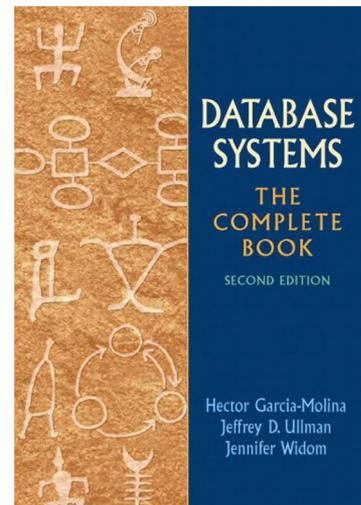CS 4440 A

Emerging Database Technologies

Agenda

1. Index overview

2. Using Index in SQL

3. Index structure basics

Reading Materials

Database Systems: The Complete Book (2nd edition)
- Chapter 8.3: Indexes in SQL
- Chapter 14.1: Index-Structure Basics

1. Index Overview

Index Motivation

Suppose we want to search for people of a specific age

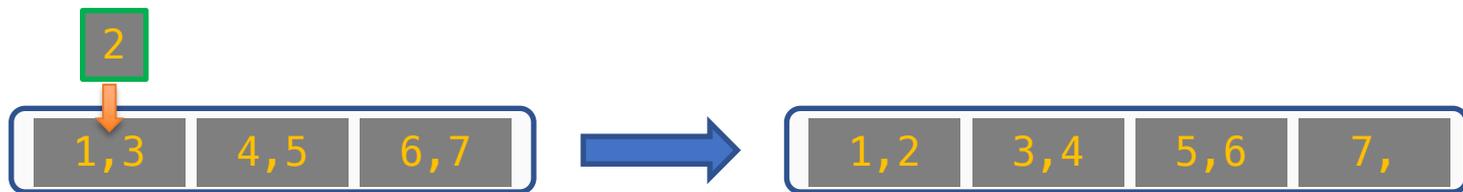*First idea:* Sort the records by age… we know how to do this fast!

How many IO operations to search over *N sorted* records?

- ○ Simple scan: *O(N)*

- ○ Binary search: $O(\log_2 N)$

Could we get even cheaper search? E.g. go from $\log_2 N \rightarrow \log_{200} N$?

Index Motivation

What about if we want to **insert** a new person, but keep the list sorted?

We would have to potentially shift *N* records, requiring up to ~ **2*N/P** IO operations (where P = # of records per page)!

○ We could leave some "slack" in the pages…

Could we get faster insertions?

Index Motivation

What about if we want to be able to search quickly along multiple attributes (e.g. not just age)?

- ○ We could keep multiple copies of the records, each sorted by one attribute set… this would take a lot of space

Can we get fast search over multiple attribute (sets) without taking too much space?

We'll create separate data structures called indexes to address all these points

Further Motivation for Indexes: NoSQL!

NoSQL engines are (basically) *just indexes!*

- ○ A lot more is left to the user in NoSQL… one of the primary remaining functions of the DBMS is still to provide index over the data records, for the reasons we just saw!

- ○ Example indexes: B+ Trees, hash indexes, LSM Tree

Indexes are critical across all DBMS types

Indexes: High-level

An *index* on a file speeds up selections on the *search key fields* for the index.

- Search key properties
    - Any subset of fields
    - is **not** the same as *key of a relation*

Example:

```
Product(name, maker, price)
```

On which attributes would you build indexes?

More precisely

An *index* is a **data structure** mapping <u>search keys</u> to <u>sets of rows in a database table</u>

- Provides efficient lookup & retrieval by search key value- usually much faster than searching through all the rows of the database table

An index can store the full rows it points to (*clustered index*) or pointers to those rows (*unclustered index*)

- We'll cover both, but mainly consider unclustered indexes

Operations on an Index

<u>Search</u>: Quickly find all records which meet some *condition on the search key attributes*

      ○   Point queries, range queries, …

<u>Insert / Remove</u> entries

      ○   Bulk Load / Delete.

