Google Technology

Vera Goebel, Ifi/UiO, 2011

Search Engine: Crawler, PageRank, indexes
MapReduce
GFS – Google File System
Data Center

Google Big Table (slides + video)

Search Engine

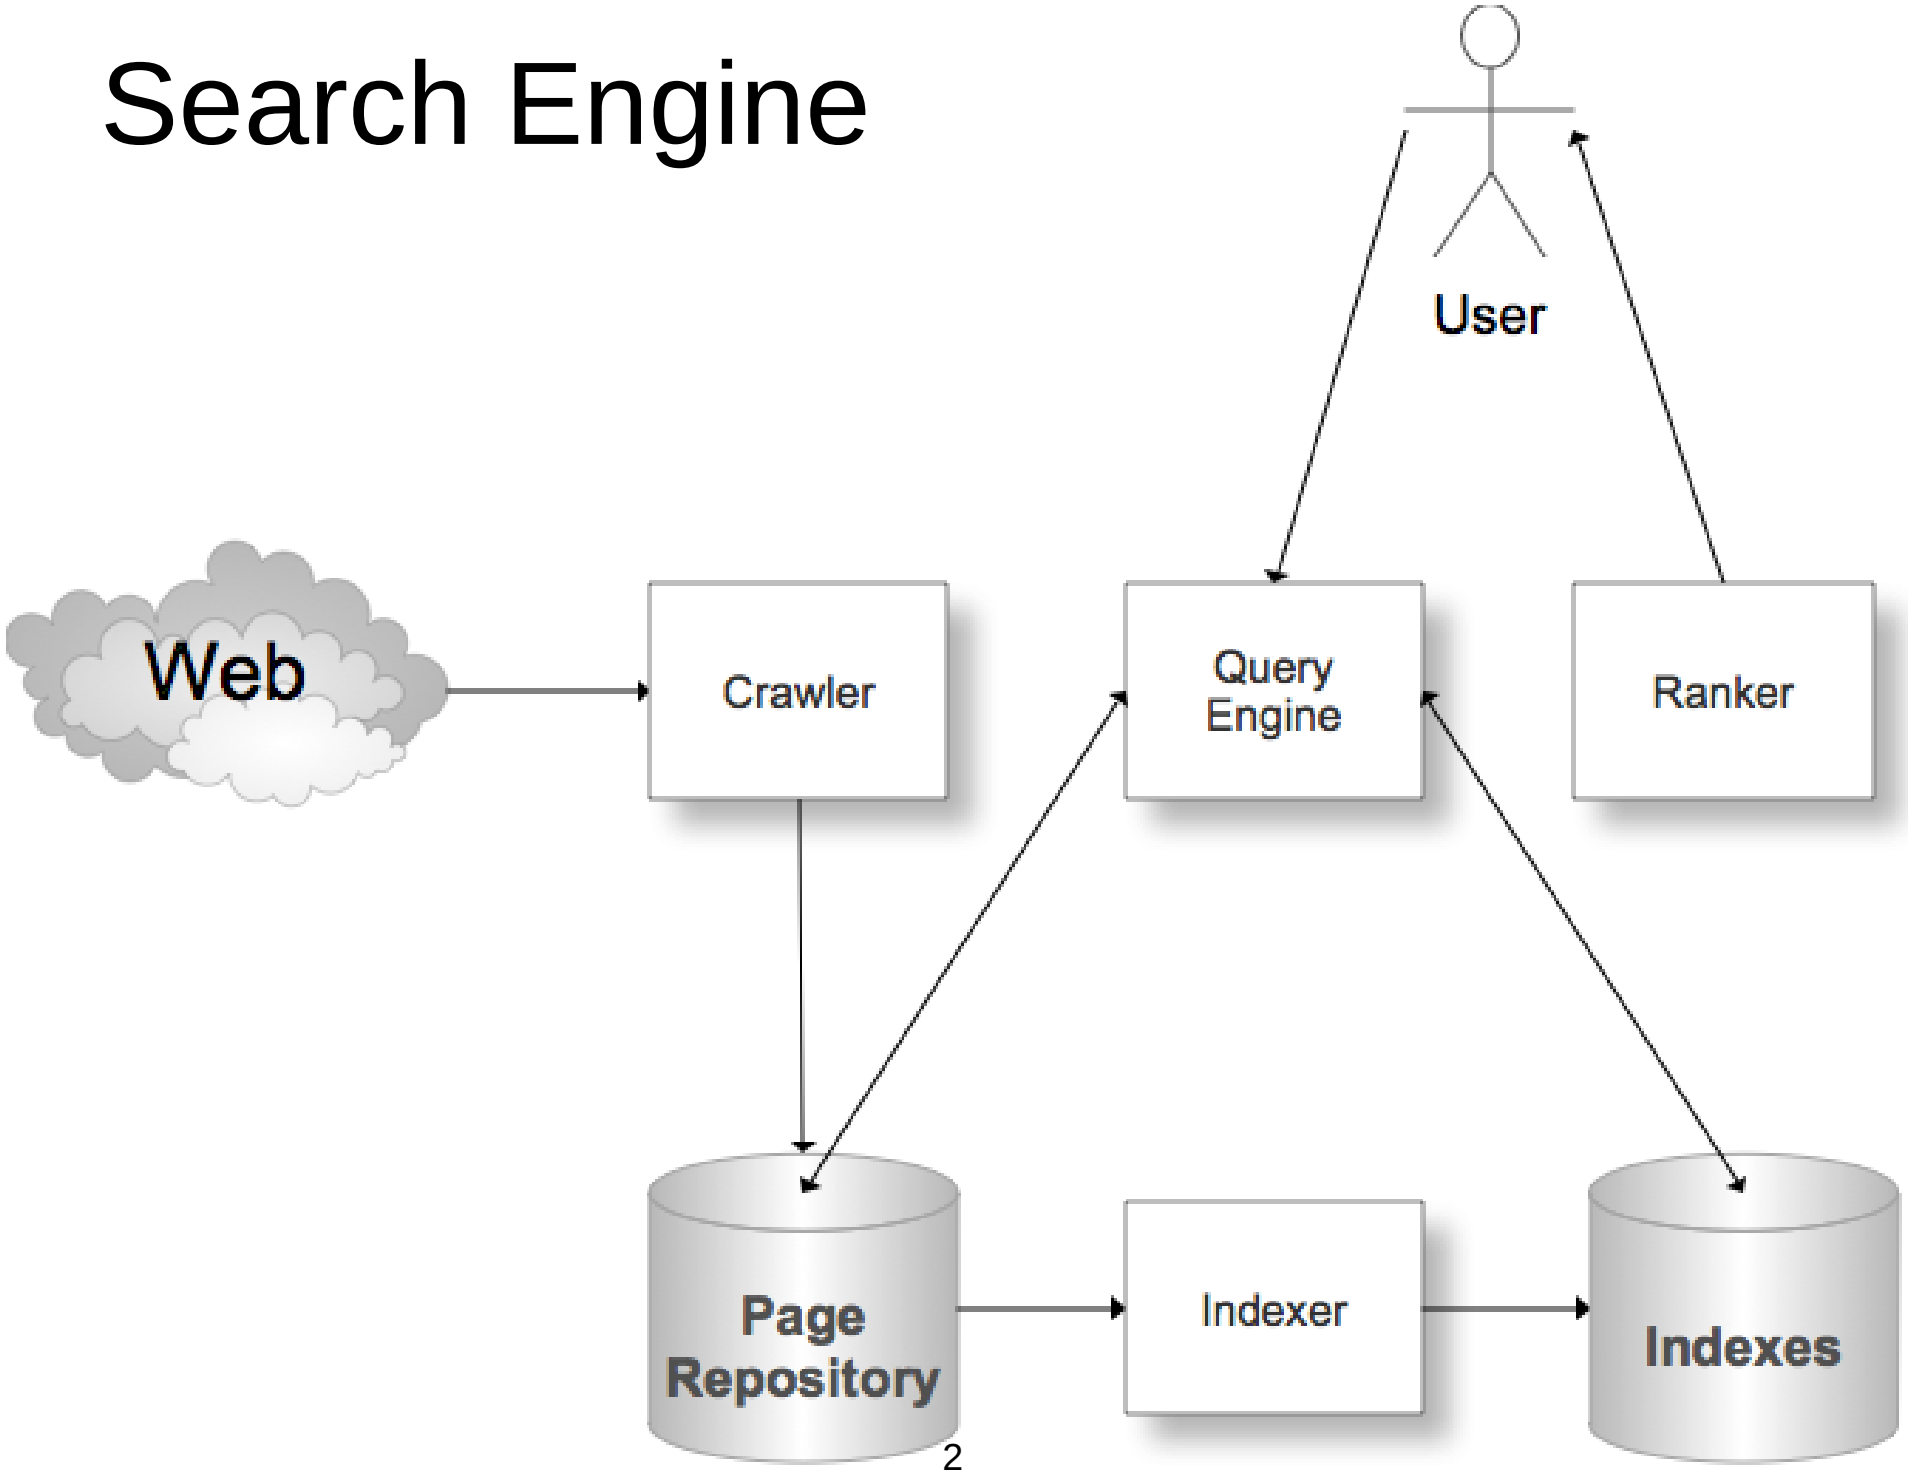

Crawler

- Process that downloads web pages to a Page Repository.
- Examine pages for links to other pages and insert the ones that are not in the Page Repository in the set for pages to be crawled.

