

CMU SCS

15-826: Multimedia Databases and Data Mining

Primary key indexing – hashing
C. Faloutsos

CMU SCS

Outline

Goal: 'Find similar / interesting things'

- Intro to DB
- ➔ • Indexing - similarity search
- Data Mining

15-826 Copyright: C. Faloutsos (2007) 2

CMU SCS

Indexing - Detailed outline

- primary key indexing
 - B-trees and variants
- ➔ • (static) hashing
 - extendible hashing
- secondary key indexing
- spatial access methods
- text
- ...

15-826 Copyright: C. Faloutsos (2007) 3

CMU SCS

(Static) Hashing

Problem: “find EMP record with ssn=123”
 What if disk space was free, and time was at premium?

15-826 Copyright: C. Faloutsos (2007) 4

CMU SCS

Hashing

A: Brilliant idea: key-to-address transformation:

15-826 Copyright: C. Faloutsos (2007) 5

CMU SCS

Hashing

Since space is NOT free:

- use M , instead of 999,999,999 slots
- hash function: $h(key) = slot-id$

15-826 Copyright: C. Faloutsos (2007) 6

CMU SCS

Hashing

Typically: each hash bucket is a page, holding many records:

15-826 Copyright: C. Faloutsos (2007) 7

CMU SCS

Hashing

Notice: could have **clustering**, or non-clustering versions:

15-826 Copyright: C. Faloutsos (2007) 8

CMU SCS

Hashing

Notice: could have clustering, or **non-clustering** versions:

15-826 Copyright: C. Faloutsos (2007) 9

CMU SCS

Hashing - design decisions?

- eg., IRS, 200M tax returns, by SSN

15-826 Copyright: C. Faloutsos (2007) 10

CMU SCS

Indexing- overview

- B-trees
- hashing
 - ➔ – hashing functions
 - size of hash table
 - collision resolution
- Hashing vs B-trees
- Indices in SQL

15-826 Copyright: C. Faloutsos (2007) 11

CMU SCS

Design decisions

- 1) formula $h()$ for hashing function
- 2) size of hash table M
- 3) collision resolution method

15-826 Copyright: C. Faloutsos (2007) 12

CMU SCS

Design decisions - functions

- Goal: **uniform** spread of keys over hash buckets
- Popular choices:
 - Division hashing
 - Multiplication hashing

15-826 Copyright: C. Faloutsos (2007) 13

CMU SCS

Division hashing

$$h(x) = (a*x+b) \bmod M$$

- eg., $h(ssn) = (ssn) \bmod 1,000$
 - gives the last three digits of ssn
- M : size of hash table - choose a prime number, defensively (why?)

15-826 Copyright: C. Faloutsos (2007) 14

CMU SCS

Division hashing

- eg., $M=2$; hash on driver-license number (dln), where last digit is 'gender' (0/1 = M/F)
- in an army unit with predominantly male soldiers
- Thus: avoid cases where M and keys have common divisors - prime M guards against that!

15-826 Copyright: C. Faloutsos (2007) 15

CMU SCS

Multiplication hashing

$h(x) = [\text{fractional-part-of} (x * \phi)] * M$

- ϕ : golden ratio ($0.618\dots = (\text{sqrt}(5)-1)/2$)
- in general, we need an irrational number
- advantage: M need not be a prime number
- but ϕ must be irrational

15-826 Copyright: C. Faloutsos (2007) 16

CMU SCS

Other hashing functions

- quadratic hashing (bad)
- ...
- conclusion: use division hashing

15-826 Copyright: C. Faloutsos (2007) 17

CMU SCS

Design decisions

- 1) formula $h()$ for hashing function
- ➡ 2) size of hash table M
- 3) collision resolution method

15-826 Copyright: C. Faloutsos (2007) 18

CMU SCS

Size of hash table

- eg., 50,000 employees, 10 employee-records / page
- Q: $M=??$ pages/buckets/slots