> Indexing is one the most important features provided by a database for performance

Overview of Common Index Types

B-Tree and its variants *(next Monday)*
- ○ The default index in most RDBMS
- ○ very good for range queries, sorted data

Hash Tables *(next Wednesday)*
- ○ These are **persistent hash tables** designed specifically for disk storage
- ○ Good for point queries

Other specialized indexes:
- ○ Bitmap indexes: for low-cardinality data
- ○ Multi-dimensional indexes: KD-tree, R-tree for GIS data
- ○ LSM-tree: commonly used in NoSQL systems
- ○ Vector indexes: HNSW, IVF-PQ for high-dimensional vector data

2. Using Index in SQL

Using Indexes in SQL

An index is used to efficiently find tuples with certain values of attributes
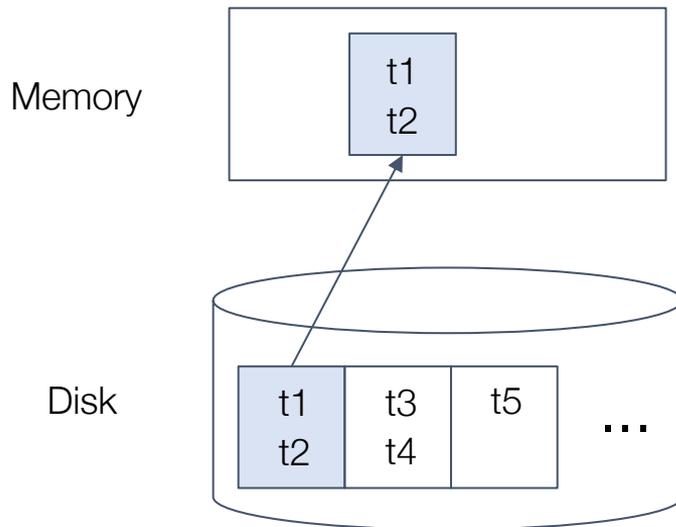
Selection of indexes
- An index may speed up lookups and joins
- However, every built index makes insertions, deletions, and updates to relation more complex and time-consuming

```
CREATE INDEX KeyIndex ON Movies(title, year);
```

```
DROP INDEX KeyIndex;
```

Recall: Simple I/O based cost model

- Multiple tuples are stored in blocks on disk
- Every block needed is always retrieved from disk
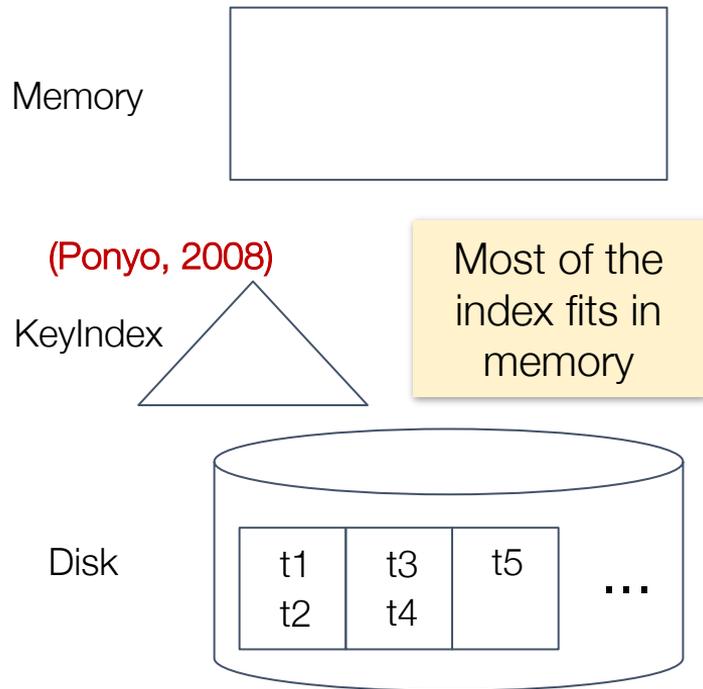- Disk I/Os are expensive

Index on a key

- An index on a key is often useful
- Retrieve at most one block to memory for tuple
  - Possibly other blocks for the index itself

```
CREATE INDEX KeyIndex ON Movies(title, year);
```

```
SELECT *
FROM Movies
WHERE title = 'Ponyo' AND year = 2008;
```

Memory

(Ponyo, 2008)

