

Cloud Computing


# **Link Analysis in the Cloud**

Dell Zhang

Birkbeck, University of London

2018/19

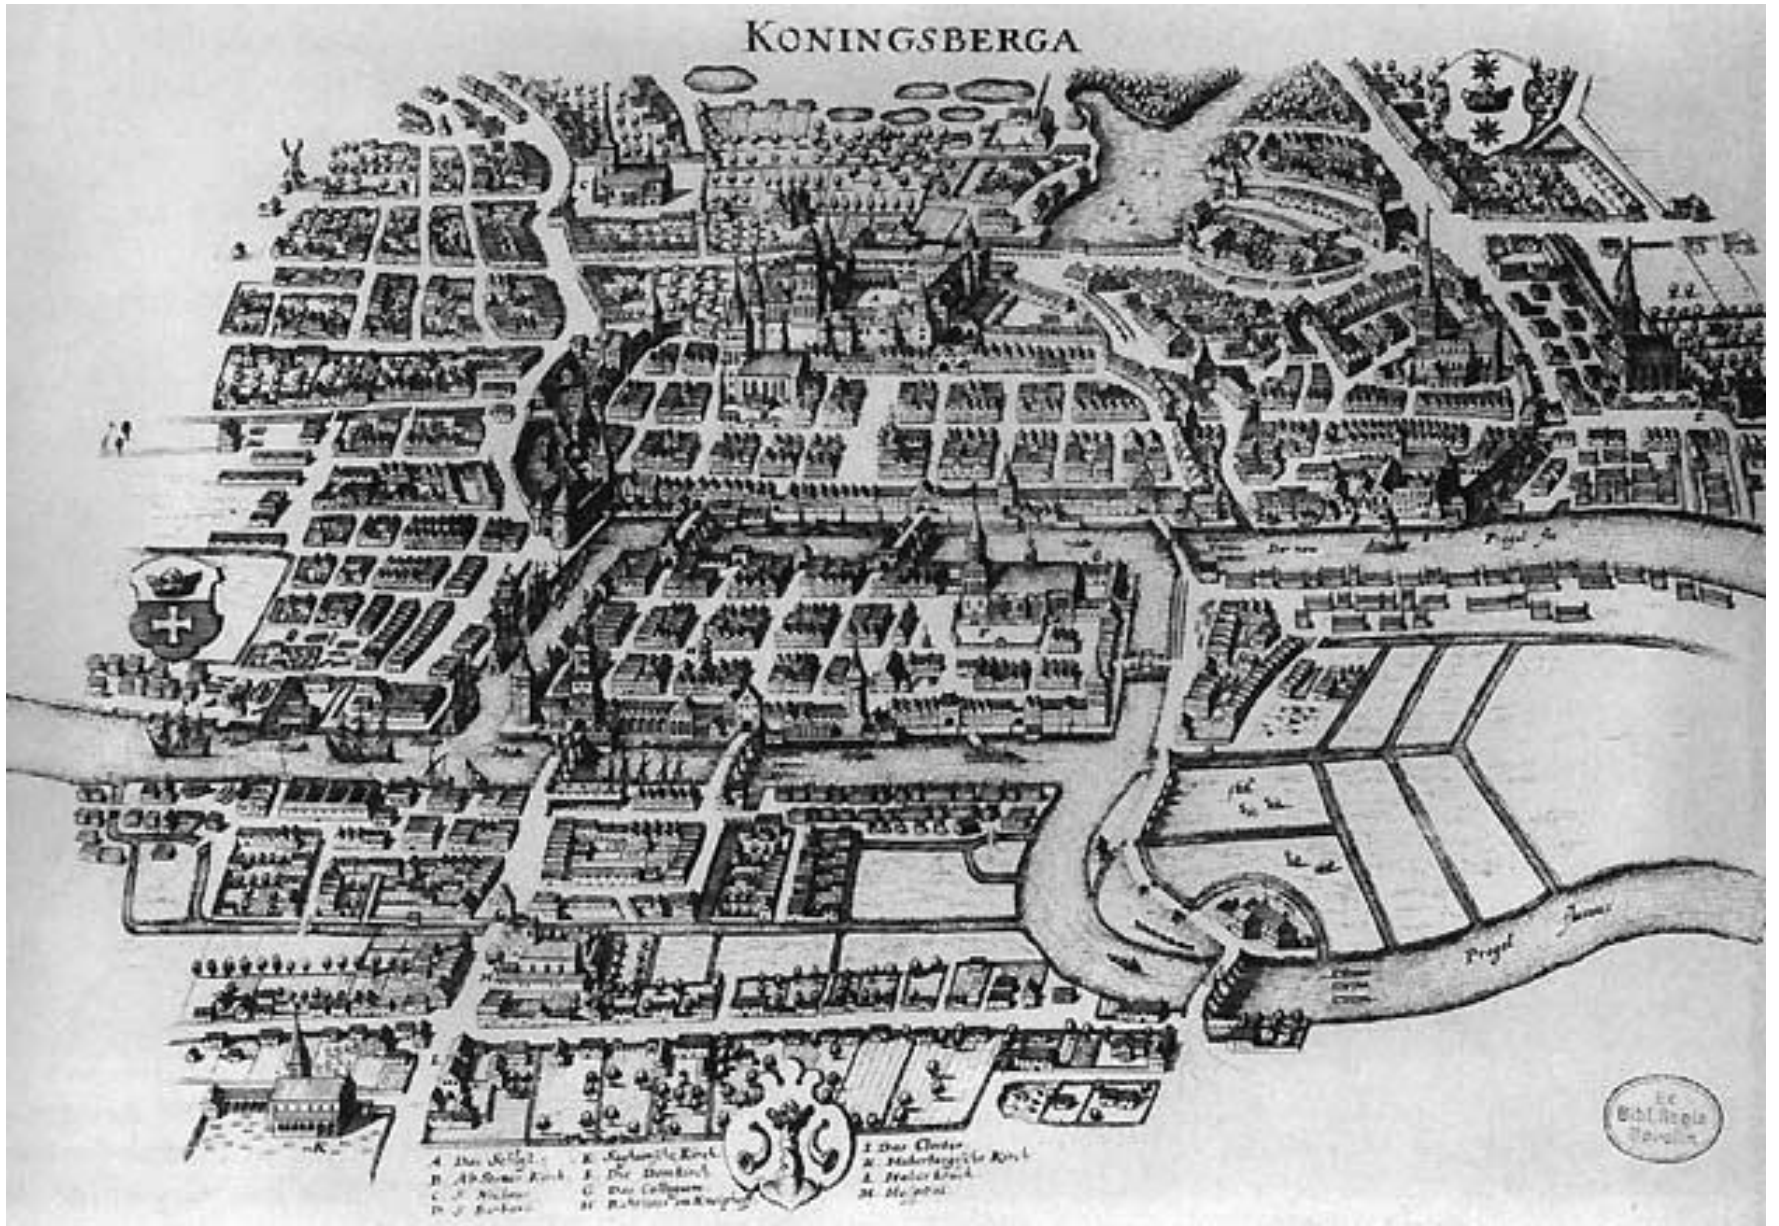
# Graph Problems & Representations



# What is a Graph?


- $G = (V, E)$ , where
  - $V$  represents the set of vertices (nodes)
  - $E$  represents the set of edges (links)
  - Both vertices and edges may contain additional information (e.g., edge weights)
- Different types of graphs:
  - directed vs. undirected edges
  - presence or absence of cycles

# KONINGSBERGA

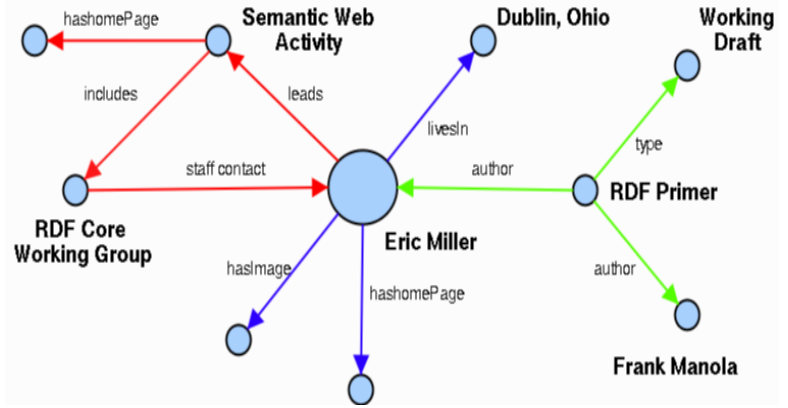


# We See Graphs Everywhere

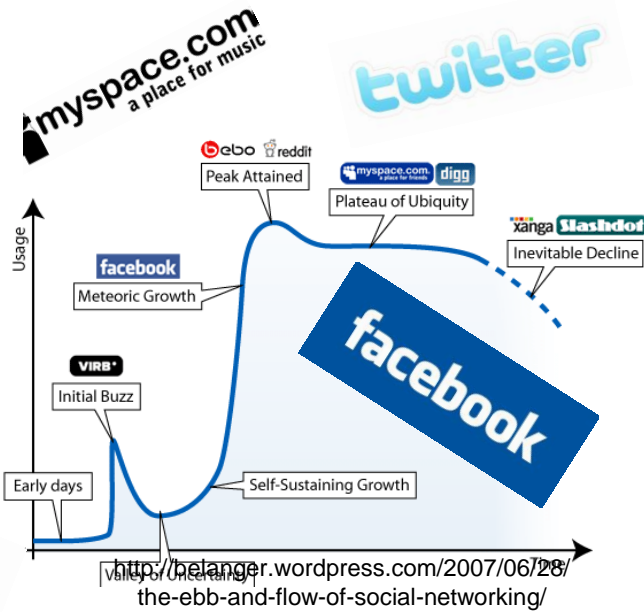
- Ubiquitous network (graph) data
  - Technological Network
    - Internet
  - Information Network
    - WWW, Sematic Web/Ontologies, XML/RD
  - Social network
  - Biological Network
  - Financial Network
  - Transportation Network