15-826 Copyright: C. Faloutsos (2007) 19

CMU SCS

Size of hash table

- eg., 50,000 employees, 10 employees/page
- Q: $M=??$ pages/buckets/slots
- A: utilization ~ 90% and
– M : prime number

Eg., in our case: $M = \text{closest prime to } 50,000/10 / 0.9 = 5,555$

15-826 Copyright: C. Faloutsos (2007) 20

CMU SCS

Design decisions

- 1) formula $h()$ for hashing function
- 2) size of hash table M
- ➡ 3) collision resolution method

15-826 Copyright: C. Faloutsos (2007) 21

CMU SCS

Collision resolution

- Q: what is a 'collision'?
- A: ??

15-826 Copyright: C. Faloutsos (2007) 22

CMU SCS

Collision resolution

123; Smith; Main str.

15-826 Copyright: C. Faloutsos (2007) 23

CMU SCS

Collision resolution

- Q: what is a 'collision'?
- A: ??
- Q: why worry about collisions/overflows?
(recall that buckets are ~90% full)

15-826 Copyright: C. Faloutsos (2007) 24

CMU SCS

Collision resolution

- Q: what is a 'collision'?
- A: ??
- Q: why worry about collisions/overflows?
(recall that buckets are ~90% full)
- A: 'birthday paradox'

15-826 Copyright: C. Faloutsos (2007) 25

CMU SCS

Collision resolution

- open addressing
 - linear probing (ie., put to next slot/bucket)
 - re-hashing
- separate chaining (ie., put links to overflow pages)

15-826 Copyright: C. Faloutsos (2007) 26

CMU SCS

Collision resolution

linear probing:

15-826 Copyright: C. Faloutsos (2007) 27

CMU SCS

Collision resolution

re-hashing

123; Smith; Main str.

$h1()$

$h2()$

#0 page

#h(123)

M

15-826 Copyright: C. Faloutsos (2007) 28

CMU SCS

Collision resolution

separate chaining

123; Smith; Main str.

FULL

15-826 Copyright: C. Faloutsos (2007) 29

CMU SCS

Design decisions - conclusions

- function: division hashing
 - $h(x) = (a*x+b) \text{ mod } M$
- size M : ~90% util.; prime number.
- collision resolution: separate chaining
 - easier to implement (deletions!);
 - no danger of becoming full

15-826 Copyright: C. Faloutsos (2007) 30

CMU SCS

Indexing- overview

- B-trees
- hashing
 - ➔ – Hashing vs B-trees
- Indices in SQL
- extendible hashing

15-826 Copyright: C. Faloutsos (2007) 31

CMU SCS

Hashing vs B-trees:

Hashing offers

- speed ! ($O(1)$ avg. search time)

..but:

15-826 Copyright: C. Faloutsos (2007) 32

CMU SCS

Hashing vs B-trees:

..but B-trees give:

- key ordering:
 - range queries
 - proximity queries
 - sequential scan
- $O(\log(N))$ guarantees for search, ins./del.
- graceful growing/shrinking

15-826 Copyright: C. Faloutsos (2007) 33

CMU SCS

Hashing vs B-trees:

thus:

- B-trees are implemented in most systems

footnotes:

- hashing is rarely implemented (why not?)
- 'dbm' and 'ndbm' of UNIX: offer one or both

15-826 Copyright: C. Faloutsos (2007) 34

CMU SCS

Indexing- overview

- B-trees
- hashing
 - Hashing vs B-trees
- ➔ • Indices in SQL
- extendible hashing

15-826 Copyright: C. Faloutsos (2007) 35

CMU SCS

Indexing in SQL

- create index **<index-name>** on **<relation-name>** (**<attribute-list>**)
- create unique index **<index-name>** on **<relation-name>** (**<attribute-list>**)
- drop index **<index-name>**

15-826 Copyright: C. Faloutsos (2007) 36

CMU SCS

Indexing in SQL

- eg.,
create index ssn-index
on STUDENT (ssn)
- or (eg., on *TAKES(ssn, cid, grade)*):
create index sc-index
on TAKES (ssn, c-id)

