

Physical File Organization

There's a big difference in search time, which underlies everything.

- hash file is best for equality search, but worst or close to worst for range search
- ordered file is good for searching (equality and range)

And –

- a file can only be ordered/hashed on one field (or group of fields)

Secondary Access Paths

A *secondary access path* is a way to access data besides just via the physical file organization.

- typically an index

Advantages:

- faster access

Disadvantages:

- take more space
- maintenance has a cost

Questions

Where are indexes stored? Individually for a user or within a database?

- indexes are associated with tables
- defined by DB admin as part of the database implementation (not by users)

Indexes

Indexes are built on a field or group of fields.

The type of index depends on

- whether the indexing field is the same as the file's ordering field
- whether the indexing field is a key

An *index record* contains a value and a pointer to the block containing the record with that value.

Index records are stored in sorted order.

Primary vs Clustering Indexes

Primary index –

- indexing field is the ordering field of the file, and *is a key*

Clustering index –

- indexing field is the ordering field of the file, but *isn't a key*

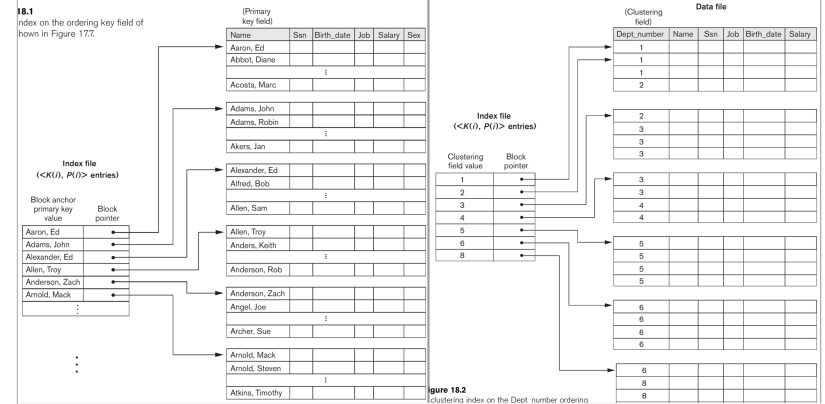
Can there be both primary and clustering indexes for the same file?

- no, at most one primary or clustering index per file

Primary vs Clustering Indexes

primary index

clustering index



one index record *per data block* of the original file contains value of indexing field for the first record on the block (*anchor record*)

one index record *per distinct value* of the indexing field contains pointer to first data block containing records with that value

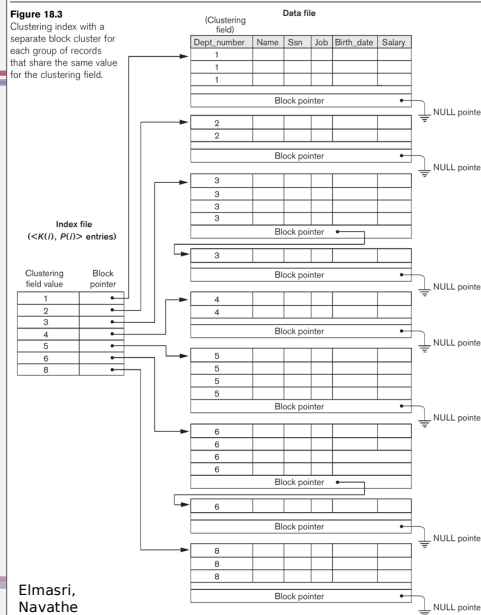
Elmasri, Navathe

Clustering Indexes

a common implementation is to instead allocate separate blocks for each value of the indexing field in the underlying data file

this simplifies insertion and deletion in the data file, as well as updating of the index

for the index, there is still one index record per distinct value of the indexing field



Primary vs Clustering Indexes

What's the point in having an index ordered by the same thing the file is already ordered by?

- index records are smaller ($b_{fr} > b_{fr_i}$) and there are fewer of them ($r_i < r$), so there are fewer blocks to load when searching ($b_i < b$)

Secondary Indexes

Organization.

- indexing field is not the ordering field of the file
 - can have any number of secondary indexes for a file, with or without primary or clustering index
- when the indexing field is a key –
 - one index record *for each data record*
 - index record contains pointer to the block containing that record
- when the indexing field is not a key –
 - option #1 – one index record for each data record
 - option #2 – variable-length index records listing all of the data blocks for each indexing field value
 - option #3 (most common) – one index record *for each distinct indexing field value*, containing pointer to an *indirect block* containing pointers to the data records with that field value
 - takes more space but avoids variable-length records


9

Secondary Index – Key

indexing field is a key –

one index record per data record

index record points to data block containing that record



CPSC 343: Database Theory and Practice • Fall 2024

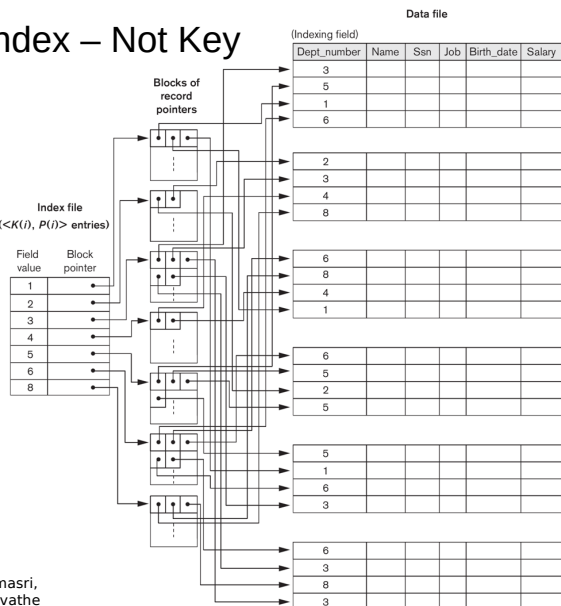
Elmasri, Navathe

Secondary Index – Not Key

indexing field is not a key, option 3 (indirect blocks) –

one index record per distinct value of the indexing field

index record points to *indirect block(s)* containing pointers to the data blocks containing the records with that value