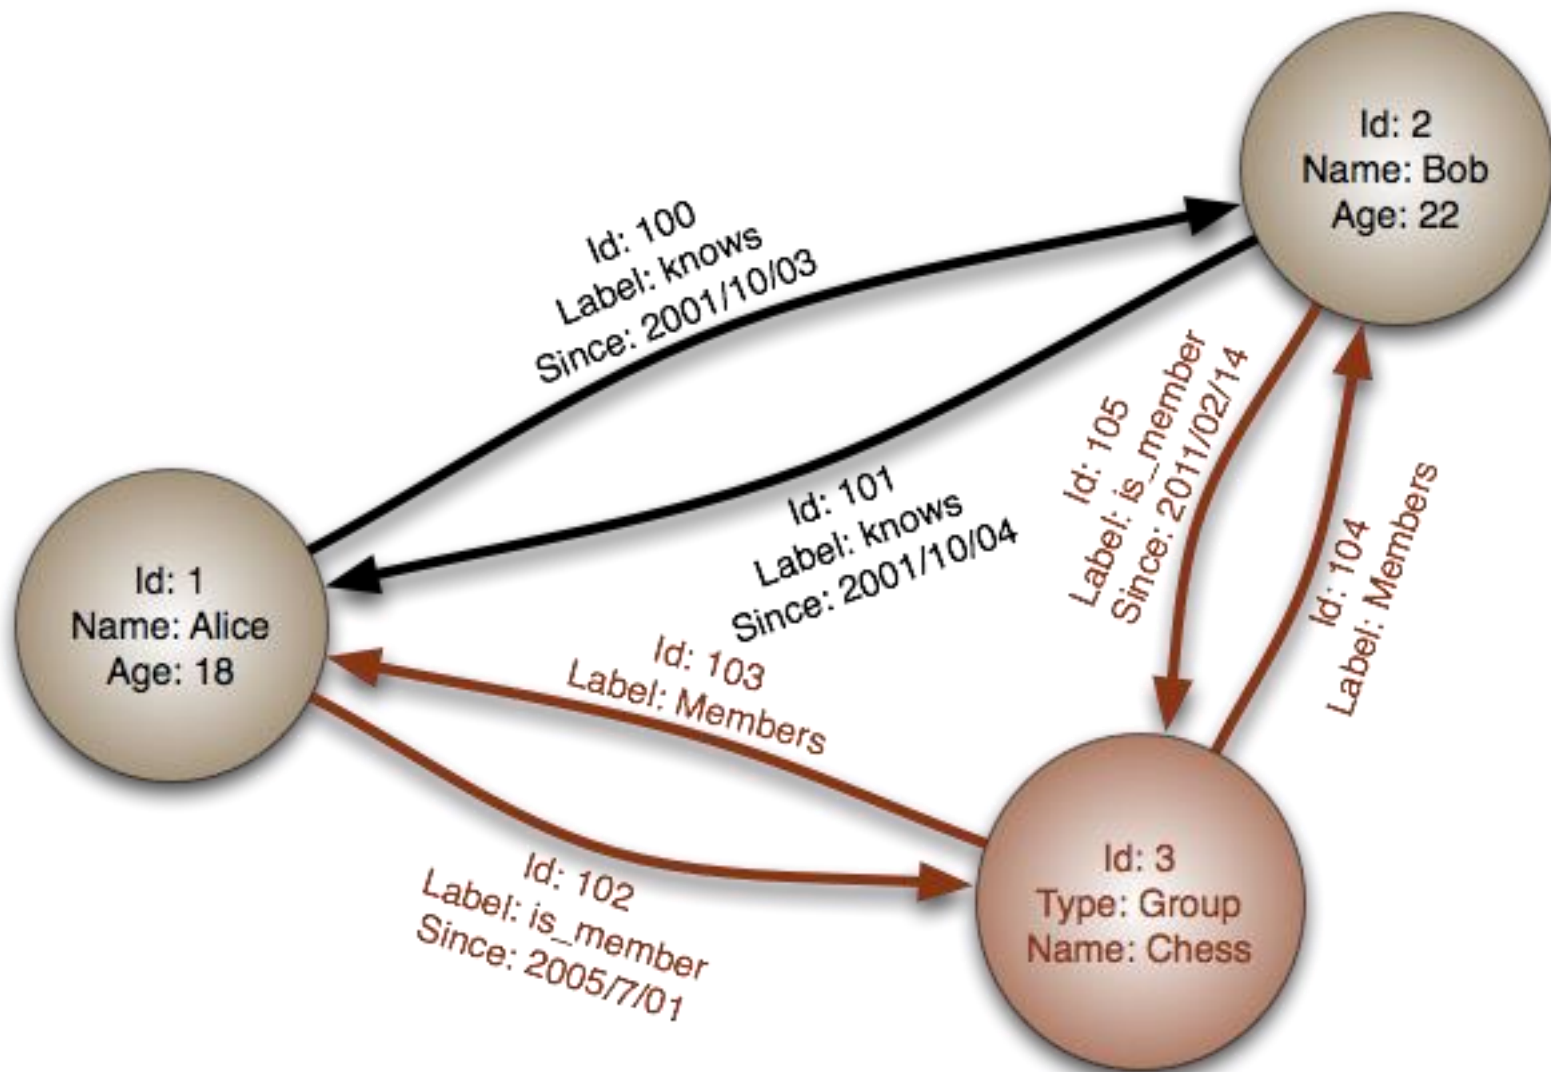


- W3C Tech Reports
- W3C Staff
- W3C Organization



Semantic Search, Guha et. al., WWW'03





# Some Graph Problems

- Finding shortest paths
  - Routing Internet traffic and UPS trucks
- Finding minimum spanning trees
  - Telecommunication companies laying down fibre
- Finding max flow
  - Airline scheduling



# Some Graph Problems

- Identify “special” nodes and communities
  - Breaking up terrorist cells, spread of avian flu
- Bipartite matching
  - Monster.com, Match.com
- And of course... PageRank

# Challenge in Dealing with Graph Data

- Flat Files
  - No query support
- RDBMS
  - Can store the graph
  - But limited support for graph query
    - Connect-By (Oracle)
    - Common Table Expressions (CTEs) (Microsoft)
    - Temporal Table

# Native Graph Databases

- An Emerging Field
  - [http://en.wikipedia.org/wiki/Graph\\_database](http://en.wikipedia.org/wiki/Graph_database)
- Storage and Basic Operators
  - **Neo4j** (an open source graph database),  
InfiniteGraph, VertexDB, ...
- Distributed Graph Processing (mostly in-memory-only)
  - Google's Pregel, GraphLab, ...

# The Graph Analytics Industry

- Status of Practice
  - Graph data in many industries
  - Graph analytics are powerful and can bring great business values/insights
  - Graph analytics not utilized enough in enterprises due to lack of available platforms/tools (except leading tech companies which have high caliber in house engineering teams and resources)

# Graphs and MapReduce

- Graph algorithms typically involve:
  - Performing computations at each node: based on node features, edge features, and local link structure
  - Propagating computations: “traversing” the graph
- Key questions:
  - How do you represent graph data in MapReduce?
  - How do you traverse a graph in MapReduce?


