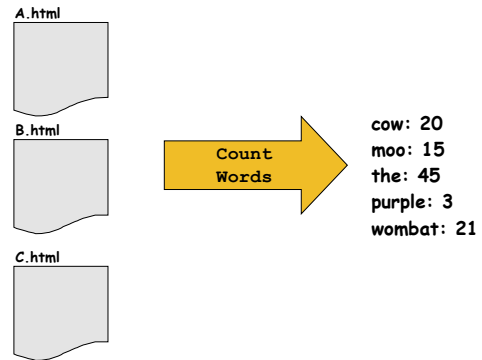


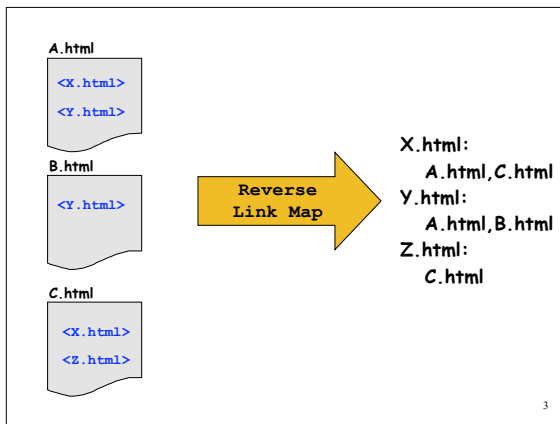
Google's MapReduce and Sawzall

CSCI 334
Stephen Freund

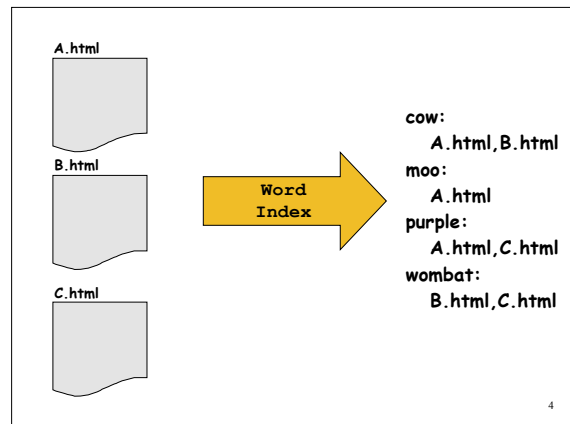
1



2



3



4

Computations Over Data

- Word Count
- Reverse Link Map
- Word Index
- Links out of a domain
- Page Rank
- But.... many terabytes of data
 - 1 terabyte = 1000 Gigabytes
 - = 1 099 511 627 776 bytes

5

Computing Infrastructure

- > 100,000 computers (as of 2004)
- Distribute network of clusters around world
- Problems:
 - need to coordinate computers
 - machines fail constantly
 - network, failure, computer/data locations, etc. should be transparent to user running analyses.

6

Dean & Ghemawat

We designed a new abstraction that allows us to express the simple computations we were trying to perform, but hides the messy details of parallelism, fault-tolerance, data distribution, and load-balancing... Our abstraction is inspired by the *map* and *reduce* primitives in Lisp and many other functional languages.

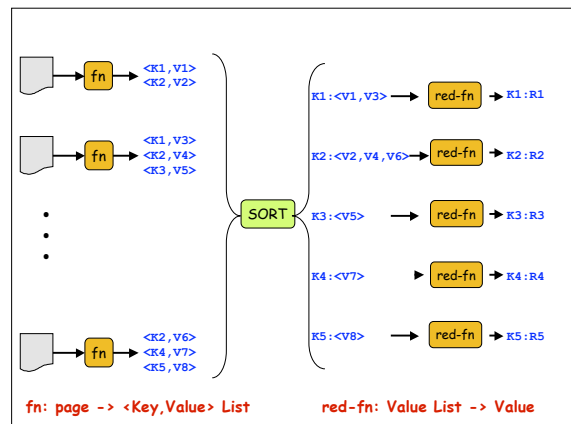
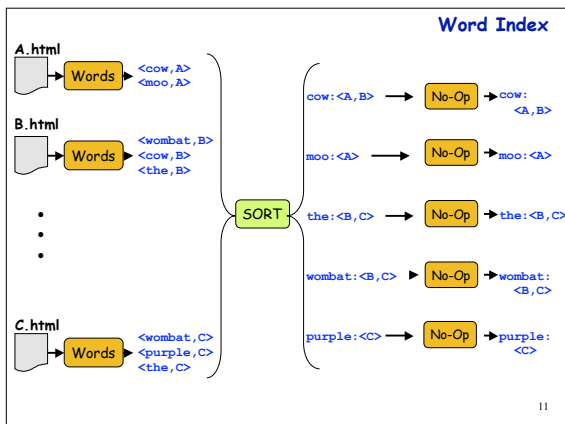
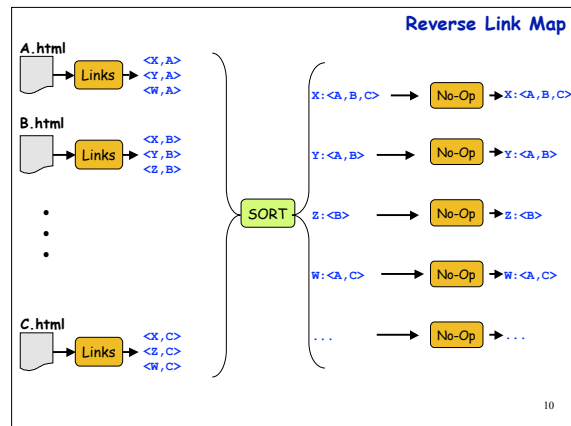
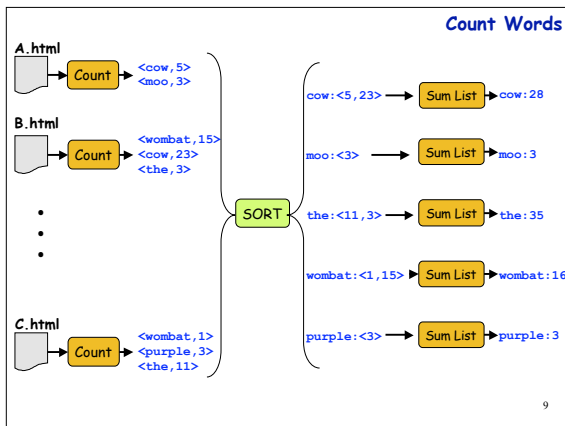
- Map: apply fn to every element of list
- Reduce: combine values in list to form a single "summary value"

7

MapReduce and Sawzall

- MapReduce (Dean and Ghemawat)
 - distributed computer management
- Sawzall (Pike et al.)
 - language for writing code to perform data analysis

8



Summary

- Performance:
 - search 10 billion unordered records in 80 sec
 - sort 10 billion records in 800 sec
- Page Rank: 24 separate map-reduce operations
- Sawzall/MapReduce execution model:
 - specify data set, map fn, reduce fn
 - most map/reduce functions < 50 lines of code
 - hides details of distributed system
 - fault tolerant, fast, flexible architecture

13