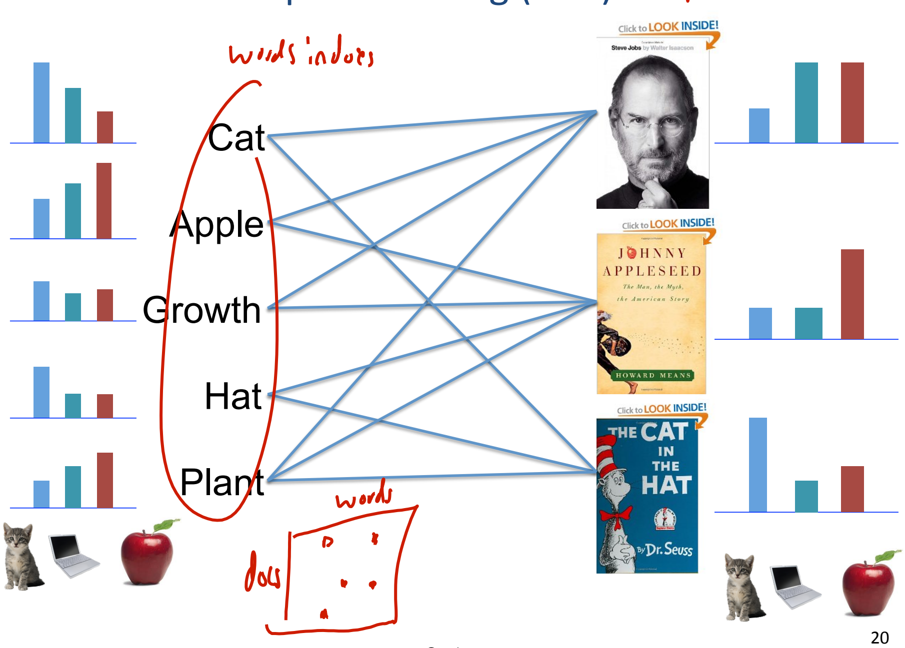

ML Tasks Beyond Data-Parallelism

Data-Parallel — Graph-Parallel

Map Reduce

Feature Extraction

Cross Validation

Computing Sufficient Statistics

Graphical Models
Gibbs Sampling
Belief Propagation
Variational Opt.

Semi-Supervised Learning
Label Propagation
CoEM

Collaborative Filtering
Tensor Factorization

Graph Analysis
PageRank
Triangle Counting

Example of a Graph-Parallel Algorithm

PageRank

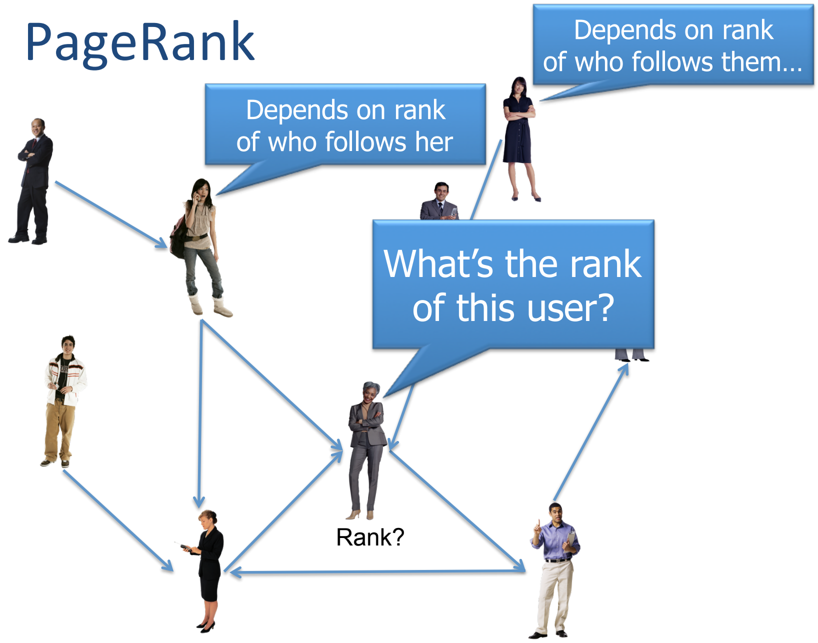

Loops in graph → Must iterate!