KeyIndex

Most of the index fits in memory

Disk

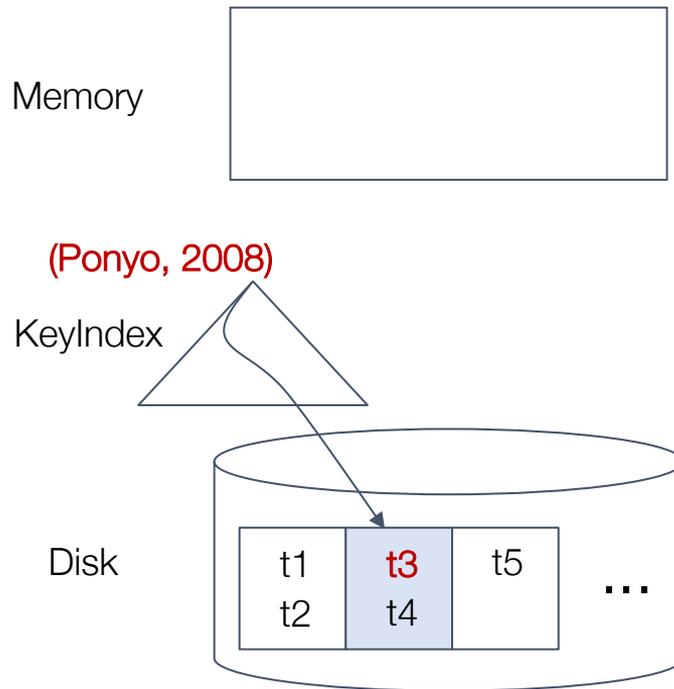| t1 t2 | t3 t4 | t5 | ... |

Index on a key

- An index on a key is often useful
- Retrieve at most one block to memory for tuple
  - Possibly other blocks for the index itself

```
CREATE INDEX KeyIndex ON Movies(title, year);
```

```
SELECT *
FROM Movies
WHERE title = 'Ponyo' AND year = 2008;
```

Memory

(Ponyo, 2008)

KeyIndex

Disk

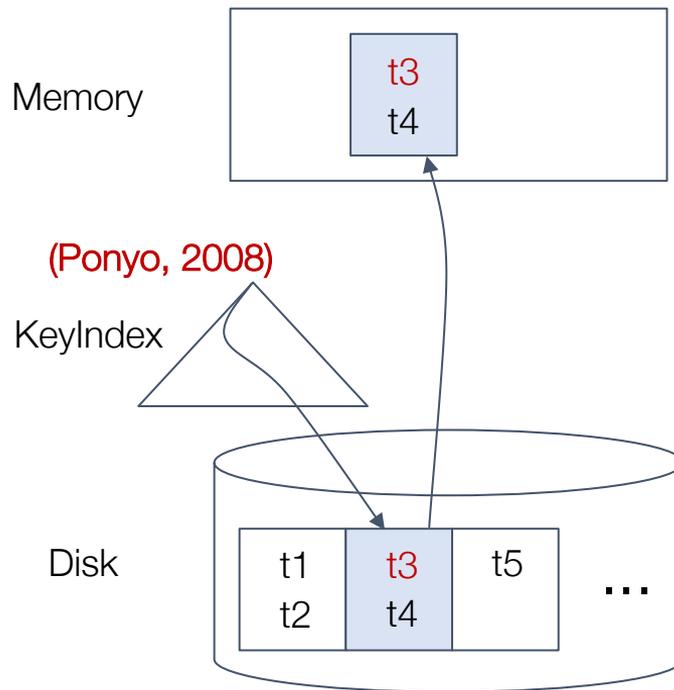| t1 | t3 | t5 |
| t2 | t4 | |

...

Index on a key

- An index on a key is often useful
- Retrieve at most one block to memory for tuple
  - Possibly other blocks for the index itself

```
CREATE INDEX KeyIndex ON Movies(title, year);
```

```
SELECT *
FROM Movies
WHERE title = 'Ponyo' AND year = 2008;
```

Memory

(Ponyo, 2008)

KeyIndex

Disk

t3
t4

t1
t2

t3
t4

t5

...

Indexes can be used in joins

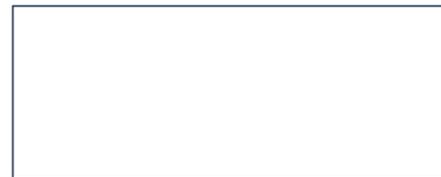With the right indexes, the join below only requires 2 page reads for the tuples
  ○ And possibly a small number of other pages for accessing the indexes

```
CREATE INDEX MIndex ON Movies(title, year);
```

```
CREATE INDEX MEIndex ON MovieExec(cert#);
```

```
SELECT name
FROM Movies, MovieExec
WHERE title = 'Ponyo' AND year = 2008
  AND producerC# = cert#;
```
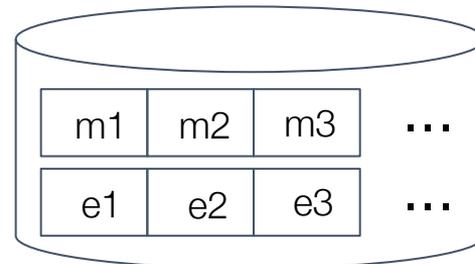
Memory

(Ponyo, 2008)