15-826 Copyright: C. Faloutsos (2007) 37

CMU SCS

Indexing- overview

- B-trees
- hashing
- Indices in SQL
- extensible hashing
- ➔ – ‘extensible’ hashing [Fagin, Pipenger +]
- ‘**linear**’ hashing [Litwin]

15-826 Copyright: C. Faloutsos (2007) 38

CMU SCS

Problem with static hashing

- problem: overflow?
- problem: underflow? (underutilization)

15-826 Copyright: C. Faloutsos (2007) 39

CMU SCS

Solution: Dynamic/extendible hashing

- idea: shrink / expand hash table on demand..
- ..dynamic hashing

Details: how to grow gracefully, on overflow?

Many solutions - One of them: 'extendible hashing' [Fagin et al]

15-826 Copyright: C. Faloutsos (2007) 40

CMU SCS

Extendible hashing

123; Smith; Main str.

→

15-826 Copyright: C. Faloutsos (2007) 41

CMU SCS

Extendible hashing

solution:
split the bucket in two

123; Smith; Main str.

→

15-826 Copyright: C. Faloutsos (2007) 42

CMU SCS

Extendible hashing

in detail:

- keep a directory, with ptrs to hash-buckets
- Q: how to divide contents of bucket in two?
- A: hash each key into a very long bit string; keep only as many bits as needed

Eventually:

15-826 Copyright: C. Faloutsos (2007) 43

CMU SCS

Extendible hashing

directory

00...
01...
10...
11...

0001...
0111...

10101...
10011...
10110...

1101...

101001...

15-826 Copyright: C. Faloutsos (2007) 44

CMU SCS

Extendible hashing

directory

00...
01...
10...
11...

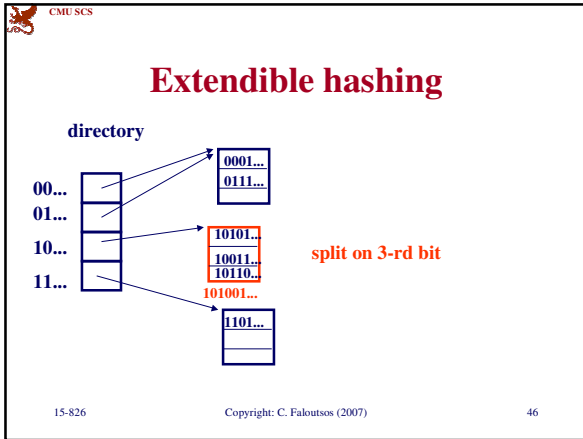
0001...
0111...

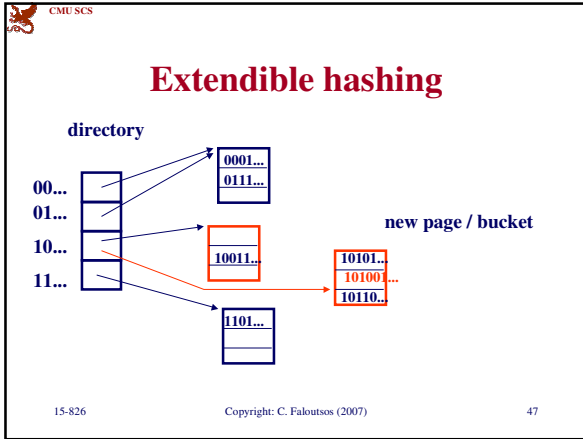
10101...
10011...
10110...

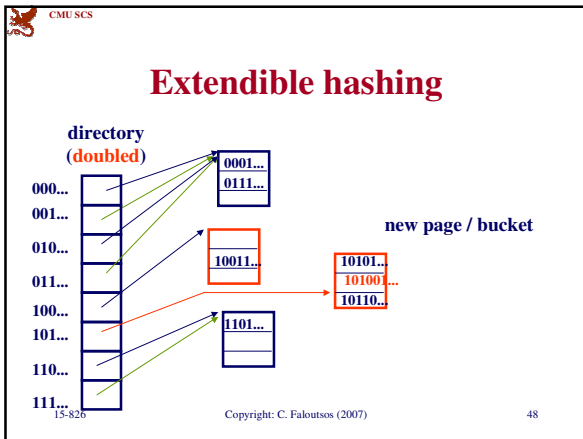
1101...