PageRank Iteration

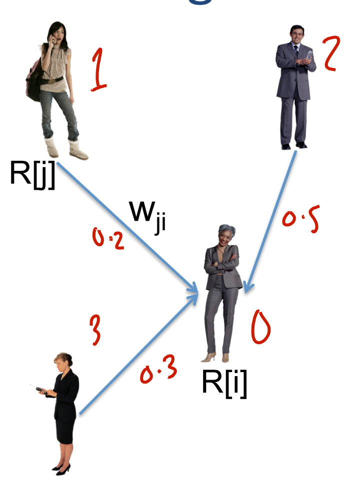

rank of user i

$$R[i] = \alpha + (1-\alpha) \sum_{(j,i)\in E} w_{ji} R[j]$$

$$R[0] = 0.15 + 0.85\left[0.2R[1] + 0.5R[2] + 0.3R[3]\right]$$

0.15

- α is the random reset probability
- w_{ji} is the prob. transitioning (similarity) from j to i

Properties of Graph Parallel Algorithms

Dependency Graph	Local Updates	Iterative Computation
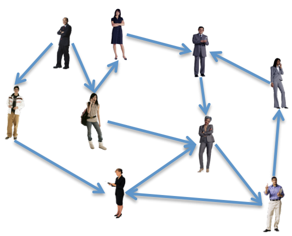	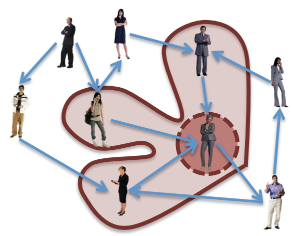	

Addressing Graph-Parallel ML

Data-Parallel ⟷ Graph-Parallel

Map Reduce	Graph-Parallel Abstraction
Feature Extraction	**Graphical Models**: Gibbs Sampling, Belief Propagation, Variational Opt.
Cross Validation	**Semi-Supervised Learning**: Label Propagation, CoEM
Computing Sufficient Statistics	**Collaborative Filtering**: Tensor Factorization
	Data-Mining: PageRank, Triangle Counting

Graph Computation:

Synchronous

v.

Asynchronous

Bulk Synchronous Parallel Model: Pregel (Giraph)

[Valiant '90]

