

DS256:Jan17 (3:1)

L5,6:MapReduce Algorithm Design

Yogesh Simmhan

31 Jan & 2 Feb, 2017

©Department of Computational and Data Science, IISc, 2016

This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/)

Copyright for external content used with attribution is retained by their original authors



Map-Only Design

Filtering: Distributed Grep

- Input
 - Lines of text from HDFS
 - “Search String” (e.g. regex), input parameter to job
- Mapper
 - Search line for string/pattern
 - Output matching lines
- Reducer
 - Identity function (output = input), or none at all



Accumulation (Histogram)

- List of courses with number of students enrolled in each (Gol scheme with citizens enrolled in each)
- Input
 - `<StudentID, CourseID>`
 - `<2482, SE256> <6427, SE252> <1635, E0 259>`
- Mapper
 - Emit `<CourseID, 1>`
 - `<SE256, 1>, <SE252, 1>, <E0 259, 1>`
- Partition
 - By Course ID
- Sort `<E0 259, 1>, <SE252, 1>, <SE256, 1>`
- Reduce `<E0 259, [1,1]>, <SE252, [1]>, <SE256, [1,1,1]>`
 - Count number of students per Course.
 - Output `<Course ID, Count>`
 - `<SE256, 2>, <SE252, 1>, <E0 259, 3>`

