



L24: MapReduce and DFS

Jeff M. Phillips


April 16, 2018

Don't put too much, or small font!


- ▶ Succinct title (and names)
- ▶ What is the problem and data you worked on?
- ▶ What were the key ideas in your approach?
- ▶ What techniques from the class did you use?
- ▶ What did you learn?

Big Data \leftrightarrow Data Science

Classic Scientific Paradigm



Data Science



Distributed File System (HDFS)

everything: $\langle \text{key}, \text{value} \rangle$ pair.

key

log id

web address

doc id

word

value

log

html, out-going-links

list of words

how many times (and where) it appears in set.

Really big data \rightarrow store in blocks (64 MB)

- 2-3 copies of each block.
- store each copy on different computers.
 \rightarrow at random.

Why little locality

- Resiliency: if computer goes down
↳ another that can take its place.
- Redundancy: split read time across cluster.
- Heterogeneous: different type of data.

infrequent updates
mainly appends.

Map Reduce

Queries over DFS

usually, output is large.

1. Map: Transformation $\langle k, v \rangle \rightarrow \langle k', v' \rangle$
when its useful to group (sort on k')

1.5. Combine: Preprocessing $\langle k', v_1 \rangle, \langle k', v_2 \rangle$ on same computer before shuffle.

2. Shuffle: Put all $\langle k', v_1 \rangle$ and $\langle k', v_2 \rangle$ with same k' on same computers.

3. Reduce: Take $\langle k', v_1 \rangle, \langle k', v_2 \rangle, \dots$
Map to set $\langle k'', U = f(v_1, v_2, v_3, \dots, v_n) \rangle$

Word Count

Need at least 20,50 computers.

Consider as input all of English Wikipedia stored in DFS. Goal is to count how many times each word is used.

Map: Read text for each
 $\rightarrow \langle w, 1 \rangle$

Combiner: $\langle w, v_1 \rangle, \langle w, v_2 \rangle, \dots$
 $\rightarrow \langle w, \sum_{i=1}^n v_i \rangle$

Reduce: For some w
 $\langle w, v_1 \rangle, \langle w, v_2 \rangle, \langle w, v_3 \rangle, \dots$

$\rightarrow \langle w, \sum_{i=1}^n v_i \rangle$

word $w \in \text{text}$

word	
"the"	7%
"on"	3.5%
"a"	2.8%

"the" wins by
last reducer"

Rounds

Rounds are slow.

Overhead. + write everything
to disk each round.

↳ Minimize # rounds

$\log_2 N$ too many.
 ≈ 30

⇒ Spork (Berkley)

Inverted Index

Consider as input all of English Wikipedia stored in DFS. Goal is to build an index, so each word has a list of pages it is in.

Map : For all words w on page p .
 $\hookrightarrow \langle w, p \rangle$

Reduce : $\langle w, p_1 \rangle, \langle w, p_2 \rangle, \dots$

$\hookrightarrow \langle w, \bigcup_{i=1}^m p_i \rangle$

\uparrow maybe sort by page rank.

Phrases

Consider as input all of English Wikipedia stored in DFS. Goal is to build an index, on 3-grams (sequence of 3 words) that appears on exactly one page, with link to page.

Map \rightarrow $\langle w, p \rangle$ $w = 3\text{-gram}$

Combine $: \langle w, p \rangle, \langle w, p \rangle \rightarrow \langle w, p \rangle$

Reduce $: \text{if } \langle w, p_1 \rangle, \langle w, p_2 \rangle, \dots$
has ≥ 1 $p_i \rightarrow \emptyset$
o.w. $\langle w, p \rangle$

Page Rank

probability transition matrix M

column $M_i = \begin{bmatrix} 0 \\ 1/2 \\ \vdots \\ 0 \\ \vdots \\ 1/n \end{bmatrix}$ ← index of page j
 $\frac{j \text{ links to}}{\text{out of } k \text{ pages}}$

$$\beta = 0.85$$

$$P = \beta M + (1 - \beta) B$$

$$B = \begin{bmatrix} \frac{1}{m} & & \\ & \ddots & \\ & & \frac{1}{m} \end{bmatrix} \downarrow m$$

$$g_{i+1} = P g_i = (\beta M + (1 - \beta) B) g_i \xrightarrow{m}$$

Block (set of web pages)
(= set of columns M)

Map: Block $M[a, b]$, g_{i-1} g_i
 $g_i[a, b] \quad \langle k, g_{i+1} = M[a, b] \cdot g_i[a, b] \rangle$

Shuffle

Reduce $g_{i+1} = \beta \cdot \sum_{j=1}^n (g_{i+1})_j + (1-\beta) \begin{bmatrix} \frac{1}{n} \\ \frac{1}{n} \\ \vdots \\ \frac{1}{n} \end{bmatrix}$

Example PageRank

$$M = \begin{bmatrix} 0 & 1/2 & 0 & 0 \\ 1/3 & 0 & 1 & 1/2 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 1/2 & 0 & 0 \end{bmatrix}$$

out
url/page

Example PageRank

What if

g_i too big
for 1 corp.

$$M = \begin{bmatrix} 0 & 1/2 & 0 & 0 \\ 1/3 & 0 & 1 & 1/2 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 1/2 & 0 & 0 \end{bmatrix}$$

Stripes:

$$M_1 = \begin{bmatrix} 0 \\ 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1/2 \\ 0 \\ 0 \\ 1/2 \end{bmatrix}$$

$$M_3 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$$

$$M_4 = \begin{bmatrix} 0 \\ 1/2 \\ 1/2 \\ 0 \end{bmatrix}$$

These are stored as $(1 : (1/3, 2), (1/3, 3), (1/3, 4)),$
 $(2 : (1/2, 1)(1/2, 4)),$ $(3 : (1, 3)),$ and $(4 : (1/2, 2), (1/2, 3)).$

Example PageRank

$$M = \begin{bmatrix} 0 & 1/2 & 0 & 0 \\ 1/3 & 0 & 1 & 1/2 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 1/2 & 0 & 0 \end{bmatrix}$$

Blocks:

$$M_{1,1} = \begin{bmatrix} 0 & 1/2 \\ 1/3 & 0 \end{bmatrix} \quad M_{1,2} = \begin{bmatrix} 0 & 0 \\ 1 & 1/2 \end{bmatrix} \quad M_{2,1} = \begin{bmatrix} 1/3 & 0 \\ 1/3 & 1/2 \end{bmatrix} \quad M_{2,2} = \begin{bmatrix} 0 & 1/2 \\ 0 & 0 \end{bmatrix}$$

These are stored as $(1 : (1/2, 2))$, $(2 : (1/3, 1))$, as $(2 : (1, 3), (1/2, 4))$, as $(3 : (1/3, 1))$, $(4 : (1/3, 1), (1/2, 2))$, and as $(3 : (1/2, 4))$.