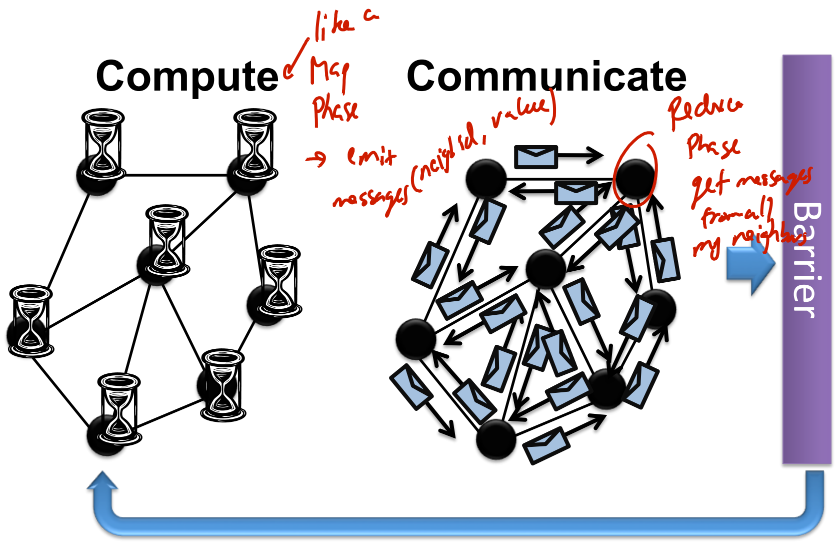

Map-Reduce – Execution Overview

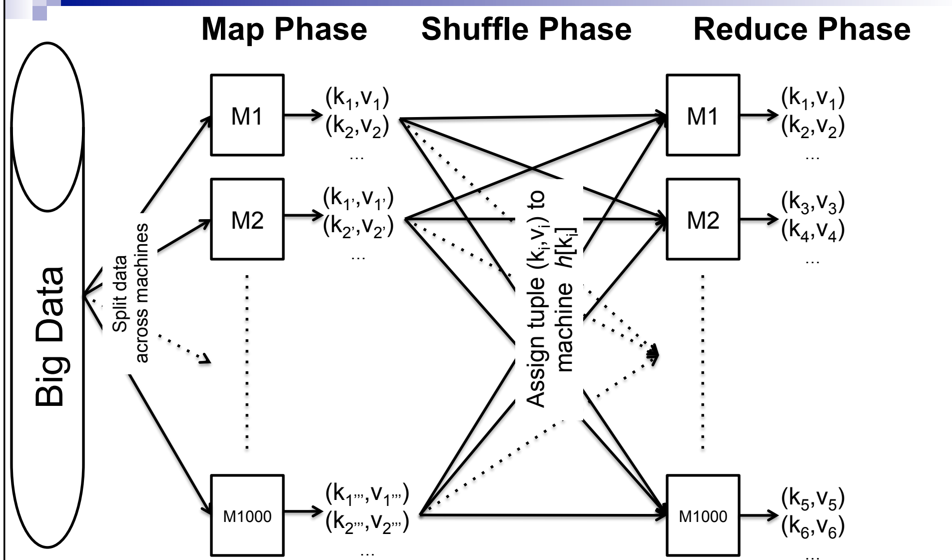

BSP – Execution Overview

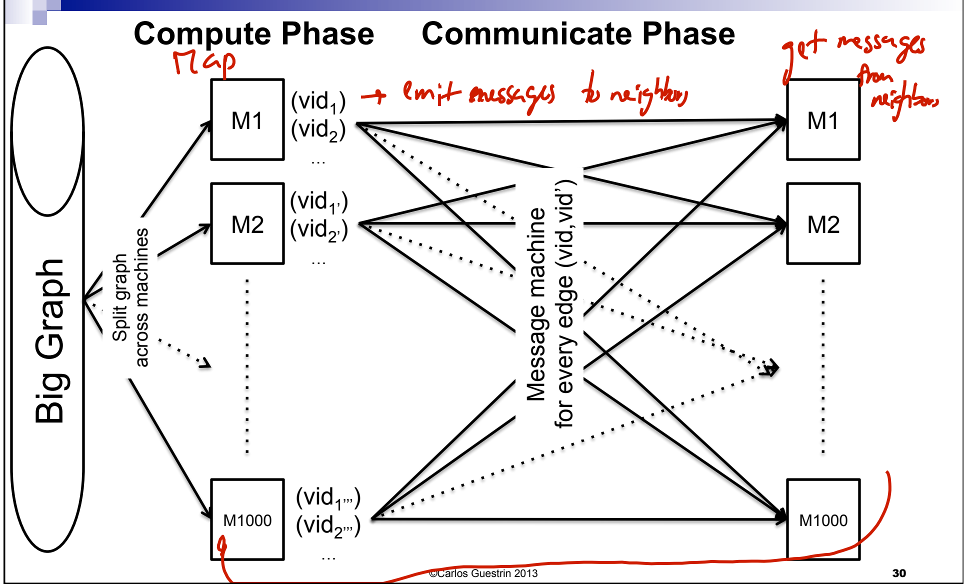

Bulk synchronous parallel model
provably inefficient
for some ML tasks

Analyzing Belief Propagation

[Gonzalez, Low, G. '09]