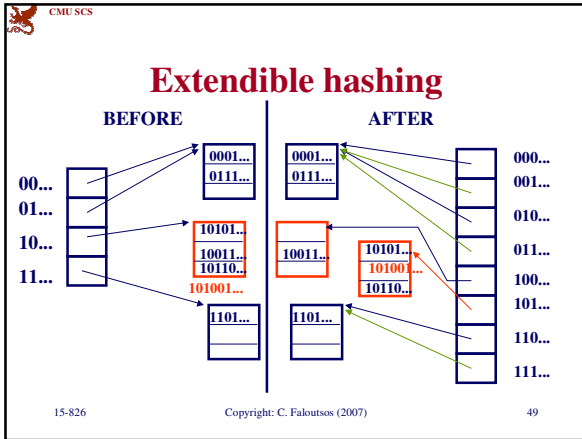
101001...

15-826 Copyright: C. Faloutsos (2007) 45









Extensible hashing

- Summary: directory doubles on demand
- or halves, on shrinking files
- needs 'local' and 'global' depth

15-826 Copyright: C. Faloutsos (2007) 50

Indexing- overview

- B-trees
- hashing
- Hashing vs B-trees
- Indices in SQL
- extensible hashing
 - 'extensible' hashing [Fagin, Pipenger +]
 - ➔ – 'linear' hashing [Litwin]

15-826 Copyright: C. Faloutsos (2007) 51

CMU SCS

Linear hashing - overview

- Motivation
- main idea
- search algo
- insertion/split algo
- deletion
- performance analysis
- variations

15-826 Copyright: C. Faloutsos (2007) 52

CMU SCS

Linear hashing

Motivation: ext. hashing needs directory etc etc; which doubles (ouch!)

Q: can we do something simpler, with smoother growth?

15-826 Copyright: C. Faloutsos (2007) 53

CMU SCS

Linear hashing

Motivation: ext. hashing needs directory etc etc; which doubles (ouch!)

Q: can we do something simpler, with smoother growth?

A: split buckets from left to right, **regardless** of which one overflowed ('crazy', but it works well!) - Eg.:

15-826 Copyright: C. Faloutsos (2007) 54

CMU SCS

Linear hashing

Initially: $h(x) = x \text{ mod } N$ (N=4 here)
 Assume capacity: 3 records / bucket
 Insert key '17'

bucket- id 0 1 2 3

4	8	5	9	6	7	11
		13				

15-826 Copyright: C. Faloutsos (2007) 55

CMU SCS

Linear hashing

Initially: $h(x) = x \text{ mod } N$ (N=4 here)
 overflow of bucket#1

bucket- id 0 1 2 3

4	8	5	9	6	7	11
		13				

17
↓

15-826 Copyright: C. Faloutsos (2007) 56

CMU SCS

Linear hashing

Initially: $h(x) = x \text{ mod } N$ (N=4 here)
 overflow of bucket#1
Split #0, anyway!!!

bucket- id 0 1 2 3

4	8	5	9	6	7	11
		13				

17
↓

15-826 Copyright: C. Faloutsos (2007) 57

CMU SCS

Linear hashing

Initially: $h(x) = x \bmod N$ (N=4 here)
 Split #0, anyway!!!