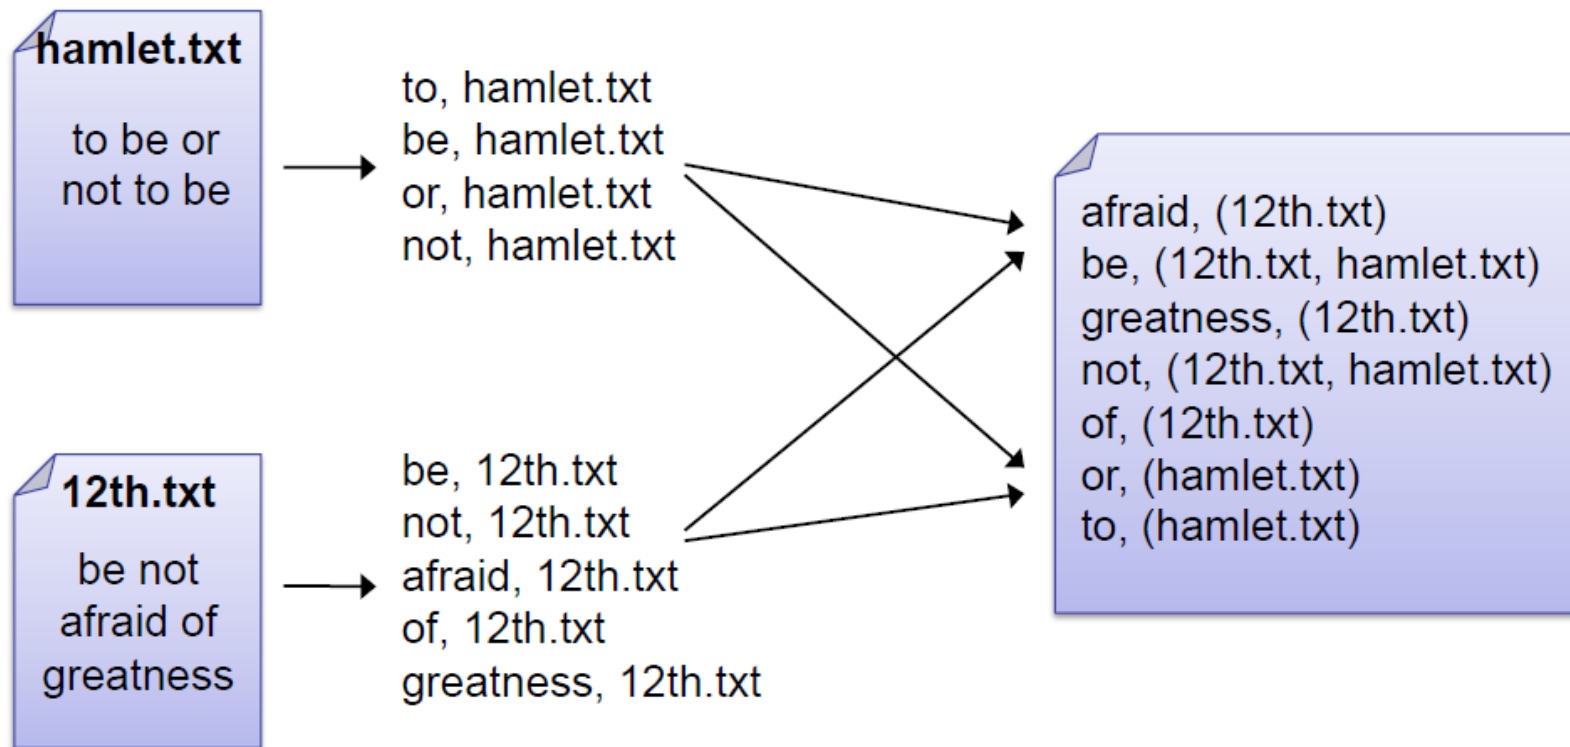


Inverted Index

- Convert from Key:Values to Value:Keys form
 - E.g. `<URL, Lines>` \rightarrow `<Word:URL[]>`
 - Useful for building search index
- Input: `<URL, Line>`
- Map: `foreach(Word in Line) emit(Word, URL)`
- Combiner: `Combine URLs for same Word`
- Reduce: `emit(Word, sort(URL[]))`



Inverted Index Example





Join

Customers

cfirstname	clastname	cphone	cstreet	czipcode
Tom	Jewett	714-555-1212	10200 Slater	92708
Alvaro	Monge	562-333-4141	2145 Main	90840
Wayne	Dick	562-777-3030	1250 Bellflower	90840

Orders

cfirstname	clastname	cphone	orderdate	soldby
Alvaro	Monge	562-333-4141	2003-07-14	Patrick
Wayne	Dick	562-777-3030	2003-07-14	Patrick
Alvaro	Monge	562-333-4141	2003-07-18	Kathleen
Alvaro	Monge	562-333-4141	2003-07-20	Kathleen

Customers joined to Orders

cfirstname	clastname	cphone	cstreet	czipcode	orderdate	soldby
Alvaro	Monge	562-333-4141	2145 Main	90840	2003-07-14	Patrick
Wayne	Dick	562-777-3030	1250 Bellflower	90840	2003-07-14	Patrick
Alvaro	Monge	562-333-4141	2145 Main	90840	2003-07-18	Kathleen
Alvaro	Monge	562-333-4141	2145 Main	90840	2003-07-20	Kathleen

<http://www.tomjewett.com/dbdesign/dbdesign.php?page=join.php>



Join


- Given two sets of files, combine the lines having the same key in each file
- Input:
 - `<customer_data>, <order_data>`
- Mapper:
 - `emit <cell, <t1,customer_data>>, <cell, <t2,order_data>>`
- Reduce:
 - If only one table ID (customer or order value) present, skip
 - If 2 values present, one from each tables ID
 - Just concatenate and emit the pair
 - `<cell, [customer_data, order_data]>`
 - If multiple values present for each table ID,
 - Emit cross product of `customer_data*` and `order_data*` values, i.e., local join for each `cell` key
 - `<cell, [customer_data*, order_data*]>`



Reverse graph edge directions & output in node order

- Input: adjacency list of graph (e.g. 3 nodes and 4 edges)

(3, [1, 2]) (1, [3])
(1, [2, 3]) → (2, [1, 3])
(3, [1])



- `node_ids` in the output values are also sorted.
But Hadoop only sorts on keys!
- MapReduce format
 - Input: (3, [1, 2]), (1, [2, 3]).
 - Intermediate: (1, [3]), (2, [3]), (2, [1]), (3, [1]). (reverse edge direction)
 - Out: (1,[3]) (2, [1, 3]) (3, [[1]).