http://goo.gl/gG3s

Crawler

Challenge	Description	Solution
Terminating search	Dynamically generated pages could create a forever loop	Limit number of pages to crawl with a "depth" limit per site
Managing the repository	1. Duplication of URL to be crawled 2. Duplicated pages due to mirror sites, different routes, plagiarism, etc.	1. An efficient index for checking stored pages 2. Minhash and locality-sensitive hashing signatures
Selecting the next page	How to prioritise next page to be crawled?	Give priority to "important" pages
Speeding up the crawl	1. How many processes should be simultaneously run? 2. How to synchronise them to avoid they crawl the same site. 3. Avoid DoS attack	1. Scale to several machines 2. Assign processes to entire hosts or sites 3. Do not issue frequent requests to a single site. Several processes in a single machine due to idle states.

Query Processing in Search Engines

- Search engine queries are not like SQL queries
- Require inverted indices
- Disk access is very expensive to offer the user acceptable response time
- Matched records are ranked before showing to the user

Recursive Formulation of Page Rank

Algorithm for identifying "important" pages: Web page is important if many important pages link to it.

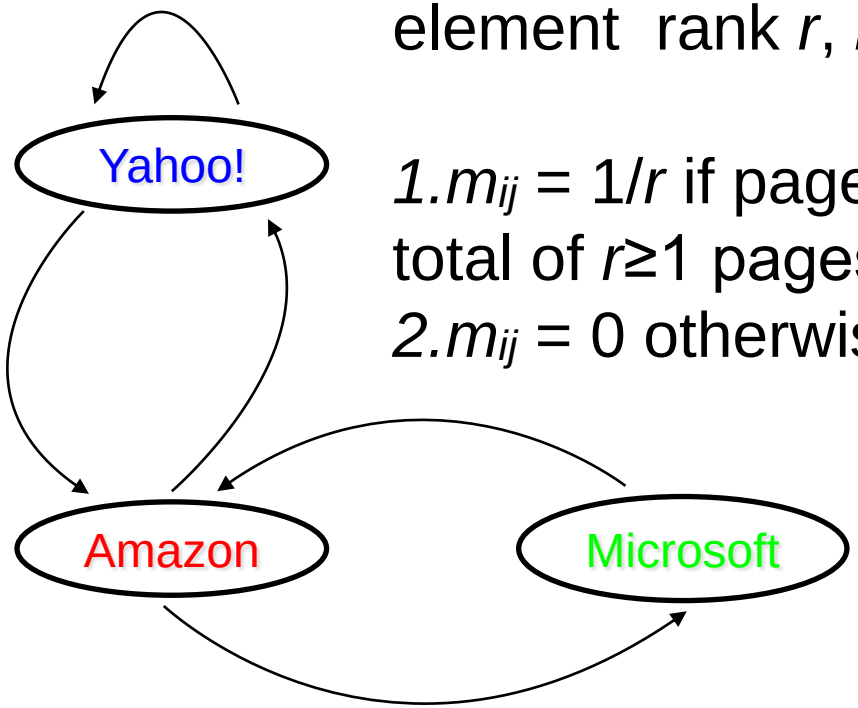

The Web in 1839

The Matrix M, the transition matrix of the Web has element rank r, m_{ij} in row i and column j, where

1. $m_{ij} = 1/r$ if page j has a link to page i, and there are a total of $r \geq 1$ pages that j links to

2. $m_{ij} = 0$ otherwise

Transition Matrix

$$M = \begin{array}{c|ccc|c} & 1/2 & 1/2 & 0 & \text{Yahoo!} \\ & 1/2 & 0 & 1 & \text{Amazon} \\ & 0 & 1/2 & 0 & \text{Microsoft} \\ \hline & \text{Yahoo!} & \text{Amazon} & \text{Microsoft} & \end{array}$$

Spider Traps and Dead Ends

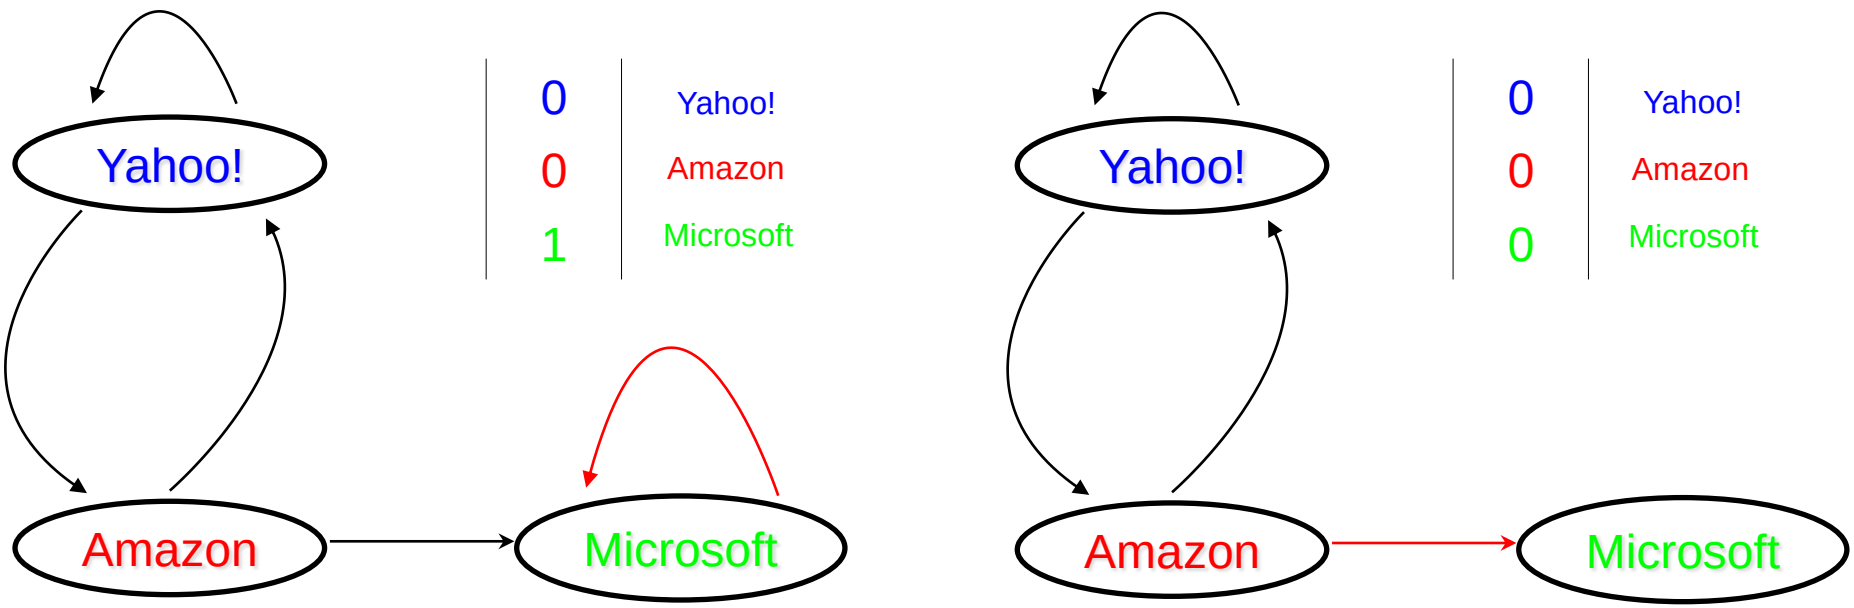

Microsoft becomes a spider trap

Microsoft becomes a dead end

Link Spam

- Spam farming in order to accumulate and concentrate PageRank on a few pages
- Links to the spam farm from pulicly accessible blogs, with messages like “I agree with you. See x1234.mySpam.Farm.com”