Q: But, how?

bucket- id 0 1 2 3

4	8	5	9	6	7	11
		13				

17
↓

15-826 Copyright: C. Faloutsos (2007) 58

CMU SCS

Linear hashing

A: use two h.f.: $h0(x) = x \bmod N$
 $h1(x) = x \bmod (2*N)$

17
↓

bucket- id 0 1 2 3

4	8	5	9	6	7	11
		13				

15-826 Copyright: C. Faloutsos (2007) 59

CMU SCS

Linear hashing - after split:

A: use two h.f.: $h0(x) = x \bmod N$
 $h1(x) = x \bmod (2*N)$

bucket- id 0 1 2 3 4

8	5	9	6	7	11	4
		13				

17

15-826 Copyright: C. Faloutsos (2007) 60

CMU SCS

Linear hashing - after split:

A: use two h.f.: $h0(x) = x \bmod N$
 $h1(x) = x \bmod (2*N)$

bucket- id 0 1 2 3 4

8	5 9 13	6	7 11	4
---	-----------	---	------	---

↓

17	overflow
----	----------

15-826 Copyright: C. Faloutsos (2007) 61

CMU SCS

Linear hashing - after split:

A: use two h.f.: $h0(x) = x \bmod N$
 $h1(x) = x \bmod (2*N)$

split ptr

bucket- id 0 1 2 3 4

8	5 9 13	6	7 11	4
---	-----------	---	------	---

↓

17	overflow
----	----------

15-826 Copyright: C. Faloutsos (2007) 62

CMU SCS

Linear hashing - overview

- Motivation
- main idea
- ➔ • search algo
- insertion/split algo
- deletion
- performance analysis
- variations

15-826 Copyright: C. Faloutsos (2007) 63

CMU SCS

Linear hashing - searching?

$h0(x) = x \text{ mod } N$ (for the un-split buckets)
 $h1(x) = x \text{ mod } (2*N)$ (for the splitted ones)

split ptr

bucket-id	0	1	2	3	4
	8	5 9 13	6	7 11	4

17 overflow

15-826 Copyright: C. Faloutsos (2007) 64

CMU SCS

Linear hashing - searching?

Q1: find key '6'? Q2: find key '4'?
 Q3: key '8'?

split ptr

bucket-id	0	1	2	3	4
	8	5 9 13	6	7 11	4

17 overflow

15-826 Copyright: C. Faloutsos (2007) 65

CMU SCS

Linear hashing - searching?

Algo to find key 'k':

- compute $b = h0(k)$;
 - if $b < \text{split-ptr}$, compute $b = h1(k)$
- search bucket b

15-826 Copyright: C. Faloutsos (2007) 66

CMU SCS

Linear hashing - overview

- Motivation
- main idea
- search algo
- ➔ insertion/split algo
- deletion
- performance analysis
- variations

15-826 Copyright: C. Faloutsos (2007) 67

CMU SCS

Linear hashing - insertion?

Algo: insert key ' k '

- compute appropriate bucket ' b '
- if the **overflow criterion** is true
 - split the bucket of ' $split\text{-}ptr$ '
 - $split\text{-}ptr++$ (*)

15-826 Copyright: C. Faloutsos (2007) 68

CMU SCS

Linear hashing - insertion?

notice: overflow criterion is up to us!!

Q: suggestions?

15-826 Copyright: C. Faloutsos (2007) 69

CMU SCS

Linear hashing - insertion?

notice: overflow criterion is up to us!!
 Q: suggestions?
 A1: space utilization \geq u-max

15-826 Copyright: C. Faloutsos (2007) 70

CMU SCS

Linear hashing - insertion?

notice: overflow criterion is up to us!!
 Q: suggestions?
 A1: space utilization $>$ u-max
 A2: avg length of ovf chains $>$ max-len
 A3:

15-826 Copyright: C. Faloutsos (2007) 71

CMU SCS

Linear hashing - insertion?

Algo: insert key ' k '

- compute appropriate bucket ' b '
- if the **overflow criterion** is true
 - split the bucket of 'split-ptr'
 - split-ptr ++ (*)

what if we reach the right edge??

15-826 Copyright: C. Faloutsos (2007) 72

CMU SCS

Linear hashing - split now?

$h0(x) = x \text{ mod } N$ (for the un-split buckets)
 $h1(x) = x \text{ mod } (2*N)$ for the splitted ones

split ptr

15-826 Copyright: C. Faloutsos (2007) 73

CMU SCS

Linear hashing - split now?