MIndex   MEIndex

Disk

| m1 | m2 | m3 | ··· |
| e1 | e2 | e3 | ··· |

Indexes can be used in joins

With the right indexes, the join below only requires 2 page reads for the tuples
  ○ And possibly a small number of other pages for accessing the indexes

```
CREATE INDEX MIndex ON Movies(title, year);
```

```
CREATE INDEX MEIndex ON MovieExec(cert#);
```

```
SELECT name
FROM Movies, MovieExec
WHERE title = 'Ponyo' AND year = 2008
  AND producerC# = cert#;
```
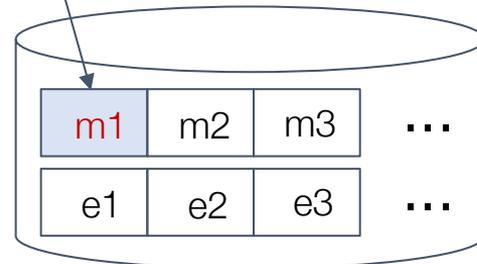
Memory

(Ponyo, 2008)

MIndex   MEIndex

Disk

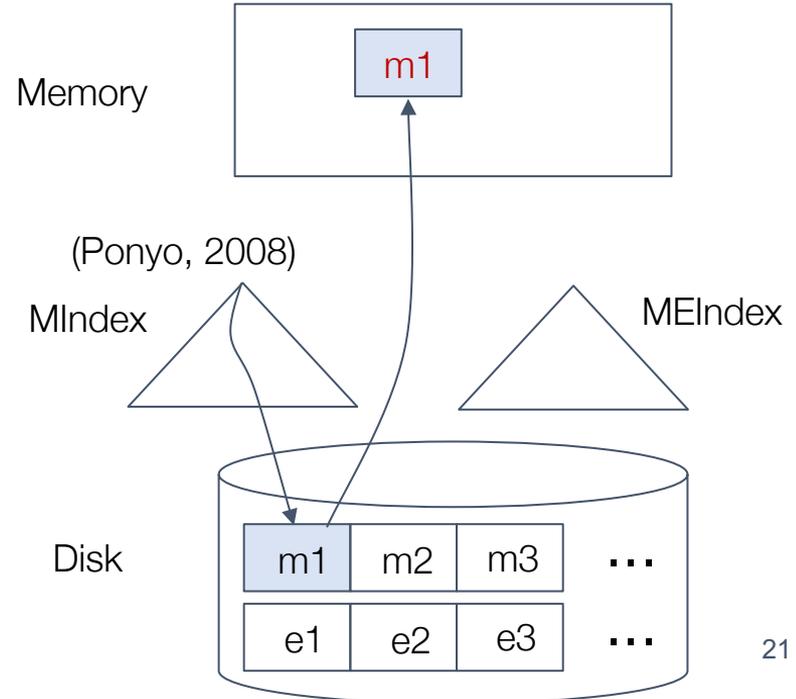| m1	m2	m3	...
e1	e2	e3	...

Indexes can be used in joins

With the right indexes, the join below only requires 2 page reads for the tuples
- And possibly a small number of other pages for accessing the indexes

```
CREATE INDEX MIndex ON Movies(title, year);
```

```
CREATE INDEX MEIndex ON MovieExec(cert#);
```

```
SELECT name
FROM Movies, MovieExec
WHERE title = 'Ponyo' AND year = 2008
  AND producerC# = cert#;
```

Memory

m1

(Ponyo, 2008)

MIndex

MEIndex

Disk

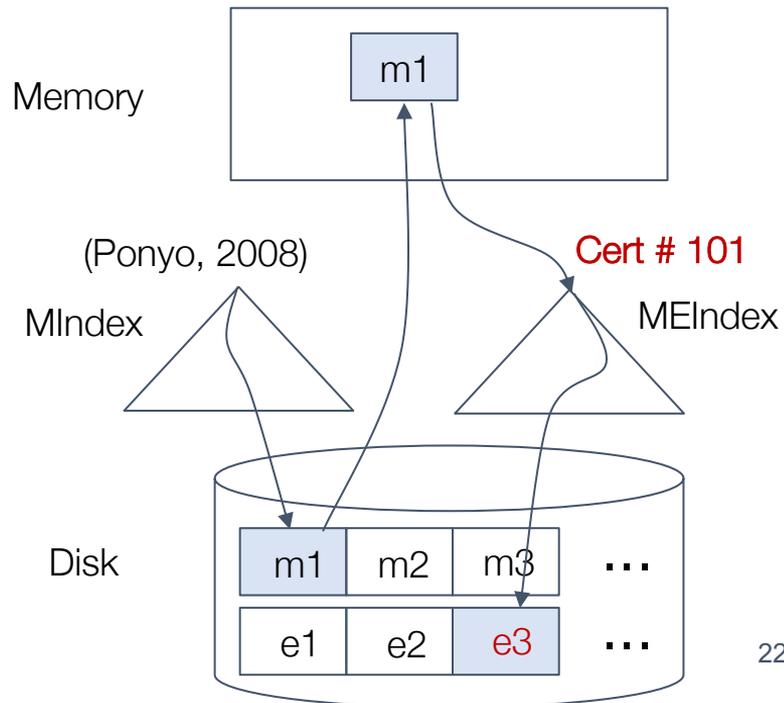| m1	m2	m3	...
e1	e2	e3	...