# Representing Graphs

- Two common representations
  - Adjacency matrix
  - Adjacency list

# Adjacency Matrices

- Represent a graph as an  $n \times n$  square matrix  $M$ 
  - $n = |\mathbf{V}|$
  - $M_{ij} = 1$  means a link from node  $i$  to  $j$

	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>
<b>1</b>	0	1	0	1
<b>2</b>	1	0	1	1
<b>3</b>	1	0	0	0
<b>4</b>	1	0	1	0



# Adjacency Matrices: Critique

- Advantages:
  - Amenable to mathematical manipulation
  - Iteration over rows and columns corresponds to computations on out-links and in-links
- Disadvantages:
  - Lots of zeros for sparse matrices
  - Lots of wasted space



# Adjacency Lists

- Take adjacency matrices... and throw away all the zeros

	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>
<b>1</b>	0	1	0	1
<b>2</b>	1	0	1	1
<b>3</b>	1	0	0	0
<b>4</b>	1	0	1	0




1: 2, 4  
2: 1, 3, 4  
3: 1  
4: 1, 3

# Adjacency Lists: Critique

- Advantages:
  - Much more compact representation
  - Easy to compute over out-links
- Disadvantages:
  - Much more difficult to compute over in-links


# Parallel Breadth-First Search




# Single Source Shortest Path

- **Problem:** find shortest path from a source node to one or more target nodes
  - “shortest” might also mean lowest weight or cost
- First, a refresher: Dijkstra’s algorithm


# Dijkstra's Algorithm




# Dijkstra's Algorithm



# Dijkstra's Algorithm




# Dijkstra's Algorithm






# Dijkstra's Algorithm



# Dijkstra's Algorithm




# Single Source Shortest Path

- **Problem:** find shortest path from a source node to one or more target nodes
  - “shortest” might also mean lowest weight or cost
- On a single machine: Dijkstra’s algorithm
- MapReduce: Parallel Breadth-First Search (BFS)
  - Consider simplest case of equal edge weights first
  - Solution to the problem can be defined inductively