Scalable Hadoop Algorithms: Themes

- Avoid object creation
 - Inherently costly operation
 - Garbage collection
- Avoid buffering
 - Limited heap size
 - Works for small datasets, but won't scale!



Importance of Local Aggregation

- Ideal scaling characteristics:
 - Twice the data, twice the running time
 - Twice the resources, half the running time
- Why can't we achieve this?
 - Synchronization requires communication
 - Communication kills performance
- Thus... avoid communication!
 - Reduce intermediate data via local aggregation
 - Combiners can help



Design Pattern for Local Aggregation

- “In-mapper combining”
 - Fold the functionality of the combiner into the mapper by preserving state across multiple map calls
- Advantages
 - Speed
 - Why is this faster than actual combiners?
- Disadvantages
 - Explicit memory management required
 - Potential for order-dependent bugs



Combiner Design

- Combiners and reducers share same method signature
 - Sometimes, reducers can serve as combiners
 - Often, not...
- Remember: combiner are optional optimizations
 - Should not affect algorithm correctness
 - May be run 0, 1, or multiple times
- Example: find average of all integers associated with the same key



Advanced Algorithms



Term Co-occurrence

- Term co-occurrence matrix for a text collection
 - $M = N \times N$ matrix ($N =$ vocabulary size)
 - M_{ij} : number of times i and j co-occur in some context
(for concreteness, let's say context = sentence)
- Why?
 - Distributional profiles as a way of measuring semantic distance
 - Semantic distance useful for many language processing tasks



MapReduce: Large Counting Problems

- Term co-occurrence matrix for a text collection = specific instance of a large counting problem
 - A large event space (number of terms)
 - A large number of observations (the collection itself)
 - Goal: keep track of interesting statistics about the events
- How do we compute using MapReduce?
 - Map Input: DocID, DocContent
- Basic approach
 - Mappers generate partial counts
 - Reducers aggregate partial counts

How do we aggregate partial counts efficiently?



First Try: "Pairs"

- Each mapper takes a sentence:
 - Generate all co-occurring term pairs
 - For all pairs, emit (a, b) → count
- Reducers sum up counts associated with these pairs
- Use combiners!



Pairs Approach

```
1: class MAPPER
2:   method MAP(docid  $a$ , doc  $d$ )
3:     for all term  $w \in \text{doc } d$  do
4:       for all term  $u \in \text{NEIGHBORS}(w)$  do
5:         EMIT(pair ( $w, u$ ), count 1)      ▷ Emit count for each co-occurrence

1: class REDUCER
2:   method REDUCE(pair  $p$ , counts [ $c_1, c_2, \dots$ ])
3:      $s \leftarrow 0$ 
4:     for all count  $c \in \text{counts } [c_1, c_2, \dots]$  do
5:        $s \leftarrow s + c$                   ▷ Sum co-occurrence counts
6:     EMIT(pair  $p$ , count  $s$ )
```

Figure 3.8: Pseudo-code for the “pairs” approach for computing word co-occurrence matrices from large corpora.

- Mapper emits many intermediate pairs (cell values)
- Combiner operates on sparse keys



Another Try: "Stripes"

- Idea: group together pairs into an associative array

$(a, b) \rightarrow 1$

$(a, c) \rightarrow 2$

$(a, d) \rightarrow 5$

$(a, e) \rightarrow 3$

$(a, f) \rightarrow 2$

$a \rightarrow \{ b: 1, c: 2, d: 5, e: 3, f: 2 \}$

- Each mapper takes a sentence:

- Generate all co-occurring term pairs

- For each term, emit $a \rightarrow \{ b: \text{count}_b, c: \text{count}_c, d: \text{count}_d, \dots \}$