Indexes can be used in joins

With the right indexes, the join below only requires 2 page reads for the tuples
  ○ And possibly a small number of other pages for accessing the indexes

```
CREATE INDEX MIndex ON Movies(title, year);
```

```
CREATE INDEX MEIndex ON MovieExec(cert#);
```

```
SELECT name
FROM Movies, MovieExec
WHERE title = 'Ponyo' AND year = 2008
  AND producerC# = cert#;
```

Memory

m1

(Ponyo, 2008)   Cert # 101

MIndex   MEIndex

Disk

| m1 | m2 | m3 | ... |

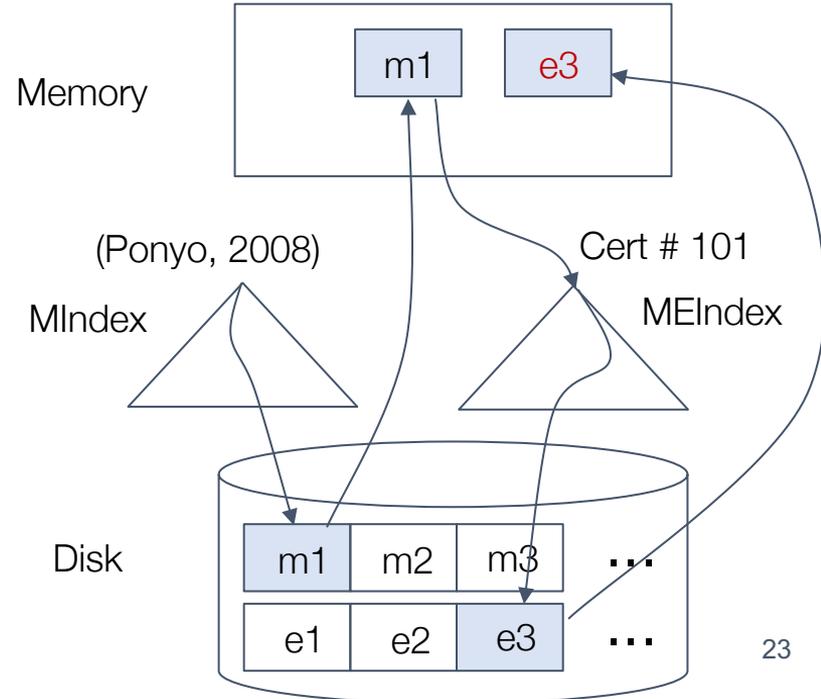| e1 | e2 | e3 | ... |

Indexes can be used in joins

With the right indexes, the join below only requires 2 page reads for the tuples
- And possibly a small number of other pages for accessing the indexes

```
CREATE INDEX MIndex ON Movies(title, year);
```

```
CREATE INDEX MEIndex ON MovieExec(cert#);
```

```
SELECT name
FROM Movies, MovieExec
WHERE title = 'Ponyo' AND year = 2008
  AND producerC# = cert#;
```

Memory

m1   e3

(Ponyo, 2008)

Cert # 101

MIndex

MEIndex

Disk

m1   m2   m3   ...

e1   e2   e3   ...

Covering Index

We say that an index is *covering* if all the needed columns are in the index.

Needed columns is the union of those in WHERE *and* SELECT.
- Means that the query can be answered using the index--without reading the table!

```
CREATE INDEX MIndex ON Movies(title, year);
```

```
CREATE INDEX MIndex ON
Movies(name, title, year);
```

```
SELECT name
FROM Movies
WHERE title = 'Ponyo' AND
year = 2008;
```

Index Activity

- Collab notebook link:
  https://tinyurl.com/mpw6jabr

- Create a COPY of the notebook

- Follow the instructions to download the database file and load it in the notebook