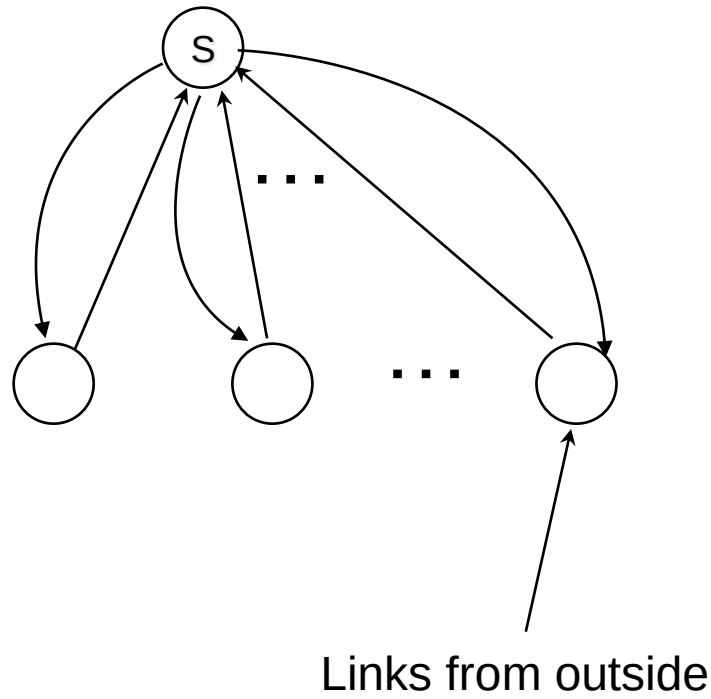

Inverted Indices

- Essential for Web Queries
- Uses indirect buckets for space efficiency

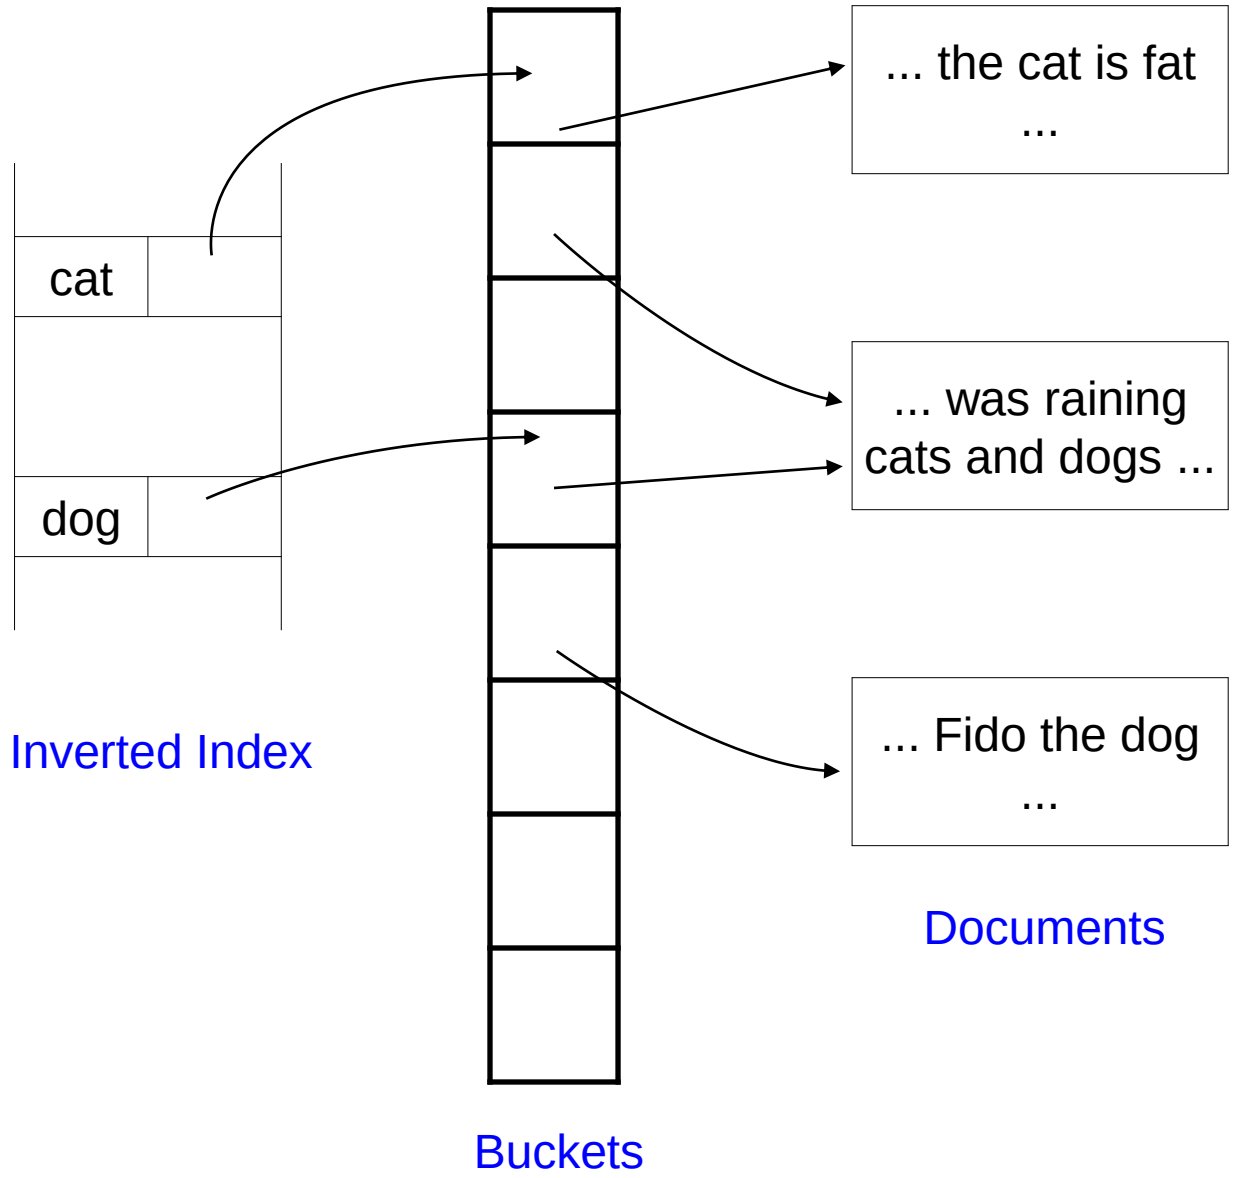

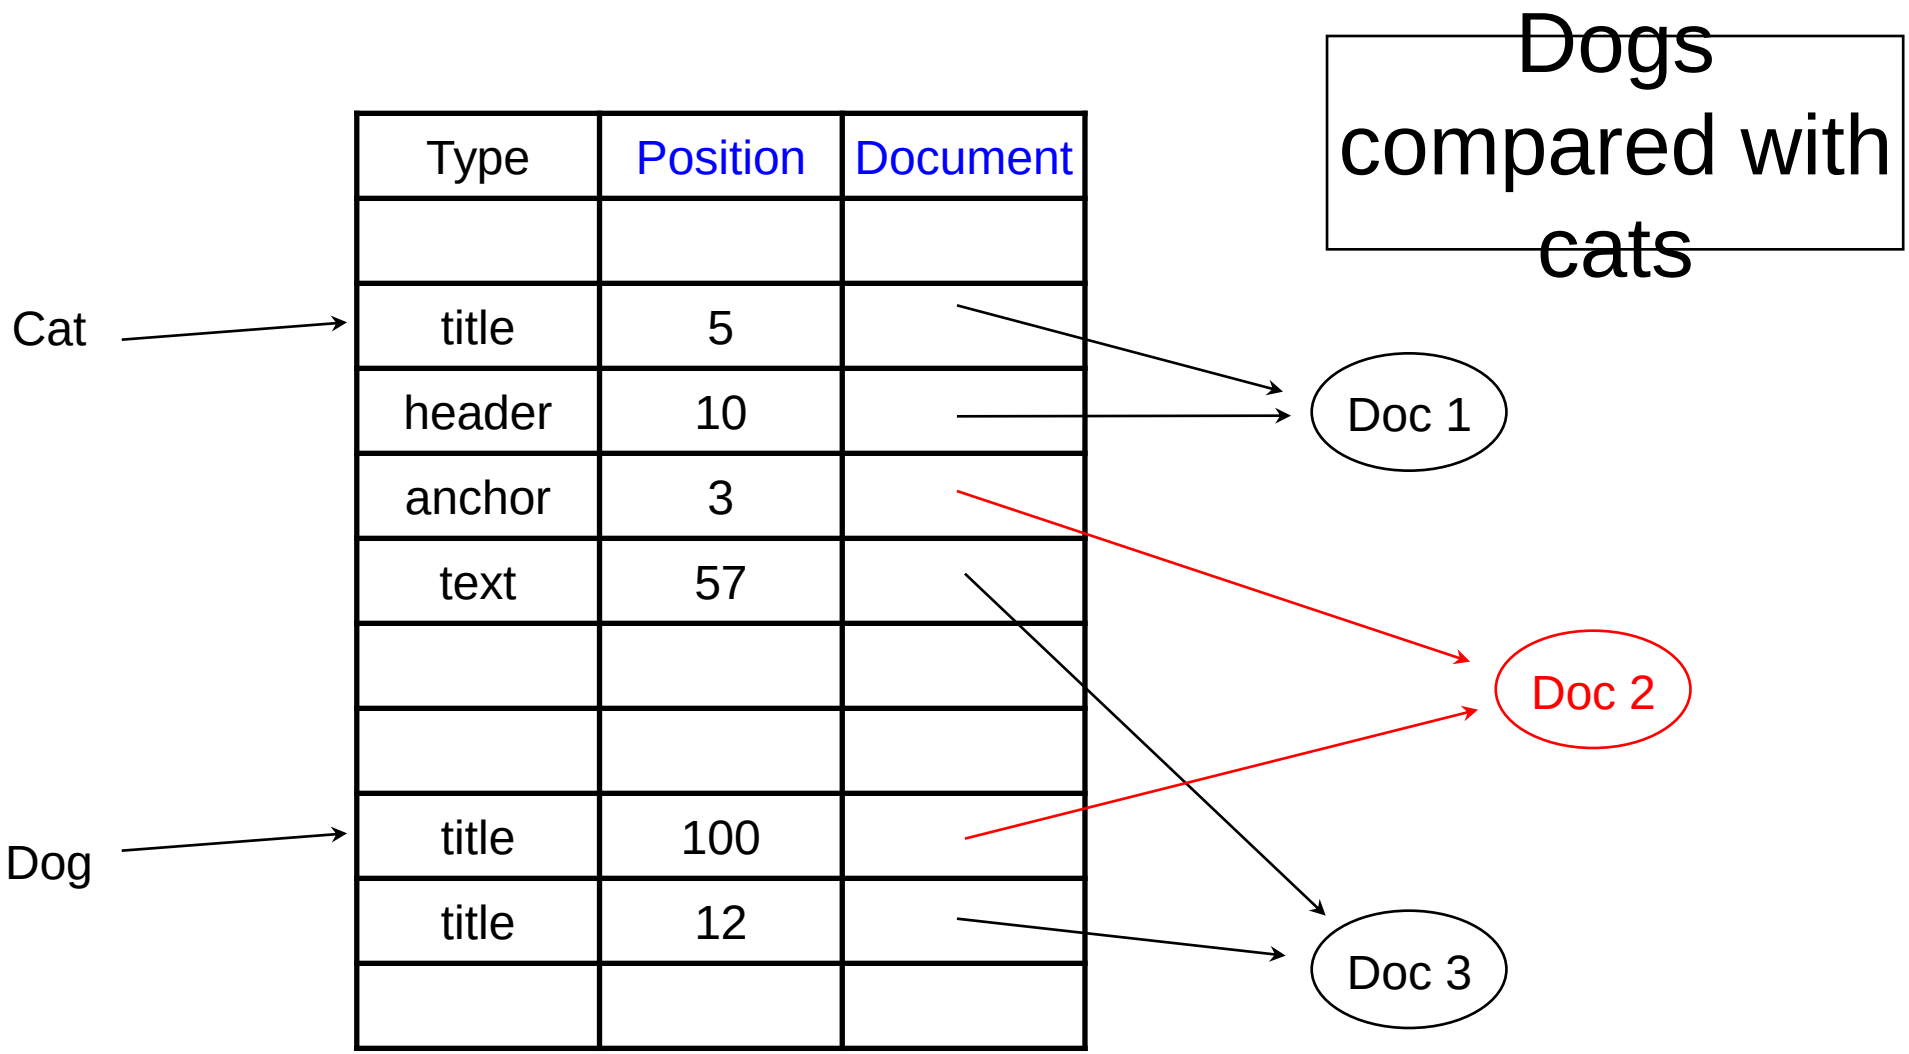

Sorting more information in the inverted index

Map-Reduce Parallelism Framework

- Large-scale parallel machines share high load operations such as joins