$a \rightarrow \{ b: 1, \quad d: 5, e: 3 \}$

$a \rightarrow \{ b: 1, c: 2, d: 2, \quad f: 2 \}$

$a \rightarrow \{ b: 2, c: 2, d: 7, e: 3, f: 2 \}$

- Reducers perform element-wise sum of associative arrays



Stripes Approach

```
1: class MAPPER
2:   method MAP(docid  $a$ , doc  $d$ )
3:     for all term  $w \in \text{doc } d$  do
4:        $H \leftarrow \text{new ASSOCIATIVEARRAY}$ 
5:       for all term  $u \in \text{NEIGHBORS}(w)$  do
6:          $H\{u\} \leftarrow H\{u\} + 1$            ▷ Tally words co-occurring with  $w$ 
7:       EMIT(Term  $w$ , Stripe  $H$ )

1: class REDUCER
2:   method REDUCE(term  $w$ , stripes [ $H_1, H_2, H_3, \dots$ ])
3:      $H_f \leftarrow \text{new ASSOCIATIVEARRAY}$ 
4:     for all stripe  $H \in \text{stripes } [H_1, H_2, H_3, \dots]$  do
5:       SUM( $H_f, H$ )                               ▷ Element-wise sum
6:     EMIT(term  $w$ , stripe  $H_f$ )
```

Figure 3.9: Pseudo-code for the “stripes” approach for computing word co-occurrence matrices from large corpora.

- Mapper emits entire row at a time
- Combiner & Reducer operate on fewer keys
- Need to store entire row in memory!



“Stripes” Analysis


■ Advantages

- ▶ Far less sorting and shuffling of key-value pairs
- ▶ Can make better use of combiners

■ Disadvantages

- ▶ More difficult to implement
- ▶ Underlying object more heavyweight
- ▶ Fundamental limitation in terms of size of event space

Comparison of "pairs" vs. "stripes" for computing word co-occurrence matrices



Cluster size: 38 cores

Data Source: Associated Press Worldstream (APW) of the English Gigaword Corpus (v3), which contains 2.27 million documents (1.8 GB compressed, 5.7 GB uncompressed)



Relative Frequencies

- How do we estimate relative frequencies from counts?

$$f(B | A) = \frac{\text{count}(A, B)}{\text{count}(A)} = \frac{\text{count}(A, B)}{\sum_{B'} \text{count}(A, B')}$$

- Why do we want to do this?
- How do we do this with MapReduce?



$f(B|A)$: “Stripes”

- Easy!
 - One pass to compute $(a, *)$
 - Another pass to directly compute $f(B|A)$



$f(B|A)$: "Pairs"

$(a, *) \rightarrow 32$

Reducer holds this value in memory

$(a, b_1) \rightarrow 3$

$(a, b_2) \rightarrow 12$

$(a, b_3) \rightarrow 7$

$(a, b_4) \rightarrow 1$

...



$(a, b_1) \rightarrow 3 / 32$

$(a, b_2) \rightarrow 12 / 32$

$(a, b_3) \rightarrow 7 / 32$

$(a, b_4) \rightarrow 1 / 32$


...

■ For this to work:

- ▶ Must emit extra $(a, *)$ for every b_n in mapper
- ▶ Must make sure all a 's get sent to same reducer (use partitioner)
- ▶ Must make sure $(a, *)$ comes first (define sort order)
- ▶ Must hold state in reducer across different key-value pairs