$h0(x) = x \text{ mod } N$ (for the un-split buckets)
 $h1(x) = x \text{ mod } (2*N)$ (for the splitted ones)

split ptr

15-826 Copyright: C. Faloutsos (2007) 74

CMU SCS

Linear hashing - split now?

~~$h0(x) = x \text{ mod } N$ (for the un-split buckets)~~
 $h1(x) = x \text{ mod } (2*N)$ (for the splitted ones)


split ptr

15-826 Copyright: C. Faloutsos (2007) 75

CMU SCS

Linear hashing - split now?

~~$h0(x) = x \text{ mod } N$ (for the un-split buckets)~~
 $h1(x) = x \text{ mod } (2*N)$ (for the splitted ones)

split ptr



0	1	2	3	4	5	6	7

15-826 Copyright: C. Faloutsos (2007) 76

CMU SCS

Linear hashing - split now?

this state is called 'full expansion'

split ptr


0	1	2	3	4	5	6	7

15-826 Copyright: C. Faloutsos (2007) 77

CMU SCS

Linear hashing - observations

In general, at any point of time, we have at **most two** h.f. active, of the form:

- $h_n(x) = x \text{ mod } (N * 2^n)$
- $h_{n+1}(x) = x \text{ mod } (N * 2^{n+1})$

(after a full expansion, we have only one h.f.)

15-826 Copyright: C. Faloutsos (2007) 78

CMU SCS

Linear hashing - overview

- Motivation
- main idea
- search algo
- insertion/split algo
- ➔ • deletion
- performance analysis
- variations

15-826 Copyright: C. Faloutsos (2007) 79

CMU SCS

Linear hashing - deletion?

- reverse of insertion:

15-826 Copyright: C. Faloutsos (2007) 80

CMU SCS

Linear hashing - deletion?

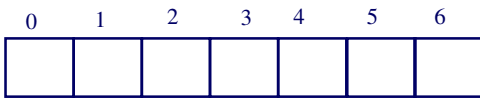
- reverse of insertion:
- if the underflow criterion is met
 - contract!

15-826 Copyright: C. Faloutsos (2007) 81

CMU SCS

Linear hashing - how to contract?

$h0(x) = \text{mod } N$ (for the un-split buckets)
 $h1(x) = \text{mod } (2*N)$ (for the splitted ones)

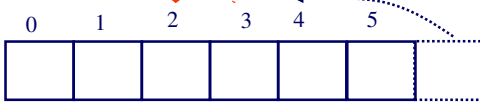
split ptr


15-826 Copyright: C. Faloutsos (2007) 82

CMU SCS

Linear hashing - how to contract?

$h0(x) = \text{mod } N$ (for the un-split buckets)
 $h1(x) = \text{mod } (2*N)$ (for the splitted ones)

split ptr


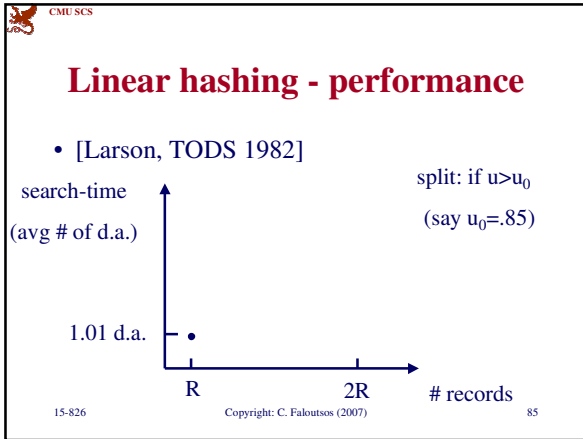
15-826 Copyright: C. Faloutsos (2007) 83