# Finding the Shortest Path

- Here's the intuition:
  - Define:  $b$  is reachable from  $a$  if  $b$  is in the adjacency list of  $a$
  - $\text{DISTANCETo}(s) = 0$
  - For all nodes  $p$  reachable from  $s$ :  
 $\text{DISTANCETo}(p) = 1$
  - For all nodes  $n$  reachable from some other set of nodes  $M$ :  
 $\text{DISTANCETo}(n) = 1 + \min_{m \in M} \text{DISTANCETo}(m)$

# Finding the Shortest Path



# Visualizing Parallel BFS





# From Intuition to Algorithm

- Data representation:
  - Key:  
node  $n$
  - Value:  
 $d$  (distance from start),  
adjacency list (list of nodes reachable from  $n$ )
  - Initialization:  
for all nodes except the start node,  $d = \infty$ .



# From Intuition to Algorithm

- Mapper:
  - $\forall m \in \text{adjacency list: emit } (m, d + 1)$
- Sort/Shuffle
  - Groups distances by reachable nodes
- Reducer:
  - Selects the minimum distance path for each reachable node
  - Additional bookkeeping needed to keep track of the actual path

# Multiple Iterations Needed

- Each MapReduce iteration advances the “known frontier” by one hop
  - Subsequent iterations include more and more reachable nodes as frontier expands
  - Multiple iterations are needed to explore entire graph
- Preserving graph structure:
  - Problem: Where did the adjacency list go?
  - Solution: mapper emits  $(n, \text{adjacency list})$  as well

# BFS Pseudo-Code

```
1: class MAPPER
2:   method MAP(nid  $n$ , node  $N$ )
3:      $d \leftarrow N.DISTANCE$ 
4:     EMIT(nid  $n$ ,  $N$ ) ▷ Pass along graph structure
5:     for all nodeid  $m \in N.ADJACENCYLIST$  do
6:       EMIT(nid  $m$ ,  $d + 1$ ) ▷ Emit distances to reachable nodes
7:
8: class REDUCER
9:   method REDUCE(nid  $m$ , [ $d_1, d_2, \dots$ ])
10:     $d_{min} \leftarrow \infty$ 
11:     $M \leftarrow \emptyset$ 
12:    for all  $d \in \text{counts } [d_1, d_2, \dots]$  do
13:      if ISNODE( $d$ ) then ▷ Recover graph structure
14:         $M \leftarrow d$  ▷ Look for shorter distance
15:      else if  $d < d_{min}$  then
16:         $d_{min} \leftarrow d$ 
17:
18:     $M.DISTANCE \leftarrow d_{min}$  ▷ Update shortest distance
19:    EMIT(nid  $m$ , node  $M$ )
```

# Stopping Criterion

- How many iterations are needed in parallel BFS (equal edge weight case)?
- Convince yourself: when a node is first “discovered”, we’ve found the shortest path
- Now answer the question...
  - Six degrees of separation?
- Practicalities of implementation in MapReduce

# Weighted Edges



- Now add positive weights to the edges
  - Why can't edge weights be negative?
- Simple change: adjacency list now includes a weight  $w$  for each edge
  - In mapper, emit  $(m, d + w_p)$  instead of  $(m, d + 1)$  for each node  $m$
- That's it?

# Stopping Criterion

- How many iterations are needed in parallel BFS (positive edge weight case)?
- Convince yourself: when a node is first “discovered”, we’ve found the shortest path

**Not true!**

# Additional Complexities



How many iterations are required to discover the shortest distances to all nodes from  $n_1$ ?

# Stopping Criterion

- How many iterations are needed in parallel BFS (positive edge weight case)?
- Practicalities of implementation in MapReduce



# Comparison to Dijkstra

- Dijkstra's algorithm is more efficient
  - At any step it only pursues edges from the minimum-cost path inside the frontier
- MapReduce explores all paths in parallel
  - Lots of “waste”
  - Useful work is only done at the “frontier”
- Why can't we do better using MapReduce?