CPSC 343: Database Theory and Practice

Elmasri, Navathe

Secondary Indexes

Can there be both secondary and primary indexes for the same file?

- yes

Can there be both secondary and clustering indexes for the same file?

- yes

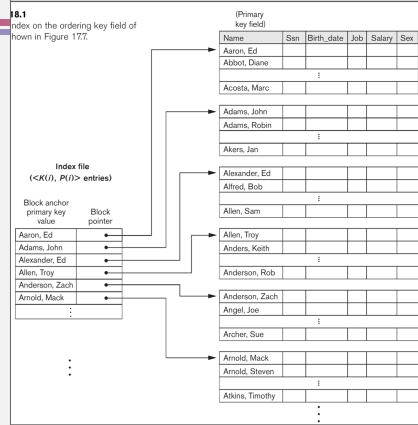
Can there be more than one secondary index per file?

- yes

CPSC 343: Database Theory and Practice • Fall 2024

32

Primary Index



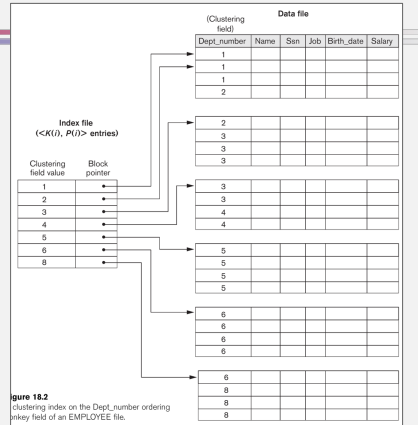
one index record *per data block* of the original file contains value of indexing field for the first record on the block (*anchor record*)

Primary Index

indexing field is a key
one index record per data block

operation			file properties		
search on indexing field	binary search on file	$\text{ceil}(\log_2(b)) = 12$	# records	r	30,000
	binary search on index + data block	$\text{ceil}(\log_2(b)) + 1 = 7$	block size	B	1024
range search on indexing field	on file (binary search to find first, then scan to get rest of the s matches)	$\text{ceil}(\log_2(b)) + \text{ceil}(s/\text{bfr}) = 12 + \text{ceil}(s/10)$	record length	R	100 bytes (fixed length)
	using index (find first using index, then scan file to get rest of the s matches)	$\text{ceil}(\log_2(b)) + \text{ceil}(s/\text{bfr}) = 6 + \text{ceil}(s/10)$	blocking factor (# records per block)	bfr	$\text{bfr} = \text{floor}(B/R)$ 10
			# data blocks	b	$b = \text{ceil}(r/\text{bfr})$ 3000
			index properties		
			ordering key length	V	9 bytes
			block pointer	P	6 bytes
			# index records	r_i	one per data block = b 3000
			index record length	R_i	$R_i = V + P$ 15 bytes
			blocking factor for index	bfr_i	$\text{bfr}_i = \text{floor}(B/R_i)$ 68
			# index blocks	b_i	$b_i = \text{ceil}(r_i/\text{bfr}_i)$ 45

Clustering Index



one index record *per distinct value* of the indexing field contains pointer to first data block containing records with that value

Clustering Index

indexing field is not a key
one index record per distinct value

operation			file properties		
search on indexing field	binary search + scan file (expect r/d matches)	$\text{ceil}(\log_2(b)) + \text{ceil}(r/d/\text{bfr}) = 12 + 1 = 13$	# records	r	30,000
	binary search on index + data block(s) (expect r/d matches)	$\text{ceil}(\log_2(b)) + \text{ceil}(r/d/\text{bfr}) = 6 + 1 = 7$	block size	B	1024
range search on indexing field	binary search + scan file (expect s matches)	$\text{ceil}(\log_2(b)) + \text{ceil}(s/\text{bfr}) = 12 + \text{ceil}(s/10)$	record length	R	100 bytes (fixed length)
	binary search on index + data block(s) (expect s matches)	$\text{ceil}(\log_2(b)) + \text{ceil}(s/\text{bfr}) = 6 + \text{ceil}(s/10)$	blocking factor (# records per block)	bfr	$\text{bfr} = \text{floor}(B/R)$ 10
			# data blocks	b	$b = \text{ceil}(r/\text{bfr})$ 3000
			# distinct values for index field	d	3000
			index properties		
			ordering key length	V	9 bytes
			block pointer	P	6 bytes
			# index records	r_i	one per distinct value 3000
			index record length	R_i	$R_i = V + P$ 15 bytes
			blocking factor for index	bfr_i	$\text{bfr}_i = \text{floor}(B/R_i)$ 68
			# index blocks	b_i	$b_i = \text{ceil}(r_i/\text{bfr}_i)$ 45