Matrix-Vector Multiply





Matrix-Matrix Multiply





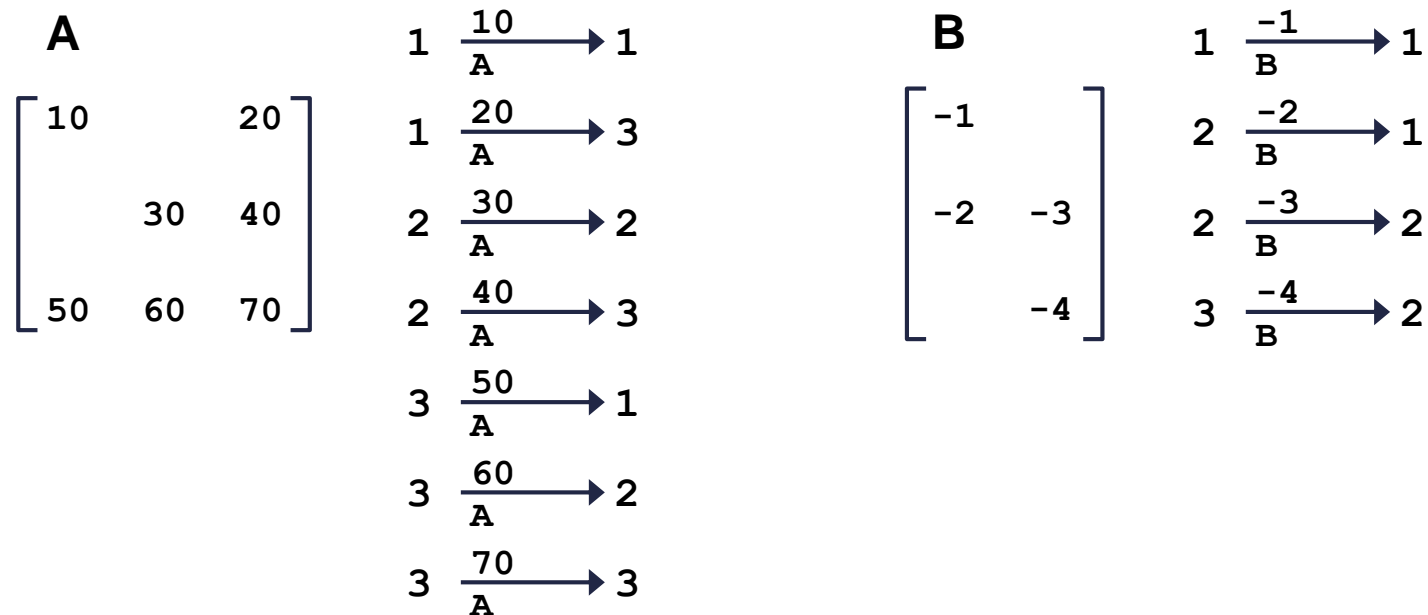
Sparse Matrix Multiplication

$$\begin{array}{c} \mathbf{A} \\ \left[\begin{array}{cc} 10 & 20 \\ & 30 & 40 \\ 50 & 60 & 70 \end{array} \right] \end{array} \quad \mathbf{X} \quad \begin{array}{c} \mathbf{B} \\ \left[\begin{array}{c} -1 \\ -2 & -3 \\ & -4 \end{array} \right] \end{array} \quad = \quad \begin{array}{c} \mathbf{C} \\ \left[\begin{array}{cc} -10 & -80 \\ -60 & -250 \\ -170 & -460 \end{array} \right] \end{array}$$

- ▶ Task: Compute product $C = A \cdot B$
- ▶ Assume most matrix entries are 0
- Motivation
 - ▶ Core problem in scientific computing
 - ▶ Challenging for parallel execution
 - ▶ Demonstrate expressiveness of Map/Reduce



Sparse Matrix Multiplication



- Represent matrix as list of nonzero entries
 - $\langle \text{row, col, value, matrixID} \rangle$
- Strategy
 - Phase 1: Compute all products $a_{i,k} \cdot b_{k,j}$
 - Phase 2: Sum products for each entry i,j
 - Each phase involves a Map/Reduce