# Implementation on Hadoop

<http://goo.gl/TEoU4>

# Graphs and MapReduce

- Generic recipe:
  - Represent graphs as adjacency lists
  - Perform local computations in mapper
  - Pass along partial results via out-links, keyed by destination node
  - Perform aggregation in reducer on in-links to a node
  - Iterate until convergence: controlled by external “driver”
  - Don’t forget to pass the graph structure between iterations

# PageRank




# Random Walks over the Web

- Random surfer model:
  - User starts at a random Web page
  - User randomly clicks on links, surfing from page to page
- PageRank
  - Characterizes the amount of time spent on any given page
  - Mathematically, a probability distribution over pages

# Random Walks over the Web

- PageRank captures the notion of page importance
  - Correspondence to human intuition?
  - One of thousands of features used in Web search
  - Note: query-independent

# PageRank: Simplified




# PageRank: Simplified

- Given page  $x$  with in-links  $t_1 \dots t_n$ , where
  - $C(t)$  is the out-degree of  $t$

$$PR(x) = \sum_{i=1}^n \frac{PR(t_i)}{C(t_i)}$$



# Example: the Web in 1839



	y	a	m
y	$1/2$	$1/2$	0
a	$1/2$	0	1
m	0	$1/2$	0

# Simulating a Random Walk

- Start with the vector  $\mathbf{v} = [1, 1, \dots, 1]$  representing the idea that each Web page is given one unit of *importance*.
- Repeatedly apply the matrix  $M$  to  $\mathbf{v}$ , allowing the importance to flow like a random walk.
- Limit exists, but about 50 iterations is sufficient to estimate final distribution.

# Example: the Web in 1839

- Equations  $\mathbf{v} = M \mathbf{v}$  :

$$y = y/2 + a/2$$

$$a = y/2 + m$$

$$m = a/2$$

$y$	=	1	1	5/4	9/8	...	6/5
$a$		1	3/2	1	11/8	...	6/5
$m$		1	1/2	3/4	1/2		3/5

# Solving the Equations

- Because there are no constant terms, these 3 equations in 3 unknowns do not have a unique solution.
- Add in the fact that  $y + a + m = 3$  to solve.
- In Web-sized examples, we cannot solve by Gaussian elimination, but we need to use the power method (= iterative solution).

# Computing PageRank



- Properties of PageRank
  - Can be computed iteratively
  - Effects at each iteration are local

# Computing PageRank

- Sketch of algorithm:
  - Start with seed  $PR_i$  values
  - Each page distributes its  $PR_i$  “credit” to all of its out-links
  - Each page adds up the “credits” from all of its in-links to compute  $PR_{i+1}$
  - Iterate until the values converge