- Distributed architectures
  - Grid, networks and corporate DBs
- MRP paradigm expresses large-scale computations

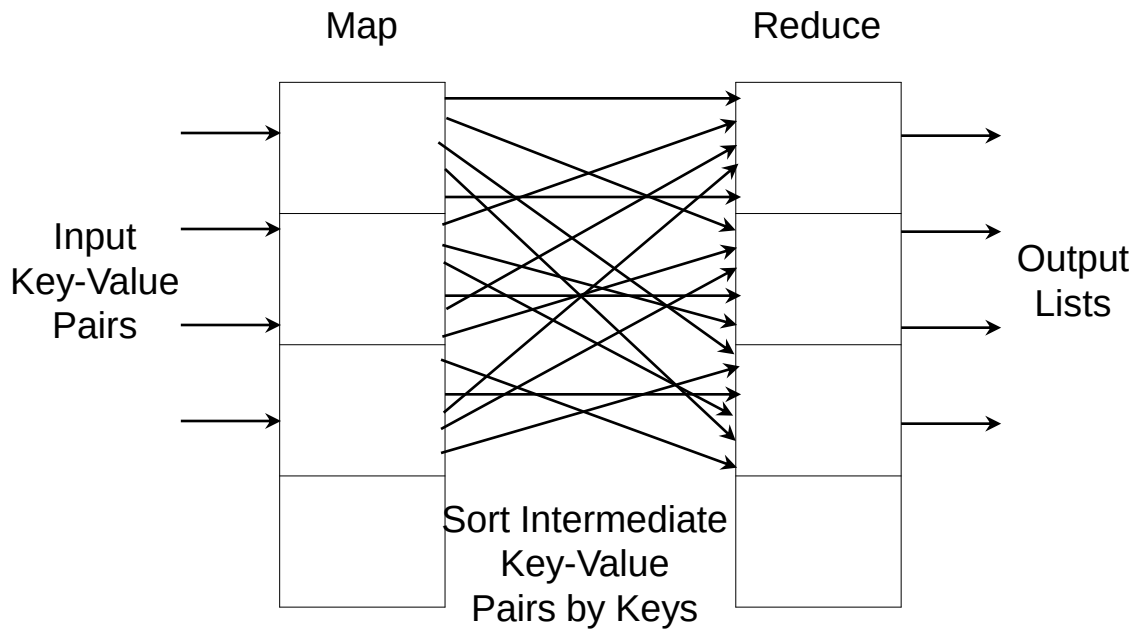

Execution of map and reduce functions

Google Search

- Uses: links, PageRank, anchors, proximity and visual presentation (e.g. bold text is weighted higher) in search logic.