Phase 1: Map

$$1 \frac{10}{A} \rightarrow 1$$

$$1 \frac{20}{A} \rightarrow 3$$

$$2 \frac{30}{A} \rightarrow 2$$

$$2 \frac{40}{A} \rightarrow 3$$

$$3 \frac{50}{A} \rightarrow 1$$

$$3 \frac{60}{A} \rightarrow 2$$

$$3 \frac{70}{A} \rightarrow 3$$

$$1 \frac{-1}{B} \rightarrow 1$$

$$2 \frac{-2}{B} \rightarrow 1$$

$$2 \frac{-3}{B} \rightarrow 2$$

$$3 \frac{-4}{B} \rightarrow 2$$

Key = row (2) $\frac{-2}{B}$ (1) Key = col

Key = 1

$$1 \frac{10}{A} \rightarrow 1$$

$$1 \frac{-1}{B} \rightarrow 1$$

$$3 \frac{50}{A} \rightarrow 1$$

Key = 2

$$2 \frac{30}{A} \rightarrow 2$$

$$2 \frac{-2}{B} \rightarrow 1$$

$$3 \frac{60}{A} \rightarrow 2$$

$$2 \frac{-3}{B} \rightarrow 2$$

Key = 3

$$1 \frac{20}{A} \rightarrow 3$$

$$2 \frac{40}{A} \rightarrow 3$$

$$3 \frac{70}{A} \rightarrow 3$$

$$3 \frac{-4}{B} \rightarrow 2$$

- Group values $a_{i,k}$ and $b_{k,j}$ according to key k



Phase 1: Reduce

Key = 1

$$1 \frac{10}{A} \rightarrow 1 \quad \times \quad 1 \frac{-1}{B} \rightarrow 1$$

$$3 \frac{50}{A} \rightarrow 1$$

Key = 2

$$2 \frac{30}{A} \rightarrow 2 \quad \times \quad 2 \frac{-2}{B} \rightarrow 1$$

$$3 \frac{60}{A} \rightarrow 2 \quad 2 \frac{-3}{B} \rightarrow 2$$

Key = 3

$$1 \frac{20}{A} \rightarrow 3 \quad \times \quad 3 \frac{-4}{B} \rightarrow 2$$

$$2 \frac{40}{A} \rightarrow 3$$

$$3 \frac{70}{A} \rightarrow 3$$

$$1 \frac{-10}{C} \rightarrow 1$$

$$3 \frac{-50}{A} \rightarrow 1$$

$$2 \frac{-60}{C} \rightarrow 1$$

$$2 \frac{-90}{C} \rightarrow 2$$

$$3 \frac{-120}{C} \rightarrow 1$$

$$3 \frac{-180}{C} \rightarrow 2$$

$$1 \frac{-80}{C} \rightarrow 2$$

$$2 \frac{-160}{C} \rightarrow 2$$

$$3 \frac{-280}{C} \rightarrow 2$$

- Generate all products $a_{i,k} \cdot b_{k,j}$



Phase 2: Map

$$1 \xrightarrow[C]{-10} 1$$

$$3 \xrightarrow[A]{-50} 1$$

$$2 \xrightarrow[C]{-60} 1$$

$$2 \xrightarrow[C]{-90} 2$$

$$3 \xrightarrow[C]{-120} 1$$

$$3 \xrightarrow[C]{-180} 2$$

$$1 \xrightarrow[C]{-80} 2$$

$$2 \xrightarrow[C]{-160} 2$$

$$3 \xrightarrow[C]{-280} 2$$

Key = row,col

$$\text{Key} = 1,1 \quad 1 \xrightarrow[C]{-10} 1$$

$$\text{Key} = 1,2 \quad 1 \xrightarrow[C]{-80} 2$$

$$\text{Key} = 2,1 \quad 2 \xrightarrow[C]{-60} 1$$

$$\text{Key} = 2,2 \quad 2 \xrightarrow[C]{-90} 2$$

$$2 \xrightarrow[C]{-160} 2$$

$$\text{Key} = 3,1 \quad 3 \xrightarrow[C]{-120} 1$$

$$3 \xrightarrow[A]{-50} 1$$

$$\text{Key} = 3,2 \quad 3 \xrightarrow[C]{-280} 2$$

$$3 \xrightarrow[C]{-180} 2$$

- ▶ Group products $a_{i,k} \cdot b_{k,j}$ with matching values of i and j



Phase 2: Reduce

Key = 1,1 1 $\xrightarrow{\frac{-10}{C}}$ 1

1 $\xrightarrow{\frac{-10}{C}}$ 1

▶ Key = 1,2 1 $\xrightarrow{\frac{-80}{C}}$ 2

final entries 1 $\xrightarrow{\frac{-80}{C}}$ 2

Key = 2,1 2 $\xrightarrow{\frac{-60}{C}}$ 1

2 $\xrightarrow{\frac{-60}{C}}$ 1

Key = 2,2 2 $\xrightarrow{\frac{-90}{C}}$ 2

2 $\xrightarrow{\frac{-250}{C}}$ 2

2 $\xrightarrow{\frac{-160}{C}}$ 2

C

-10 -80

$\begin{bmatrix} -60 & -250 \end{bmatrix}$

-170 -460

Key = 3,1 3 $\xrightarrow{\frac{-120}{C}}$ 1