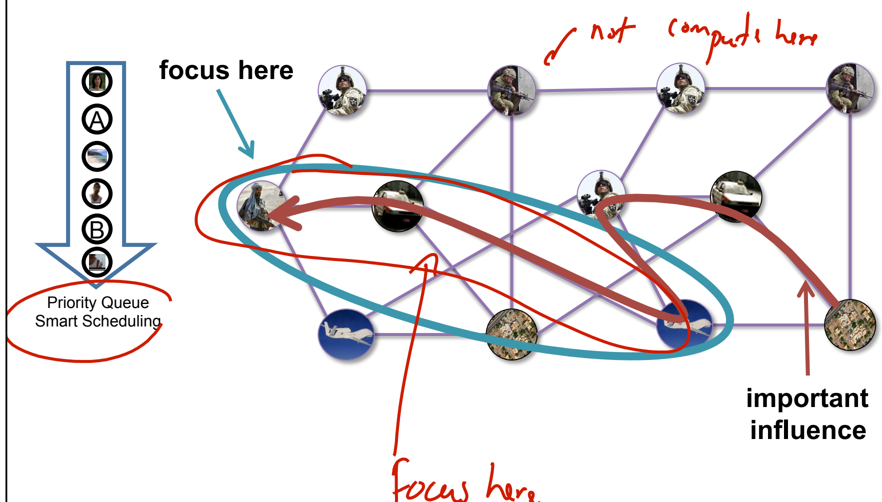

Asynchronous Belief Propagation

Challenge = Boundaries

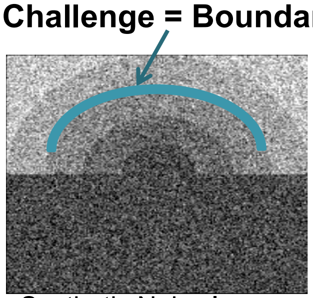

Synthetic Noisy Image

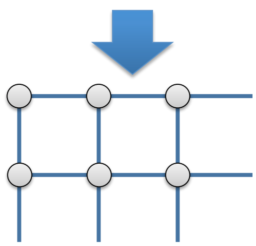

Graphical Model

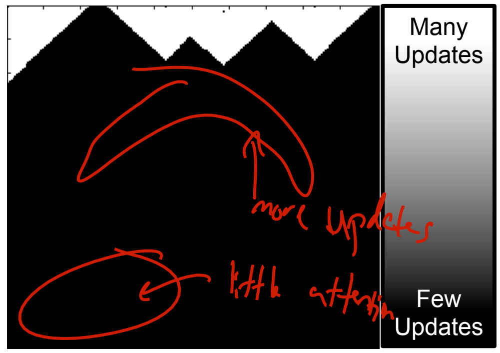

Cumulative Vertex Updates

Algorithm identifies and focuses on hidden sequential structure

BSP ML Problem: BP on real data
Synchronous Algorithms can be **Inefficient**

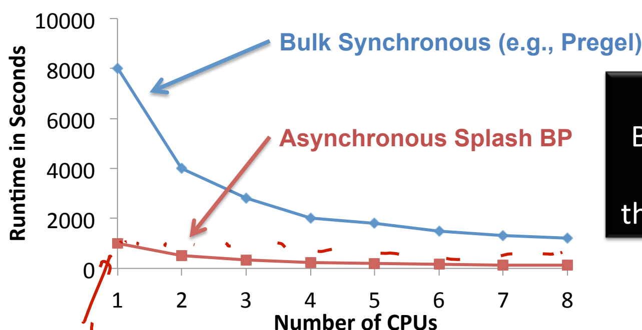

1 proc async faster than 8 proc of sync

Theorem:
Bulk Synchronous BP
O(#vertices) slower
than Asynchronous BP

Synchronous v. Asynchronous

- Bulk synchronous processing:
  - Computation in phases
    - All vertices participate in a phase
      - Though OK to say no-op
    - All messages are sent
  - Simpler to build, like Map-Reduce
    - No worries about race conditions, barrier guarantees data consistency
    - Simpler to make fault-tolerant, save data on barrier
  - Slower convergence for many ML problems
  - In matrix-land, called Jacobi Iteration
  - Implemented by Google Pregel 2010
- Asynchronous processing:
  - Vertices see latest information from neighbors
    - Most closely related to sequential execution
  - Harder to build:
    - Race conditions can happen all the time
      - Must protect against this issue
    - More complex fault tolerance
    - When are you done?
    - Must implement scheduler over vertices
  - Faster convergence for many ML problems
  - In matrix-land, called Gauss-Seidel Iteration
  - Implemented by GraphLab 2010, 2012