# Sample PageRank Iterations

Iteration 1



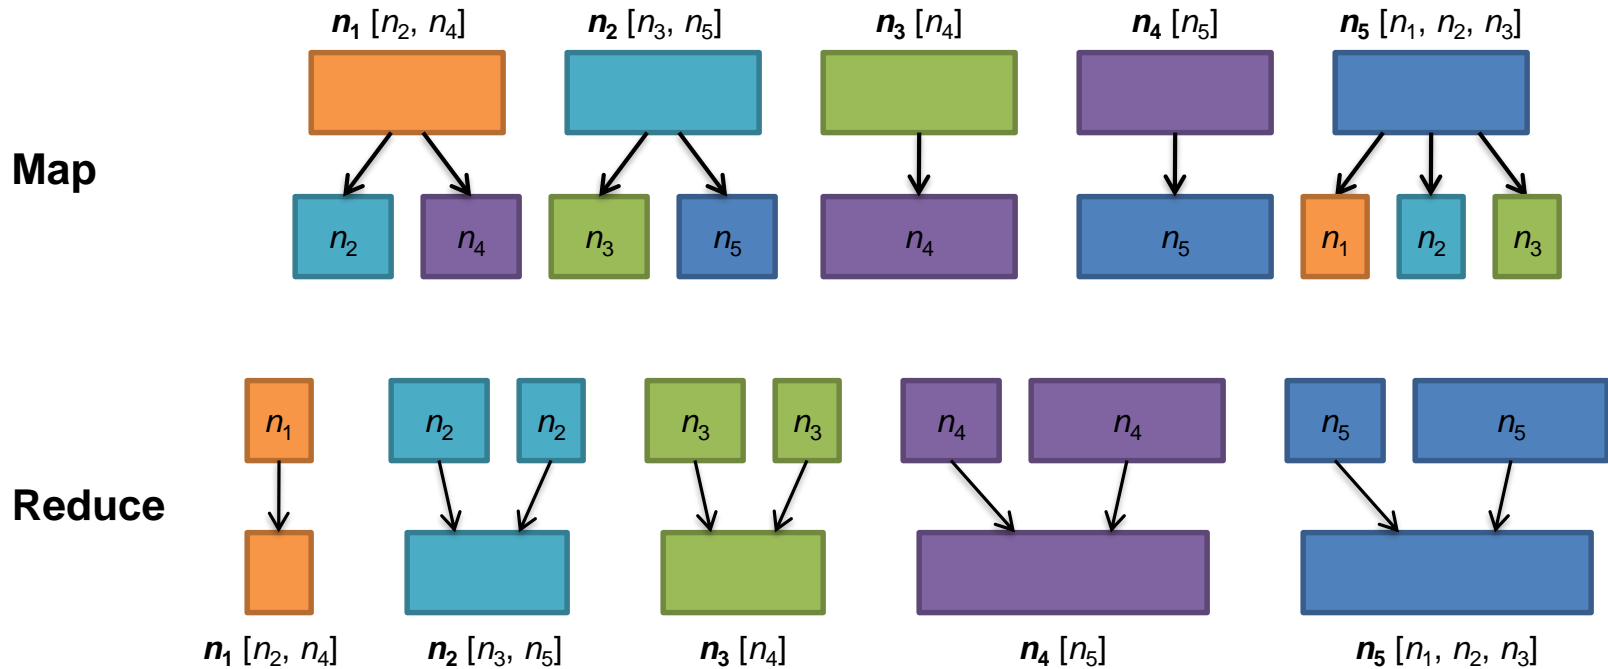
# Sample PageRank Iterations

Iteration 2





# PageRank in MapReduce




# PageRank Pseudo-Code

```
1: class MAPPER
2:   method MAP(nid  $n$ , node  $N$ )
3:      $p \leftarrow N.PAGERANK / |N.ADJACENCYLIST|$ 
4:     EMIT(nid  $n$ ,  $N$ ) ▷ Pass along graph structure
5:     for all nodeid  $m \in N.ADJACENCYLIST$  do
6:       EMIT(nid  $m$ ,  $p$ ) ▷ Pass PageRank mass to neighbors
1: class REDUCER
2:   method REDUCE(nid  $m$ , [ $p_1, p_2, \dots$ ])
3:      $M \leftarrow \emptyset$ 
4:     for all  $p \in$  counts [ $p_1, p_2, \dots$ ] do
5:       if ISNODE( $p$ ) then
6:          $M \leftarrow p$  ▷ Recover graph structure
7:       else
8:          $s \leftarrow s + p$  ▷ Sum incoming PageRank contributions
9:      $M.PAGERANK \leftarrow s$ 
10:    EMIT(nid  $m$ , node  $M$ )
```

# Real-World Problems

- Some pages are “dead ends” (no out-links).
  - Such a page causes importance to leak out.
- Some other (groups of) pages are *spider traps* (all out-links are within the group).
  - Eventually spider traps absorb all importance.

# Microsoft becomes a dead end



	y	a	m
y	1/2	1/2	0
a	1/2	0	0
m	0	1/2	0

# Microsoft becomes a dead end

- Equations  $\mathbf{v} = M \mathbf{v}$  :


$$y = y/2 + a/2$$

$$a = y/2$$

$$m = a/2$$

y	=	1	1	3/4	5/8	...	0
a	=	1	1/2	1/2	3/8	...	0
m	=	1	1/2	1/4	1/4	...	0

# Microsoft becomes a spider trap



	y	a	m
y	1/2	1/2	0
a	1/2	0	0
m	0	1/2	1

# Microsoft becomes a spider trap

- Equations  $\mathbf{v} = M \mathbf{v}$  :

$$y = y/2 + a/2$$

$$a = y/2$$

$$m = a/2 + m$$

y	=	1	1	3/4	5/8	...	0
a	=	1	1/2	1/2	3/8	...	0
m	=	1	3/2	7/4	2	...	3

# Google's Solution

- “Tax” each page a fixed percentage at each iteration.
- Add the same constant to all pages.
- Models a random walk with a fixed probability of going to a random place next.



# Example: with 20% Tax

- Equations  $\mathbf{v} = 0.8(M \mathbf{v}) + 0.2$ :

$$y = 0.8(y/2 + a/2) + 0.2$$

$$a = 0.8(y/2) + 0.2$$

$$m = 0.8(a/2 + m) + 0.2$$

$y$		1	1.00	0.84	0.776		7/11
$a$	=	1	0.60	0.60	0.536	...	5/11
$m$		1	1.40	1.56	1.688		21/11

# PageRank: Complete

- Two additional complexities
  - What is the proper treatment of dangling nodes (i.e., nodes with no out-links)?
  - How do we factor in the random jump factor?

# PageRank: Complete

- Solution:
  - Second pass to redistribute “missing PageRank mass” and account for random jumps

$$p' = \alpha \left( \frac{1}{N} \right) + (1 - \alpha) \left( \frac{m}{N} + p \right)$$

- $p$  is PageRank value from before,  $p'$  is updated PageRank value
- $N$  is the total number of nodes in the graph
- $m$  is the missing PageRank mass

# PageRank Convergence

- Alternative convergence criteria
  - Iterate until PageRank values don't change
  - Iterate until PageRank rankings don't change
  - Fixed number of iterations
- Convergence for web graphs?

# Beyond PageRank

- Link structure is important for web search
  - PageRank is one of many link analysis algorithms: HITS, SALSA, etc.
  - Used with thousands of other features in ranking...
- Adversarial nature of web search
  - Link spamming
  - Spider traps
  - Keyword stuffing
  - ...