1. Search the index
2. Analyze the web pages for relevance
3. Evaluate the site’s reputation
4. Rank the web pages

Google's System Anatomy

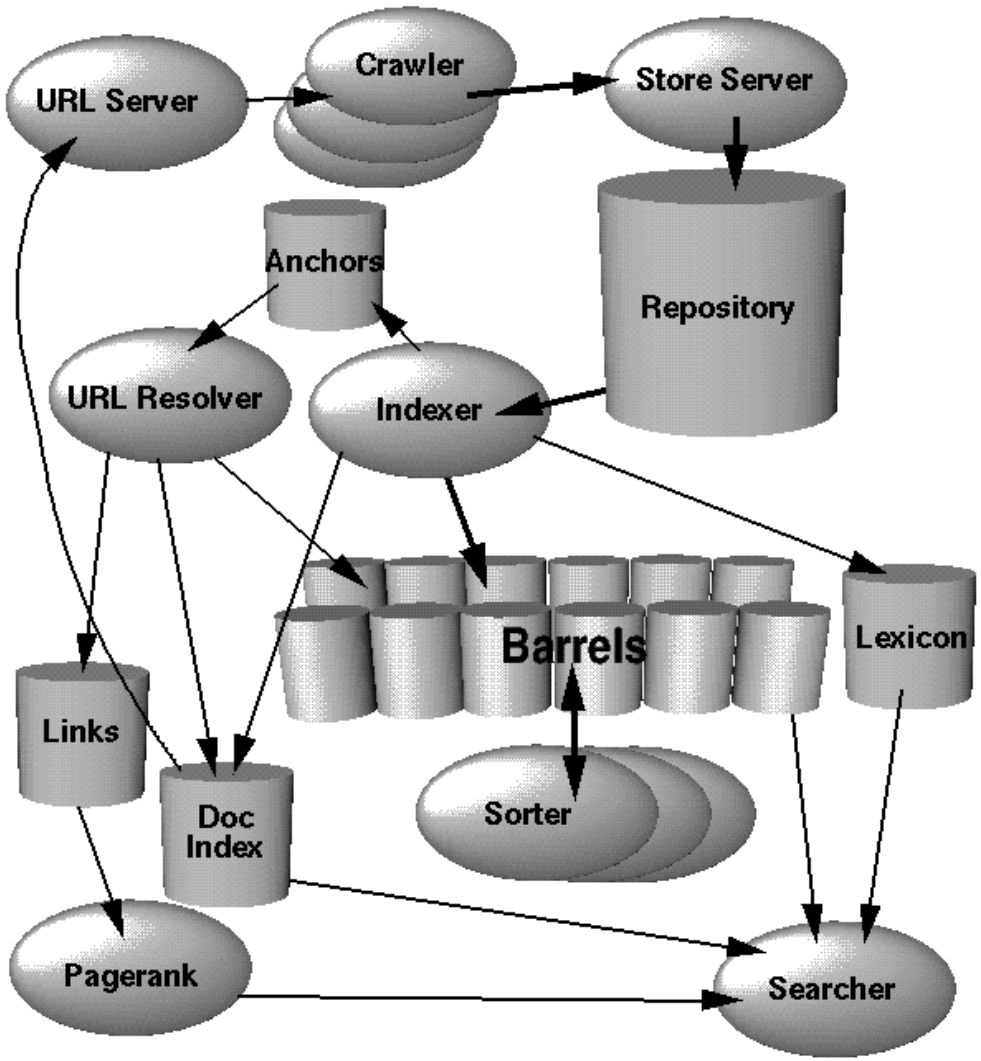

http://goo.gl/yYbb

Google File System - Motivation

- Nothing is small in Google land
  - Peta-bytes of data
  - Millions of users
  - Lots of services and servers

  → **Scalability**
- Failures are normal
  - Network connections
  - Hard disks
  - Power supplies

  → **Fault tolerance**
- Monitoring and maintenance is hard

  → **Autonomic computing**

- Clusters all over the world
- Thousands of queries served per second
  - One query reads hundreds of MB of data
  - One query consumes billions of CPU cycles
- A distributed, fault-tolerant file system is needed!

Google Data Centers

- Scaling out on commodity hardware is cheaper than scaling up on high-end servers
- Google servers:
  - > 15 000 servers (2003)
  - ~ 200 000 (2005)
  - ~ 1 M servers (2010)
- Data centers are composed of standard shipping containers with 1160 servers in each

Chunks & Chunk Servers

Chunk

- Similar to block in file systems
- Size is always 64 MB
- Less fragmentation
- Eases management
- Sent directly to clients

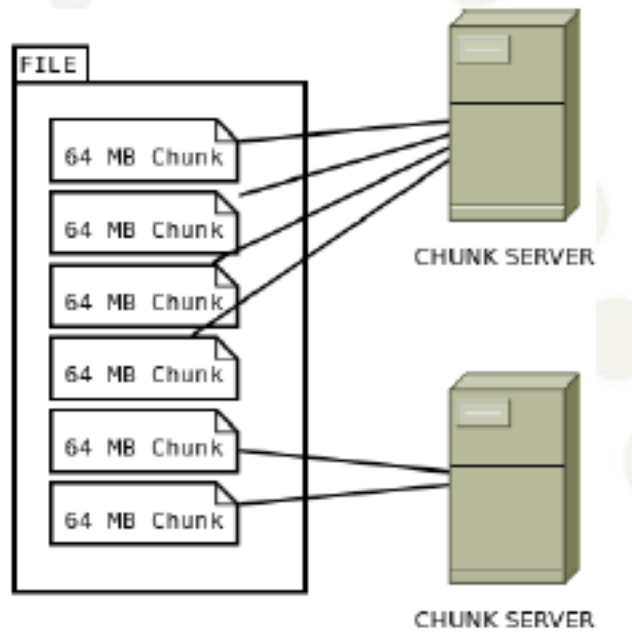

Master Servers

Master Server

- Coordinates cluster
- Updates operation log
- Stores meta-data

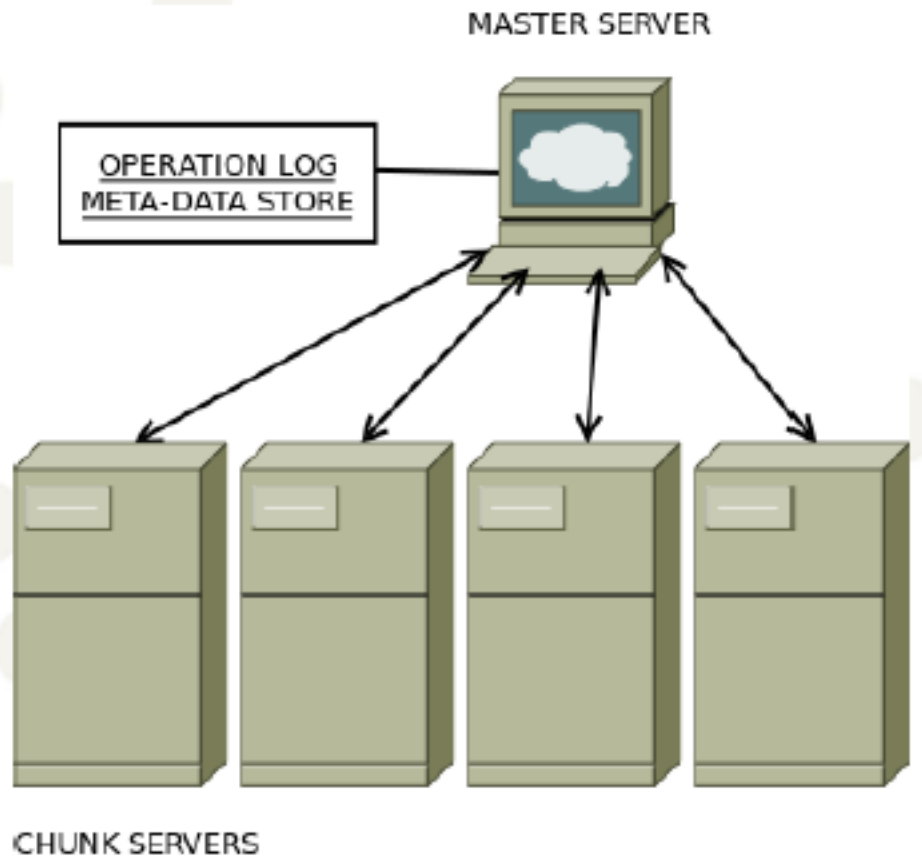

Master Server – Chunk Server Communication

State updates

- Is a chunk server down?
- Are there disk failures on a chunk server?
- Are any replicas corrupted?
- Which chunk replicas does a chunk server store?

