

# CSE 454 - Case Studies

## Design of Alta Vista

Based on a talk by Mike Burrows

5/19/2009 3:27 PM Copyright © 2000-2009 D.S.Weld 1

### CSE 454 Overview

Copyright © 2000-2009 D.S.Weld

### CSE 454 Overview

Copyright © 2000-2009 D.S.Weld

### Coming Soon...

- Today
  - (Me) **Alta Vista Case Study & Link Analysis**
- Thurs 5/21
  - (Benaloh) **Cryptography & Web Security**
- Tues 5/26
  - **No Class** **Group Meetings**
- Thurs 5/28
  - (Teevan) **Personalized Search**
- Tues 6/2
  - (Adar) **Zoetrope**
- Thurs 6/4
  - (Me) **Review & Human Computation**

5/19/2009 3:27 PM Copyright © 2000-2009 D.S.Weld 4

### Review

- **Vector Space Representation**
  - Dot Product as Similarity Metric
- **TF-IDF for Computing Weights**
  - $w_{ij} = f(i,j) * \log(N/n_i)$
  - Where  $q = \dots \text{word}_i \dots$
  - $N = |\text{docs}|$      $n_i = |\text{docs with word}_i|$
- **But How Process Efficiently?**

Copyright © 2000-2009 D.S.Weld 5

### Retrieval

#### Document-term matrix

	$t_1$	$t_2$	$\dots$	$t_j$	$\dots$	$t_m$	$nf$
$d_1$	$w_{11}$	$w_{12}$	$\dots$	$w_{1j}$	$\dots$	$w_{1m}$	$1/ d_1 $
$d_2$	$w_{21}$	$w_{22}$	$\dots$	$w_{2j}$	$\dots$	$w_{2m}$	$1/ d_2 $
$d_i$	$w_{i1}$	$w_{i2}$	$\dots$	$w_{ij}$	$\dots$	$w_{im}$	$1/ d_i $
$d_n$	$w_{n1}$	$w_{n2}$	$\dots$	$w_{nj}$	$\dots$	$w_{nm}$	$1/ d_n $

$w_{ij}$  is the weight of term  $t_j$  in document  $d_i$   
Most  $w_{ij}$ 's will be zero.

Copyright © 2000-2009 D.S.Weld 6

## Inverted Files for Multiple Documents

**LEXICON**

WORD	NDOCS	PTR
jezebel	20	→
jezer	3	→
jezerit	1	→
jeziah	1	→
jeziel	1	→
jeziah	1	→
jezoar	1	→
jezrahiah	1	→
jezreel	39	→

**OCURRENCE INDEX**

DOCID	OCCUR	POS 1	POS 2	...
34	6	1	118	2087
44	3	215	2291	3010
56	4	5	22	134
...	...	...	...	...
566	3	203	245	287
67	1	132		
...	...	...	...	...
107	4	322	354	381
232	6	15	195	248
677	1	481		
713	3	42	312	802

"jezebel" occurs  
6 times in document 34,  
3 times in document 44,  
4 times in document 56...

Copyright © 2000-2009 D.S.Weld

7

## Many Variations Possible

- **Address space (flat, hierarchical)**
  - Alta Vista uses flat approach
- **Record term-position information**
- **Precalculate TF-IDF info**
- **Stored header, font & tag info**
- **Compression strategies**

Copyright © 2000-2009 D.S.Weld

8

## AltaVista: Inverted Files

- **Map each word to list of locations where it occurs**
- **Words = null-terminated byte strings**
- **Locations = 64 bit unsigned ints**
  - Layer above gives interpretation for location
    - URL
    - Index into text specifying word number
- **Slides adapted from talk by Mike Burrows**


5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

9

## Documents

- **A document is a region of location space**
  - Contiguous
  - No overlap
  - Densely allocated (first doc is location 1)
- **All document structure encoded with words**
  - enddoc at last location of document
  - begintitle, endtitle mark document title



5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

10

## Format of Inverted Files

- **Words ordered lexicographically**
- **Each word followed by list of locations**
- **Common word prefixes are compressed**
- **Locations encoded as deltas**
  - Stored in as few bytes as possible
  - 2 bytes is common
  - Sneaky assembly code for operations on inverted files
    - Pack deltas into aligned 64 bit word
    - First byte contains continuation bits
    - Table lookup on byte => no branch instructs, no mispredicts
    - 35 parallelized instructions/ 64 bit word = 10 cycles/word
- **Index ~ 10% of text size**

5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

11

## Index Stream Readers (ISRs)

- **Interface for**
  - Reading result of query
  - Return ascending sequence of locations
  - Implemented using lazy evaluation
- **Methods**
  - loc(ISR)                   return current location
  - next(ISR)                 advance to next location
  - seek(ISR, X)             advance to next loc after X
  - prev(ISR)                 return previous location