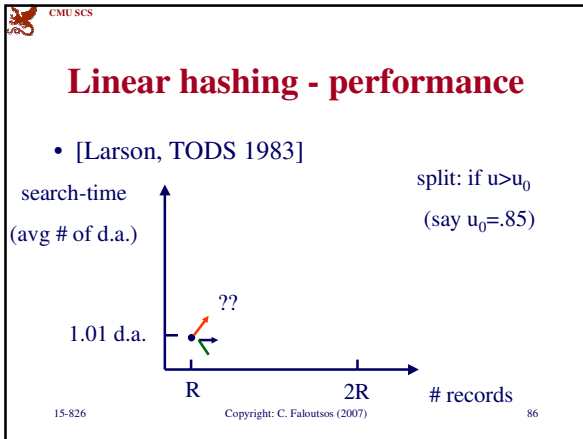
CMU SCS

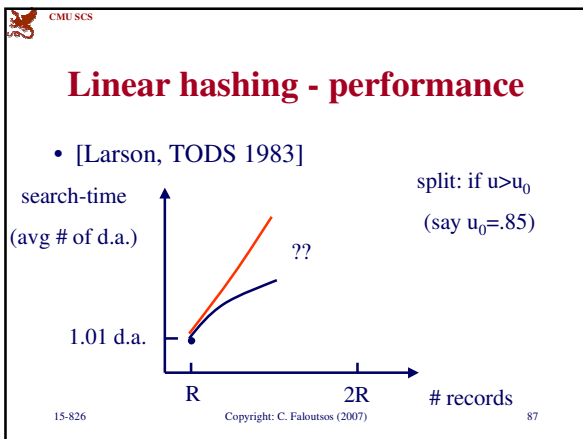
Linear hashing - overview

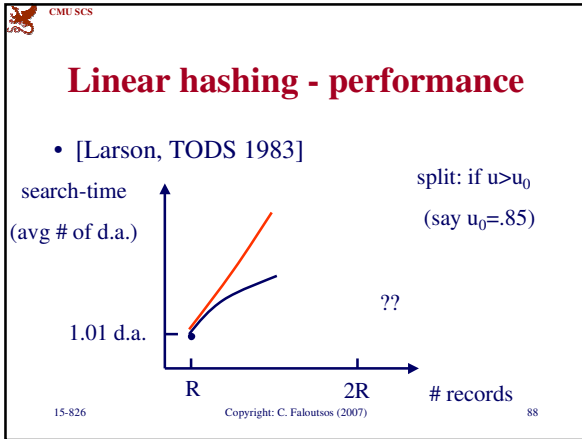
- Motivation
- main idea
- search algo
- insertion/split algo
- deletion
- ➔ • performance analysis
- variations

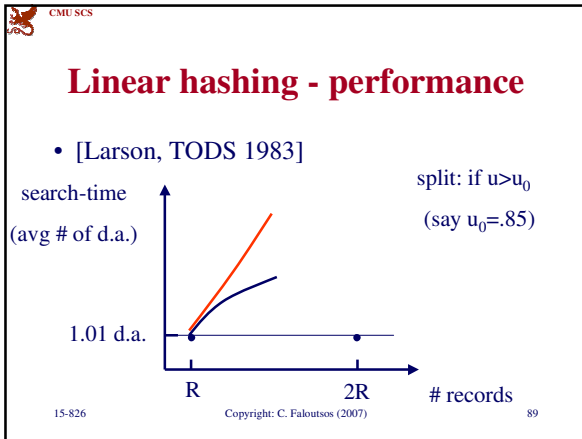
15-826 Copyright: C. Faloutsos (2007) 84

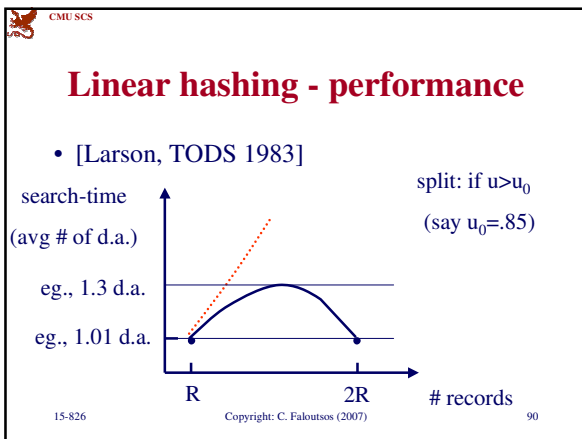












CMU SCS

Linear hashing - performance

- Q: How to shorten the maximum?