Instructions

- Create new chunk
- Delete existing chunks

GFS - Architecture

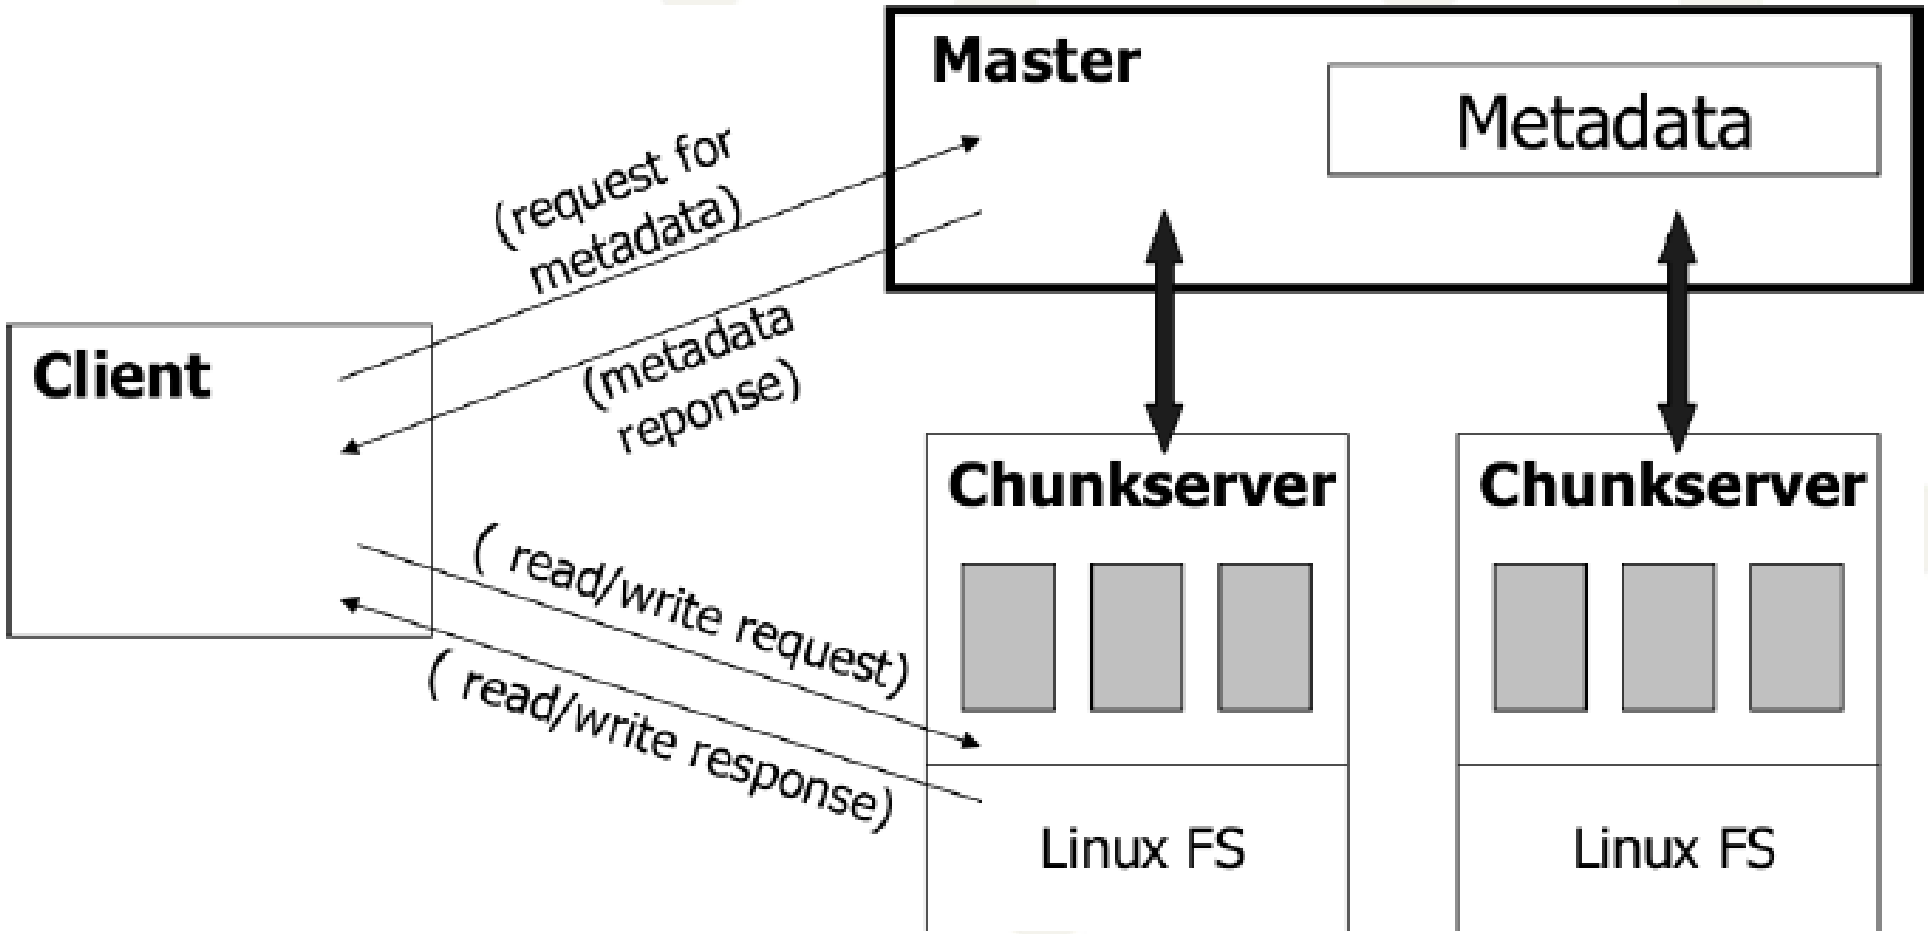

Read Operation

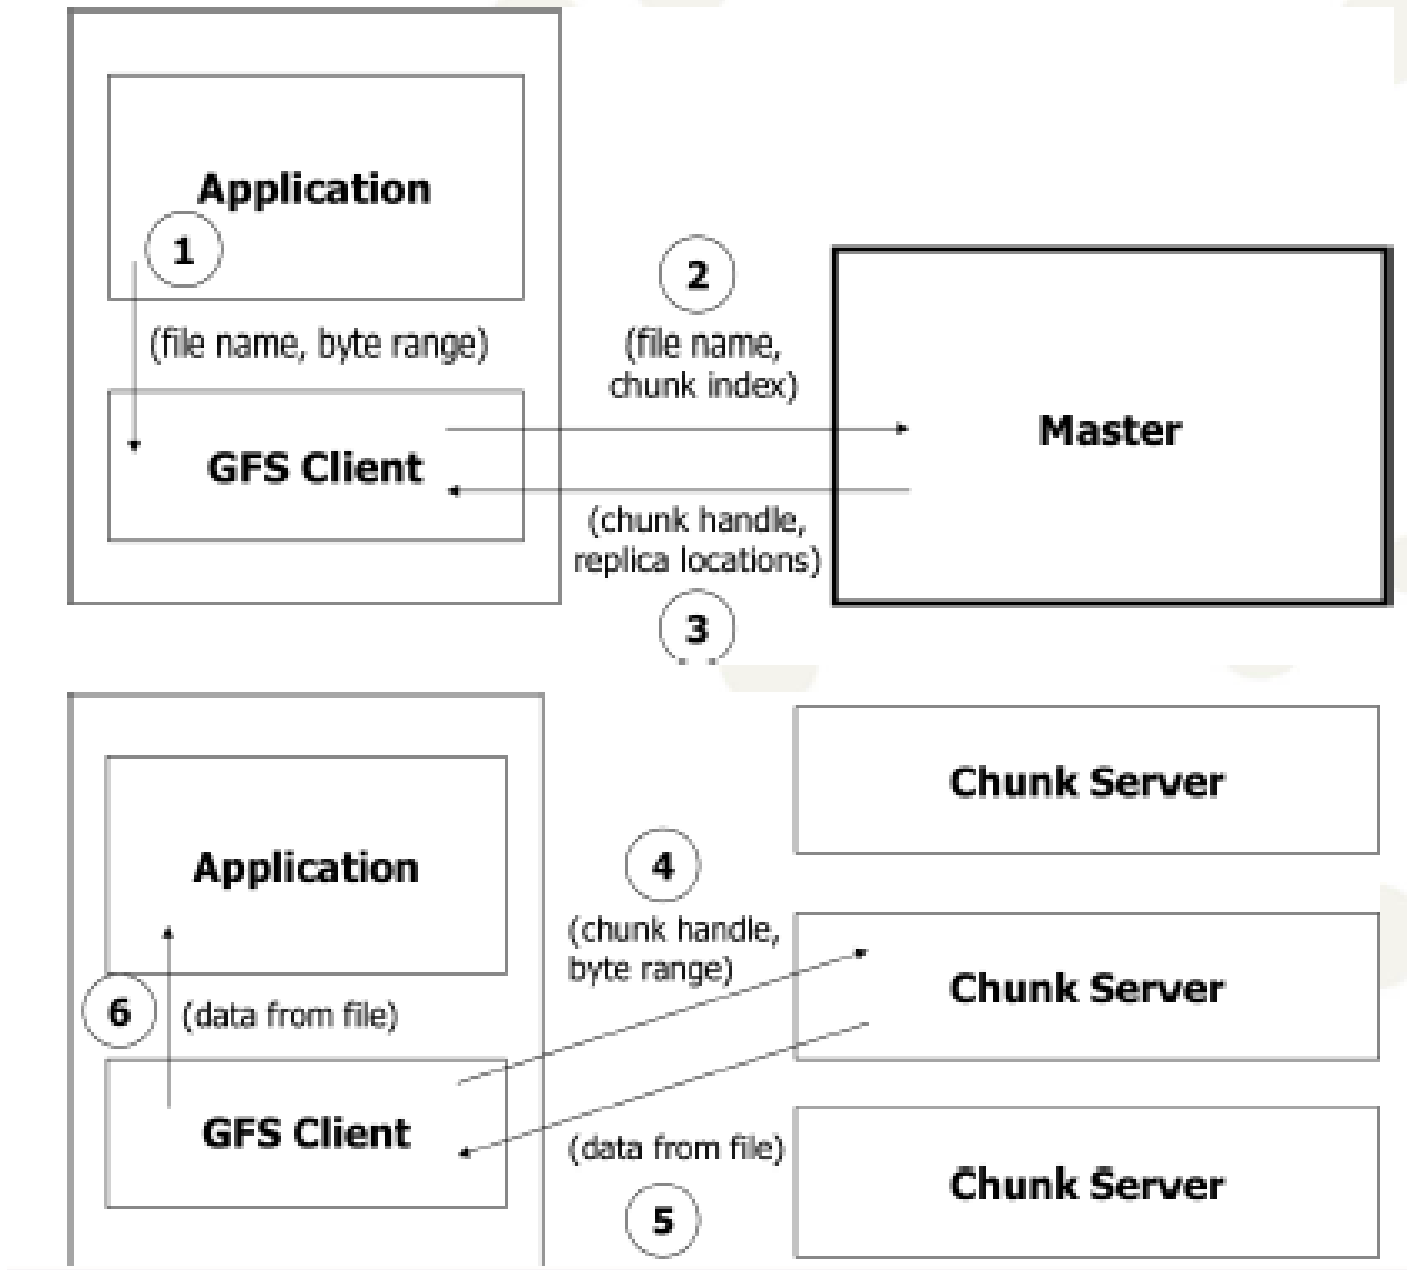

Write Operation

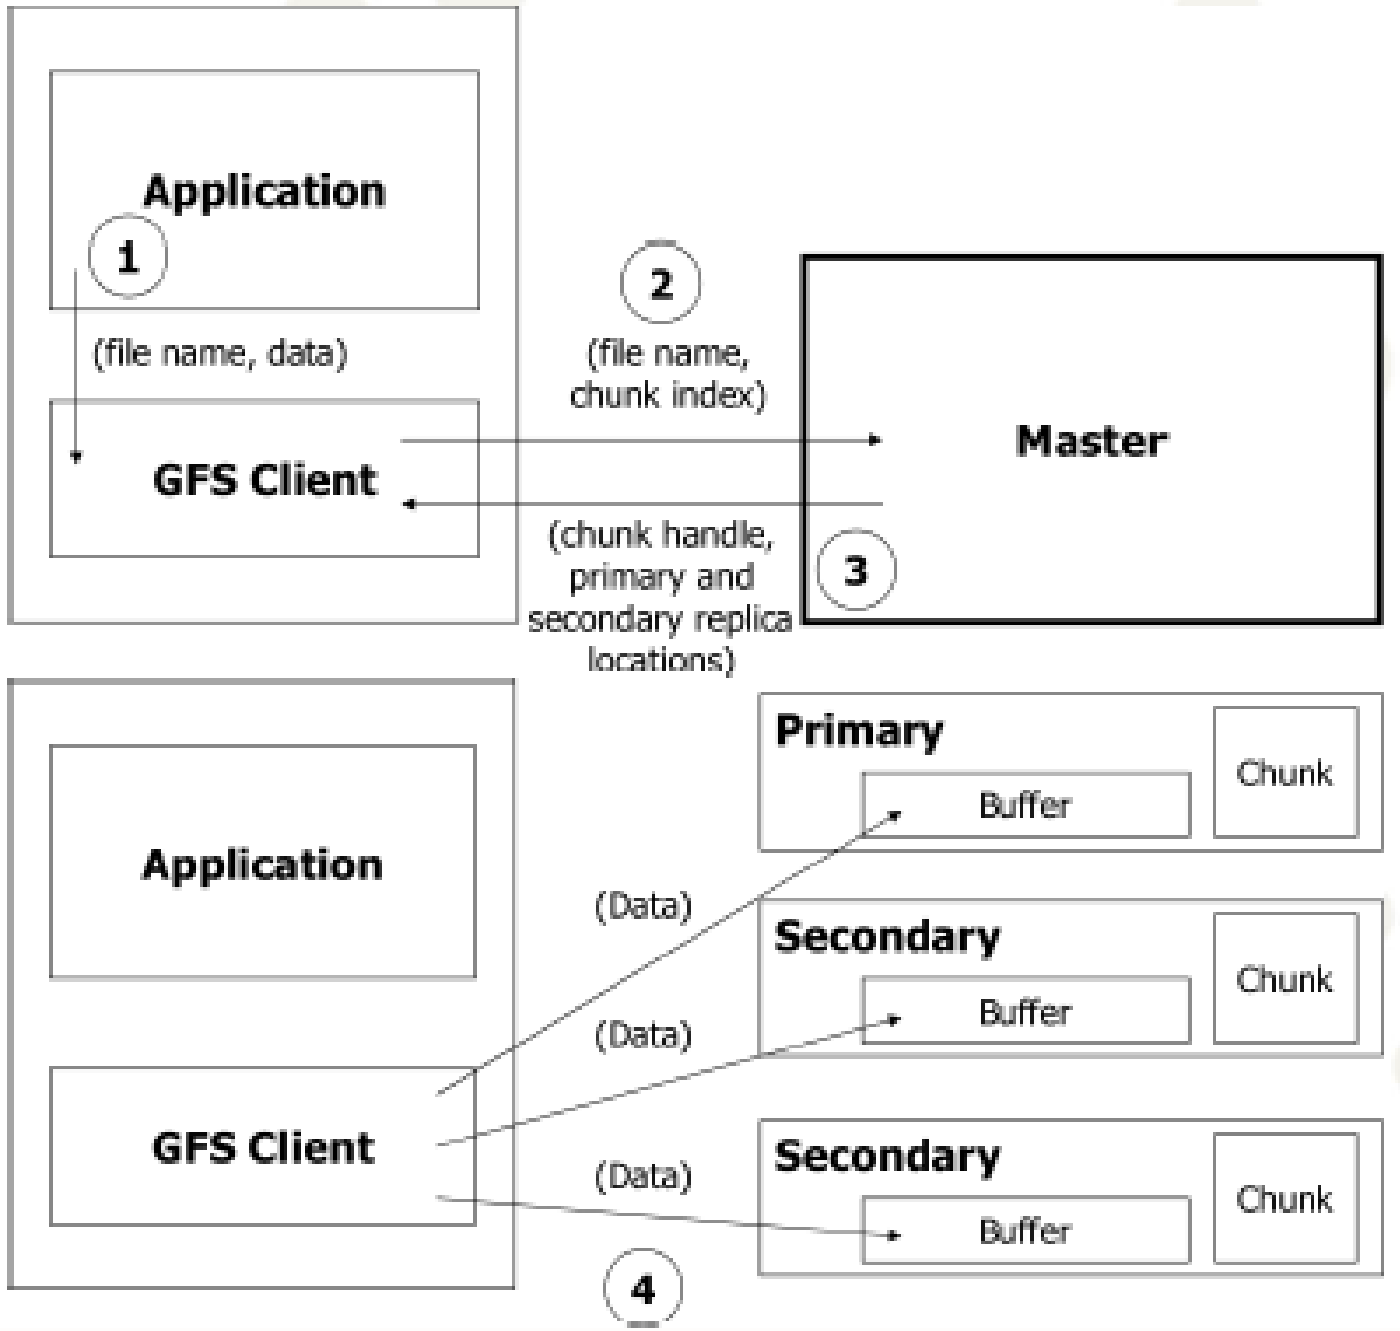

Write Operation (cont.)