3. Index Structure Basics

Sequential file

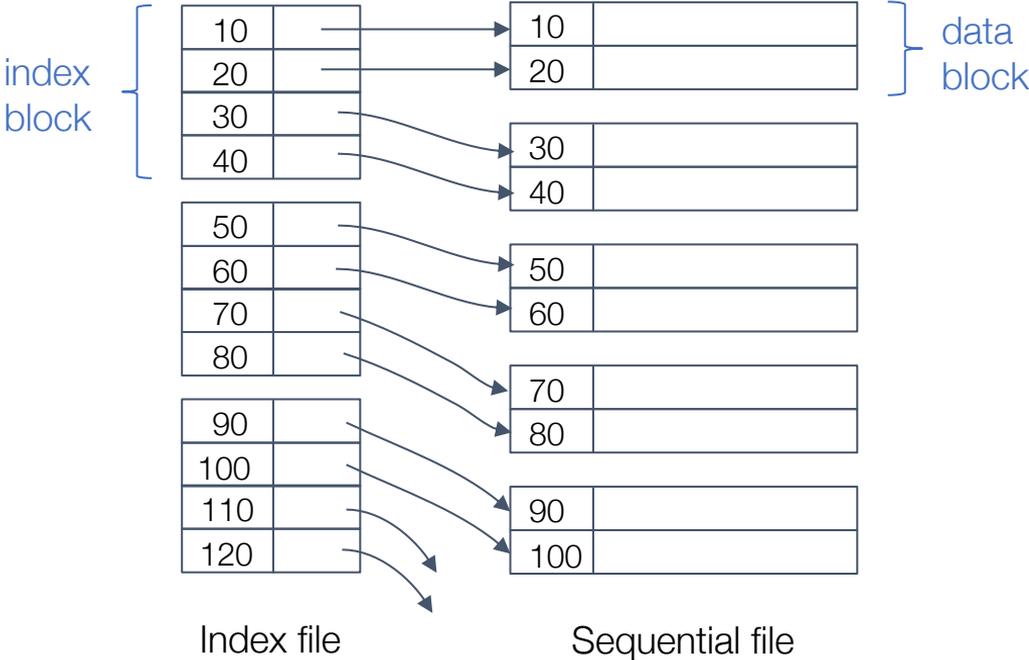- A file containing tuples of a relation sorted by their primary key

One block

10	
20	

30	
40	

50	
60	

70	
80	

90	
100	

Sequential file

Dense index

- A sequence of blocks holding keys of records and pointers to the records
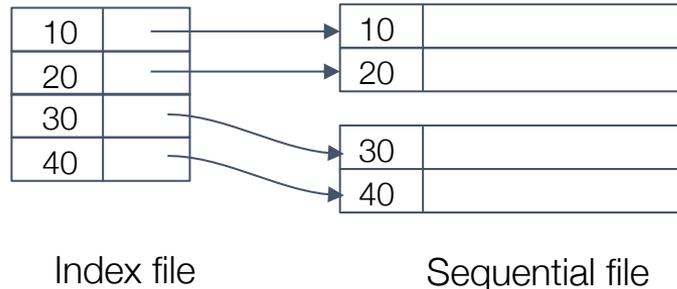

Index file   Sequential file

Dense index

Given key K, search index blocks for K, then follow associated pointer

Why is this efficient?
- Number of index blocks usually smaller than number of data blocks
- Keys are sorted, so we can use binary search
- The index may be small enough to fit in memory

Index file   Sequential file

Sparse index

- Has one key-pointer pair per block of the data file
- Uses less space than dense index, but needs more time to find a record

index block
data block
Index file
Sequential file
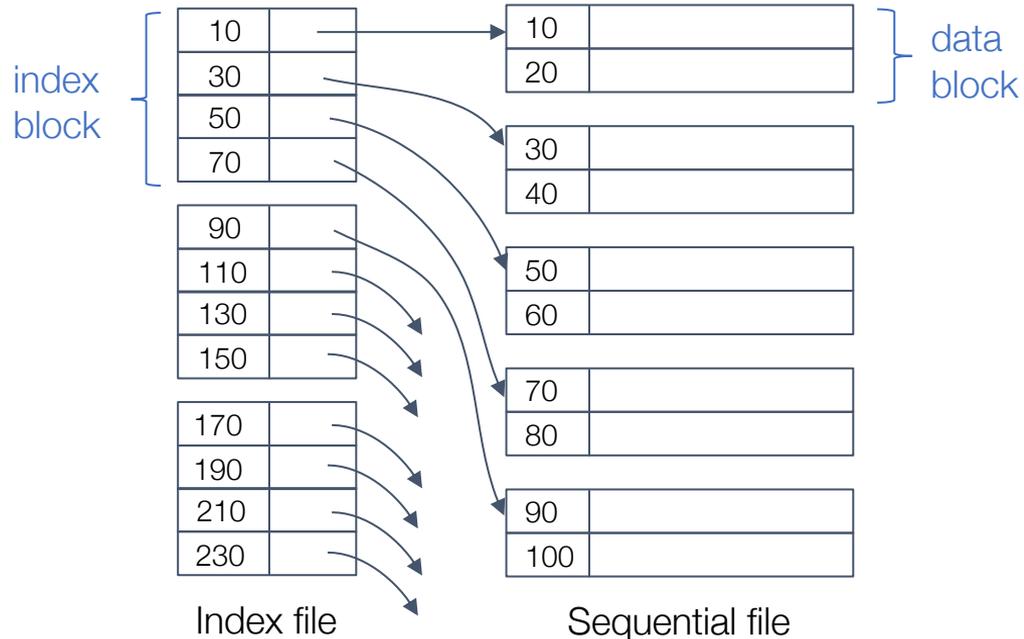
In-class Exercise