search-time

R 2R # records

15-826 Copyright: C. Faloutsos (2007) 91

CMU SCS

Linear hashing - overview

- Motivation
- main idea
- search algo
- insertion/split algo
- deletion
- performance analysis
- ➔ • variations

15-826 Copyright: C. Faloutsos (2007) 92

CMU SCS

Linear hashing - performance

- Q: How to shorten the maximum?
- A: 2-3 splits - partial expansions!

search-time

R 2R # records

15-826 Copyright: C. Faloutsos (2007) 93

CMU SCS

Linear hashing - performance

- Q: How to shorten the maximum?
- A: 2-3 splits - partial expansions!

15-826 Copyright: C. Faloutsos (2007) 94

CMU SCS

Linear hashing - variations

Two split pointers! On split:

0	1	2	3
---	---	---	---

15-826 Copyright: C. Faloutsos (2007) 95

CMU SCS

Linear hashing - variations

Two split pointers! On split:

0	1	2	3
---	---	---	---

15-826 Copyright: C. Faloutsos (2007) 96

CMU SCS

Linear hashing - variations

2nd split:

15-826 Copyright: C. Faloutsos (2007) 97

CMU SCS

Linear hashing - variations

2nd split: Partial expansion! (50% larger table)

15-826 Copyright: C. Faloutsos (2007) 98

CMU SCS

Linear hashing - variations

Q: how to do the third split?

15-826 Copyright: C. Faloutsos (2007) 99

CMU SCS

Linear hashing - variations

Q: how to do the third split?
A: 3-to-4 splits now!

0 1 2 3 4 5

15-826 Copyright: C. Faloutsos (2007) 100

CMU SCS

Linear hashing - performance

- Q1: Which of the two red peaks is higher?
- Q2: Why?

search-time

R 2R # records


15-826 Copyright: C. Faloutsos (2007) 101

CMU SCS

Linear hashing - overview

- Motivation
- main idea
- search algo
- insertion/split algo
- deletion
- performance analysis
- ➔ • variations


15-826 Copyright: C. Faloutsos (2007) 102

 CMU SCS

Other hashing variations

- ‘order preserving’
- ‘perfect hashing’ (no collisions!) [Ed. Fox, et al]


15-826 Copyright: C. Faloutsos (2007) 103

 CMU SCS

Primary key indexing - conclusions

- hashing is $O(1)$ on the average for search
- linear hashing: elegant way to grow a hash table
- B-trees: major contenders for primary-key indexing ($O(\log(N))$ w.c.!).


15-826 Copyright: C. Faloutsos (2007) 104

 CMU SCS

References for primary key indexing

- [Fagin+] Ronald Fagin, Jürg Nievergelt, Nicholas Pippenger, H. Raymond Strong: Extendible Hashing - A Fast Access Method for Dynamic Files. TODS 4(3): 315-344(1979)
- [Fox] Fox, E. A., L. S. Heath, Q.-F. Chen, and A. M. Daoud. "Practical Minimal Perfect Hash Functions for Large Databases." Communications of the ACM 35.1 (1992): 105-21.

15-826 Copyright: C. Faloutsos (2007) 105

 CMU SCS

References, cont'd

- [Knuth] D.E. Knuth. The Art Of Computer Programming, Vol. 3, Sorting and Searching, Addison Wesley
- [Larson] Per-Ake Larson Performance Analysis of Linear Hashing with Partial Expansions ACM TODS, 7,4, Dec. 1982, pp 566--587
- ➔ [Litwin] Litwin, W., (1980), Linear Hashing: A New Tool for File and Table Addressing, VLDB, Montreal, Canada, 1980

15-826 Copyright: C. Faloutsos (2007) 106