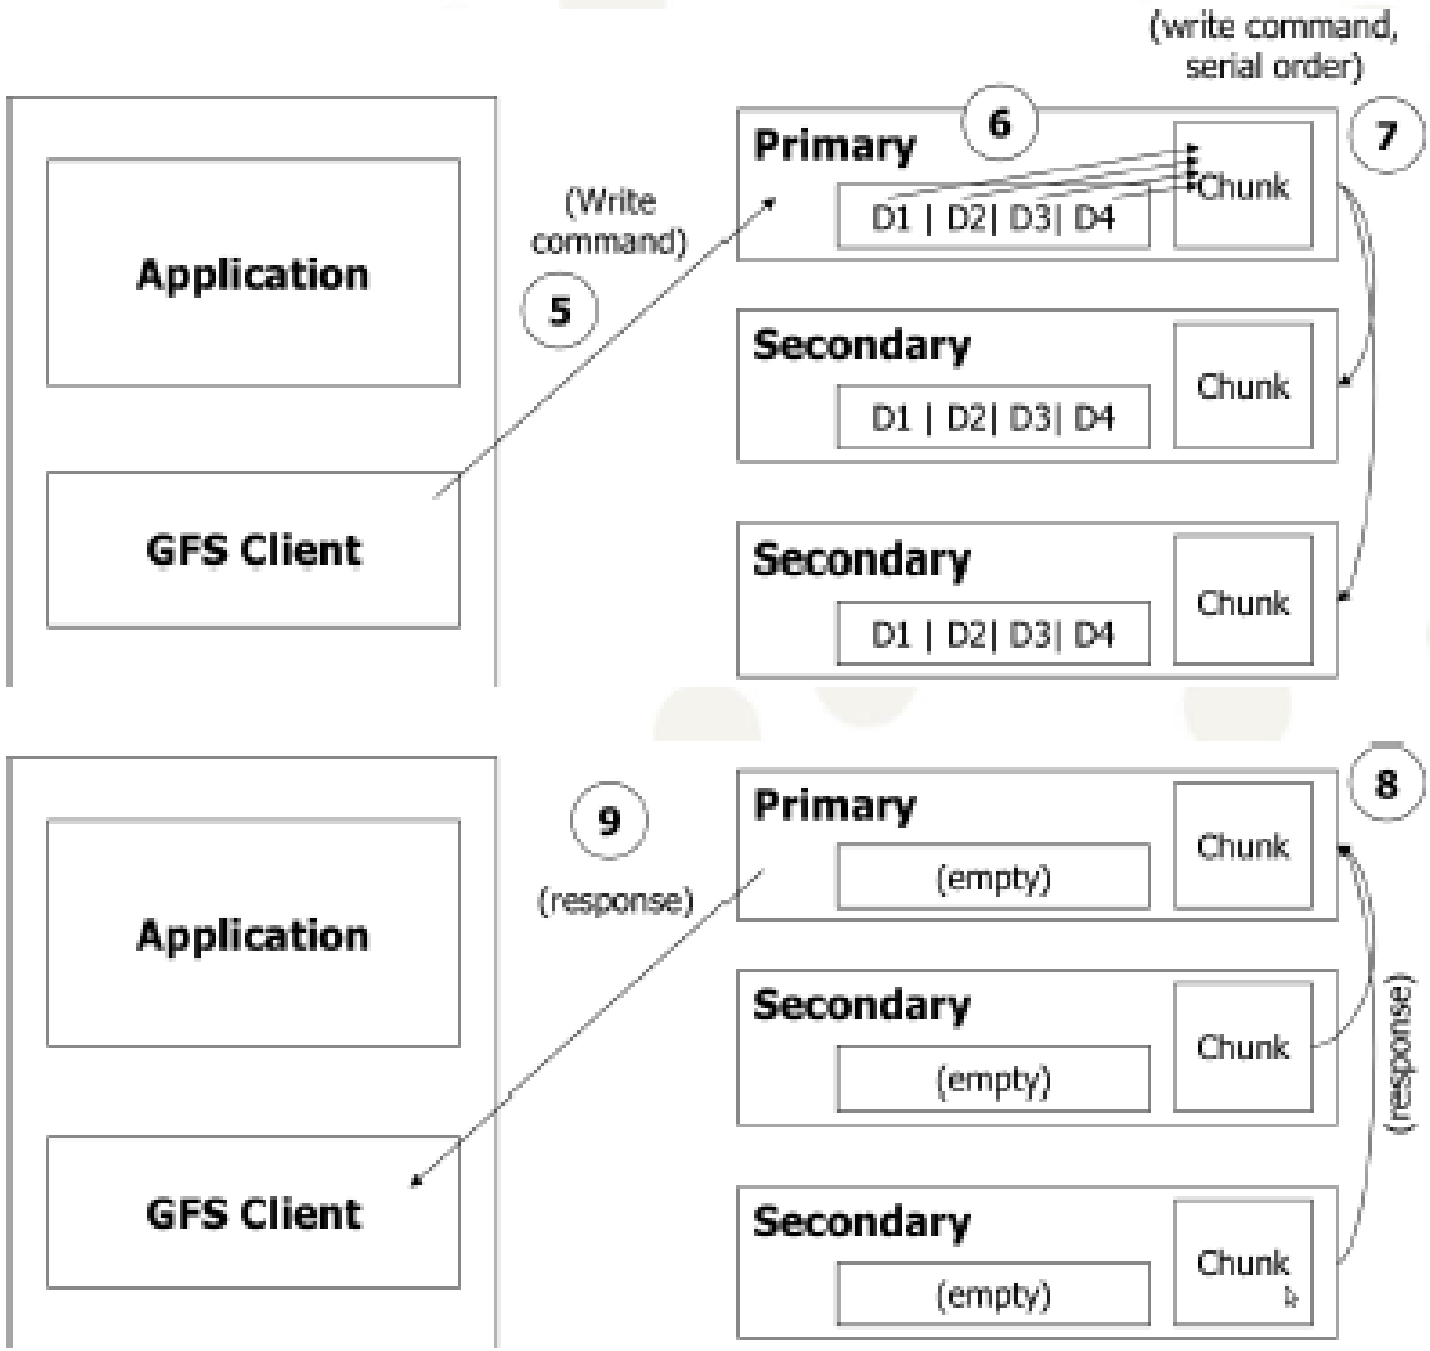

References

- The Anatomy of a Large-Scale Hypertextual Web Search Engine
- http://infolab.stanford.edu/~backrub/google.html
- Database Systems. The Complete Book. Second Edition. Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer Widom
- Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung, The Google File System, ACM Symposium on Operating Systems Principles, 2003
- Jonathan Strickland, How the Google File System Works, HowStuffWorks.com, 2010
- Wikipedia Contributors, Google File System, Wikipedia -The Free Encyclopedia, 2010