5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

12

## Processing Simple Queries

- User searches for “mp3”
- Open ISR on “mp3”
  - Uses hash table to avoid scanning entire file
- Next(), next(), next()
  - returns locations containing the word


5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

13

## Combining ISRs

- **And** Compare locs on two streams
- **Or** Merges two or more ISRs
- **Not** Returns locations not in ISR (lazily)




5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

14

## What About File Boundaries?




5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

15

## ISR Constraint Solver

- **Inputs:**
  - Set of ISRs: A, B, ...
  - Set of Constraints
- **Constraint Types**
  - $loc(A) \leq loc(B) + K$
  - $prev(A) \leq loc(B) + K$
  - $loc(A) \leq prev(B) + K$
  - $prev(A) \leq prev(B) + K$
- **For example: phrase “a b”**
  - $loc(A) \leq loc(B), loc(B) \leq loc(A) + 1$




5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

16

## Two words on one page

- Let E be ISR for word enddoc
  - Constraints for conjunction a AND b
    - $prev(E) \leq loc(A)$
    - $loc(A) \leq loc(E)$
    - $prev(E) \leq loc(B)$
    - $loc(B) \leq loc(E)$
- What if prev(E) Undefined?*




5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

17

## Advanced Search

- **Field query: a in Title of page**
- Let BT, ET be ISRP of words begintitle, endtitle
- **Constraints:**
  - $loc(BT) \leq loc(A)$
  - $loc(A) \leq loc(ET)$
  - $prev(ET) \leq loc(BT)$



5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

18

## Implementing the Solver

### Constraint Types

- $\text{loc}(A) \leq \text{loc}(B) + K$
- $\text{prev}(A) \leq \text{loc}(B) + K$
- $\text{loc}(A) \leq \text{prev}(B) + K$
- $\text{prev}(A) \leq \text{prev}(B) + K$

5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

19

## Remember: Index Stream Readers

### Methods

- $\text{loc}(\text{ISR})$  return current location
- $\text{next}(\text{ISR})$  advance to next location
- $\text{seek}(\text{ISR}, X)$  advance to next loc after X
- $\text{prev}(\text{ISR})$  return previous location

5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

20

## Solver Algorithm

```
while (unsatisfied_constraints)
    satisfy_constraint(choose_unsat_constraint())
```

Heuristic:  
Which choice  
advances a  
stream the  
furthest?

- **To satisfy:**  $\text{loc}(A) \leq \text{loc}(B) + K$ 
  - Execute:  $\text{seek}(B, \text{loc}(A) - K)$
- **To satisfy:**  $\text{prev}(A) \leq \text{loc}(B) + K$ 
  - Execute:  $\text{seek}(B, \text{prev}(A) - K)$
- **To satisfy:**  $\text{loc}(A) \leq \text{prev}(B) + K$ 
  - Execute:  $\text{seek}(B, \text{loc}(A) - K)$ ,
  - $\text{next}(B)$
- **To satisfy:**  $\text{prev}(A) \leq \text{prev}(B) + K$ 
  - Execute:  $\text{seek}(B, \text{prev}(A) - K)$
  - $\text{next}(B)$

5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

21

## Update

- **Can't insert in the middle of an inverted file**
- **Must rewrite the entire file**
  - Naïve approach: need space for two copies
  - Slow since file is huge
- **Split data along two dimensions**
  - **Buckets** solve disk space problem
  - **Tiers** alleviate small update problem


5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

22

## Buckets & Tiers

- **Each word is hashed to a bucket**
- **Add new documents by adding a new tier**
  - Periodically merge tiers, bucket by bucket




5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

23

## What if Word Removed from Doc?

- **Delete documents by adding deleted word**
- 
- deleted  
|  
□ □ □ □ □  
|  
a
- **Expunge deletions when merging tier 1**



5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

24

## Scaling

- **How handle huge traffic?**
  - AltaVista Search ranked #16
  - 10,674,000 unique visitors (Dec'99)
- **Scale across N hosts**
  1. Ubiquitous index. Query one host
  2. Split N ways. Query all, merge results
  3. Ubiquitous index. Host handles subrange of locations. Query all, merge results
  4. Hybrids


5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

25

## AltaVista Structure

- **Front ends**
  - Alpha workstations
- **Back ends**
  - 4-10 CPU Alpha servers
    - 8GB RAM, 150GB disk
  - Organized in groups of 4-10 machines
    - Each with  $1/N^{\text{th}}$  of index



5/19/2009 3:27 PM

Copyright © 2000-2009 D.S.Weld

26