3 $\xrightarrow{\frac{-170}{C}}$ 1

3 $\xrightarrow{\frac{-50}{A}}$ 1


Key = 3,2 3 $\xrightarrow{\frac{-280}{C}}$ 2

3 $\xrightarrow{\frac{-460}{C}}$ 2

3 $\xrightarrow{\frac{-180}{C}}$ 2



Block Matrix Multiply





PageRank

- Centrality measure of web page quality based on the web structure
 - How important is this vertex in the graph?
- Random walk
 - Web surfer visits a page, randomly clicks a link on that page, and does this repeatedly.
 - How frequently would each page appear in this surfing?
- Intuition
 - Expect high-quality pages to contain “endorsements” from many other pages thru hyperlinks
 - Expect if a high-quality page links to another page, then the second page is likely to be high quality too



PageRank, recursively

$$P(n) = \alpha \left(\frac{1}{|G|} \right) + (1 - \alpha) \sum_{m \in L(n)} \frac{P(m)}{C(m)}$$

- $P(n)$ is PageRank for webpage/URL 'n'
 - Probability that you're in vertex 'n'
- $|G|$ is number of URLs (vertices) in graph
- α is probability of random jump
- $L(n)$ is set of vertices that link to 'n'
- $C(m)$ is out-degree of 'm'





PageRank Iterations



$\alpha=0$

Initialize $P(n)=1/|G|$

Iteration 1



Iteration 2





PageRank using MapReduce

```
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
7:
8: class REDUCER
9:   method REDUCE(nid  $m$ , [ $p_1, p_2, \dots$ ])
10:     $M \leftarrow \emptyset$ 
11:    for all  $p \in$  counts [ $p_1, p_2, \dots$ ] do
12:      if ISNODE( $p$ ) then
13:         $M \leftarrow p$  ▷ Recover graph structure
14:      else
15:         $s \leftarrow s + p$  ▷ Sum incoming PageRank contributions
16:     $M.PAGERANK \leftarrow s$ 
17:    EMIT(nid  $m$ , node  $M$ )
```



PageRank using MapReduce

- MR run over multiple iterations (typically 30)
 - *The graph structure itself must be passed from iteration to iteration!*
- Mapper will
 - Initially, load adjacency list and initialize default PR
 - $\langle v1, \langle v2 \rangle + \rangle$
 - Subsequent iterations will load adjacency list and new PR
 - $\langle v1, \langle v2 \rangle +, pr1 \rangle$
 - Emit two types of messages from Map
 - PR messages and Graph Structure Messages
- Reduce will
 - Reconstruct the adjacency list for each vertex
 - Update the PageRank values for the vertex based on neighbour's PR messages
 - Write adjacency list and new PR values to HDFS, to be used by next Map iteration
 - $\langle v1, \langle v2 \rangle +, pr1' \rangle$



Inverted Indexes Revisited

- Each Map task parses one or more webpages
 - Input: A stream of webpages (WARC)
 - Output: A stream of (term, URL) tuples
 - (long, http://gb.com) (long, http://gb.com) (ago, http://gb.com) ...
(long, http://jn.in) (years, http://jn.in) (ago, http://jn.in) ...
- Shuffle sorts by key and routes tuples to Reducers
- Reducers convert streams of keys into streams of inverted lists
 - Sorts the values for a key (**why?**) and builds an inverted list
 - Output: (long, [http://gb.com, http://jn.in]), (ago, [http://gb.com, http://jn.in]), (years, [http://jn.in])