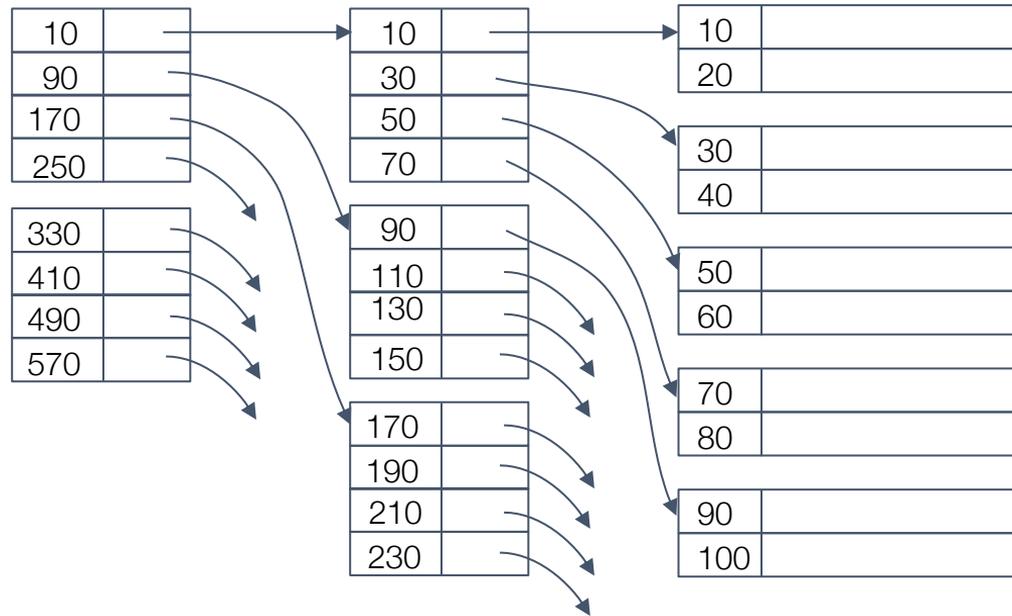Suppose a block holds 3 records or 10 key-pointer pairs

If there are n records in a data file, how many blocks are needed to hold
- ○ The data file and a dense index
- ○ The data file and a sparse index

Multiple levels of index

If the index file is still large, add another level of indexing
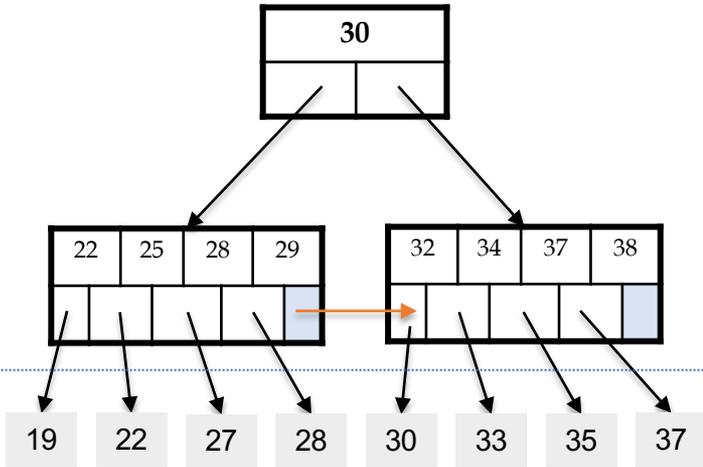- Basic idea of the B+-tree index (next lecture)

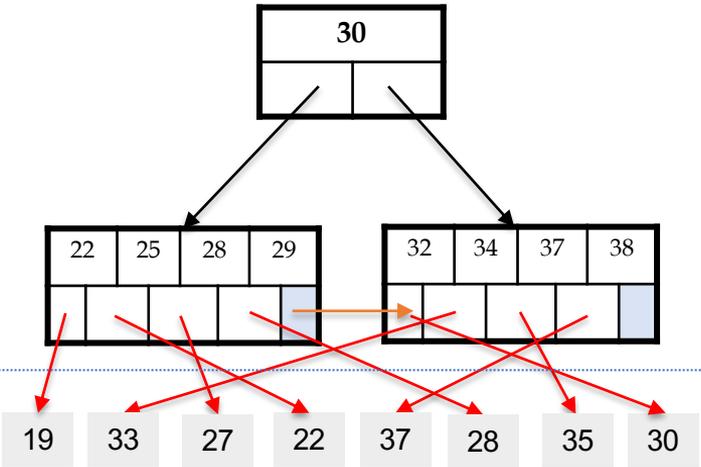Q: Should the blocks of additional levels be dense or sparse?