# Efficient Graph Algorithms

- Sparse vs. Dense Graphs
- Graph Topologies

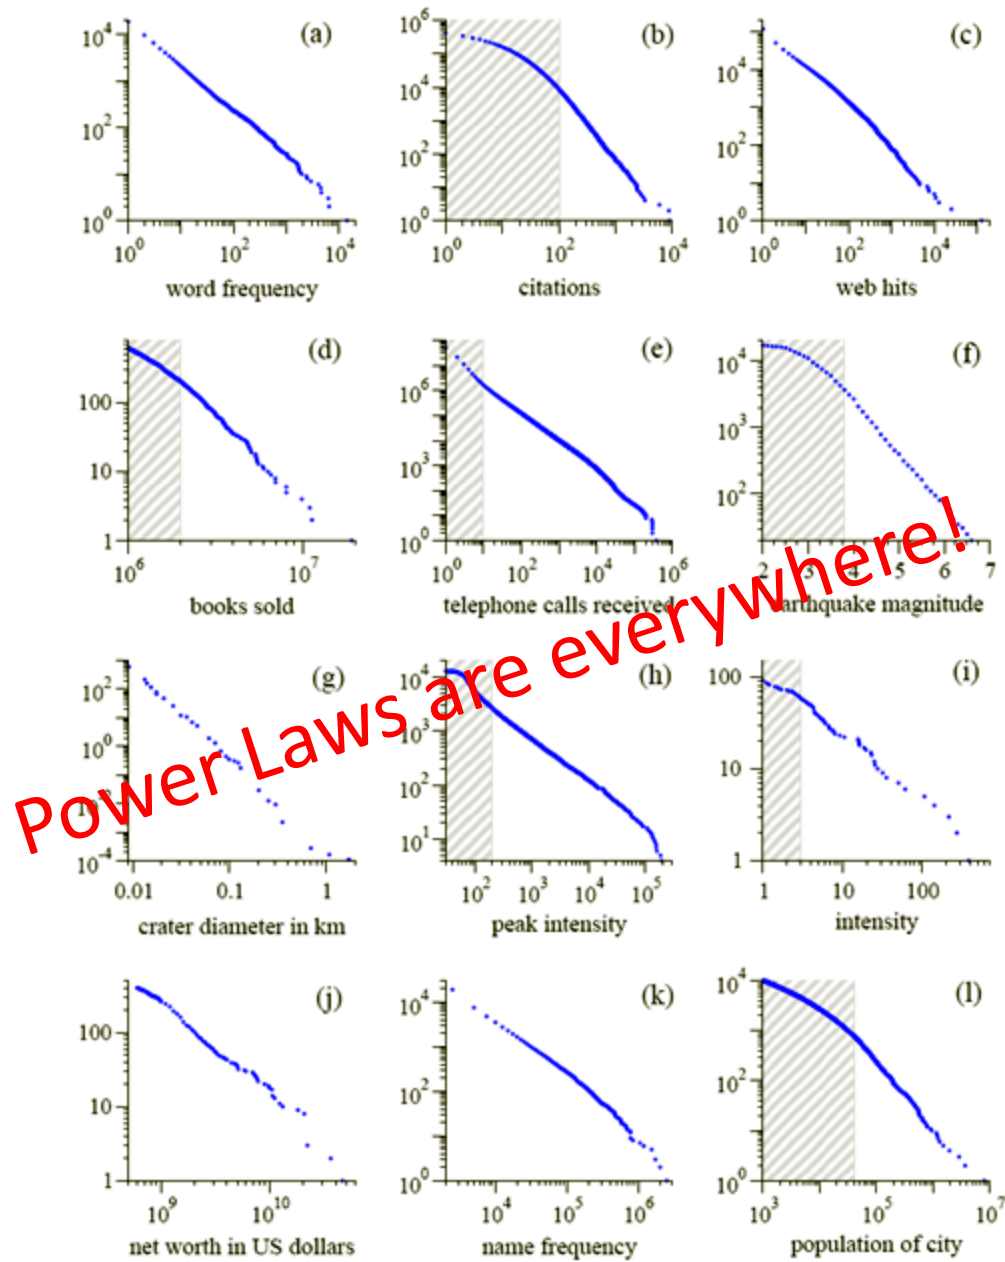


Figure from: Newman, M. E. J. (2005) "Power laws, Pareto distributions and Zipf's law." Contemporary Physics 46:323–351.

# Local Aggregation

- Use combiners!
  - In-mapper combining design pattern also applicable
- Maximize opportunities for local aggregation
  - Simple tricks: sorting the dataset in specific ways



# Take Home Messages

- Graph Problems and Representations
- Parallel Breadth-First Search
- PageRank: Simplified and Complete