Optimizations & Extensions

- URL sizes may be large
 - Replace URLs with unique longs, URL ID
 - Mapping from URL ID to URL saved as a file
 - Inverted Index has <term, [URL ID]+>
 - Skip stop words with lot of matching URLs
 - Use combiners
- Partition term by prefix alphabet(s)
 - One reducer for each term starting with “a”, “b”, etc.
 - Part file from each reducer has terms with unique a starting letter
- Additional metadata
 - **Idea**: Include a mapping from URL ID to <URL, PageRank>?
 - Include “term frequency” of term occurrence per URL ID in Inverted Index?



Challenge

- Even using URL IDs, all IDs per term *may not fit in reduce memory for sorting*
 - E.g. 17M URLs in 1% of CC data.
 - Say 1000 unique words per URL.
 - So 17B keys and values generated by Mappers.
 - Say 50,000 unique words (keys) in English
 - One key would on average have $17\text{B}/50\text{K}=340\text{K}$ URL IDs
 - Peak values would be much higher
- Use a value-to-key conversion design pattern
 - Let MR perform sorting, Reducer just emits result



```
1: class MAPPER
2:   method MAP(docid  $n$ , doc  $d$ )
3:      $H \leftarrow$  new ASSOCIATIVEARRAY
4:     for all term  $t \in$  doc  $d$  do
5:        $H\{t\} \leftarrow H\{t\} + 1$ 
6:     for all term  $t \in H$  do
7:       EMIT(tuple  $\langle t, n \rangle$ , tf  $H\{t\}$ )

1: class REDUCER
2:   method INITIALIZE
3:      $t_{prev} \leftarrow \emptyset$ 
4:      $P \leftarrow$  new POSTINGSLIST
5:   method REDUCE(tuple  $\langle t, n \rangle$ , tf  $[f]$ )
6:     if  $t \neq t_{prev} \wedge t_{prev} \neq \emptyset$  then
7:       EMIT(term  $t$ , postings  $P$ )
8:        $P.RESET()$ 
9:        $P.ADD(\langle n, f \rangle)$ 
10:     $t_{prev} \leftarrow t$ 
11:   method CLOSE
12:     EMIT(term  $t$ , postings  $P$ )
```

- Mapper emits $\langle\langle$ term, URL ID $\rangle\rangle$, tf \rangle
 - i.e. compound key
- Partitioner sends all terms to the same reducer
- Per reducer, MR sorts based on compound key \langle term, URL ID \rangle
- Only one value for each compound key
- Reduce task gets list of term and URL ID in sorted order
 - When new term seen, flush index for “prev” term and start new term
- E.g.
 - $\langle\langle$ Ago, 1 $\rangle\rangle$, tf1 \rangle
 - $\langle\langle$ Ago, 7 $\rangle\rangle$, tf7 \rangle
 - Flush \langle Ago, [\langle 1,tf1 \rangle , \langle 7,tf7 \rangle]
 - $\langle\langle$ Long, 3 $\rangle\rangle$, tf3 \rangle
 - $\langle\langle$ Long, 4 $\rangle\rangle$, tf4 \rangle
 - $\langle\langle$ Long, 6 $\rangle\rangle$, tf6 \rangle
 - Flush \langle Long, [\langle 3,tf3 \rangle , \langle 4,tf4 \rangle , \langle 6,tf6 \rangle]



Lookup of Terms

- Each Map task loads one of the index files, say, by alphabet
- Input terms e.g. “t1 & t2 & t3” passed to each Map task as AND search
- Map does lookup and sends $\langle \text{URL ID}, t_i \rangle$ to reducer
 - Optionally send $\langle \langle \text{PR}, \text{URL ID} \rangle, t_i \rangle$ for sorting by PR
- Reducer does set intersection of all t_i for a URL ID
 - If all terms match, looks up URL for the URL ID
 - If PR stored for each URL, that is returned too



Thank You!

©Department of Computational and Data Science, IISc, 2016

This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/)

Copyright for external content used with attribution is retained by their original authors