Clustered Indexes

An index is <u>clustered</u> if the underlying data is ordered in the same way as the index's data entries.

Clustered vs. Unclustered Index

Index Entries

| 30 |
| 22 | 25 | 28 | 29 | | 32 | 34 | 37 | 38 |

| 19 | 22 | 27 | 28 | 30 | 33 | 35 | 37 |

| 30 |
| 22 | 25 | 28 | 29 | | 32 | 34 | 37 | 38 |

| 19 | 33 | 27 | 22 | 37 | 28 | 35 | 30 |

Data Records

Clustered: often on primary key

Unclustered

Q: How many clustered/unclustered indexes can a table have?

Clustered vs. Unclustered Index

Recall that for a disk with block access, **sequential IO is much faster than random IO**

For point lookup, no difference between clustered / unclustered

For range search over R values: difference between **1 random IO + R sequential IO**, and **R random IO:**

- A random IO costs ~ 10ms (sequential much much faster)

- For R = 100,000 records- **difference between ~10ms and ~17min!**

Non-clustered Index

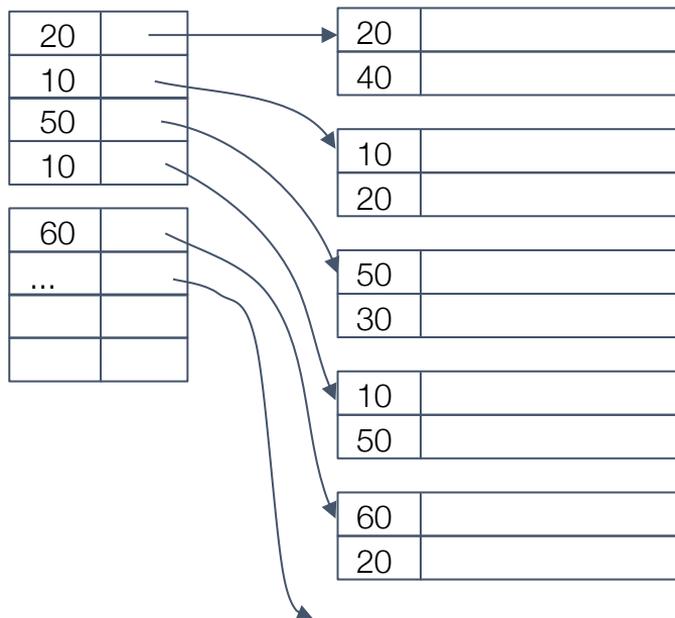Unlike a clustered index, does not determine the placement of records

20	
40	

10	
20	

50	
30	

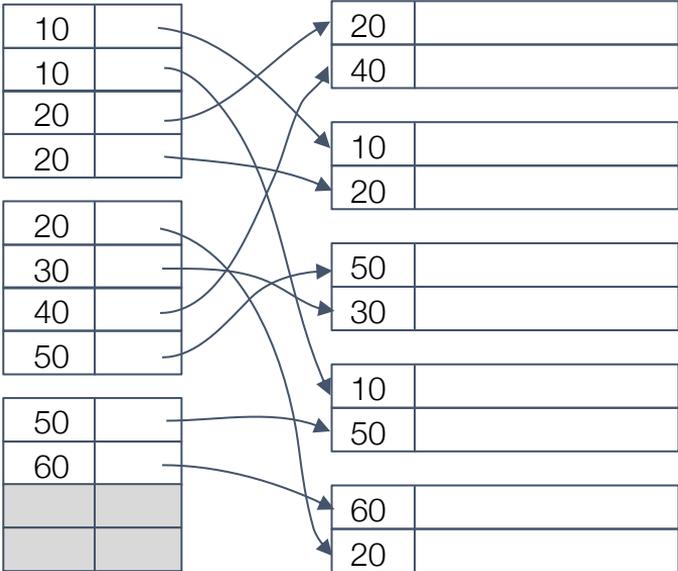|	
10	
50	

60	
20	

Non-clustered Index
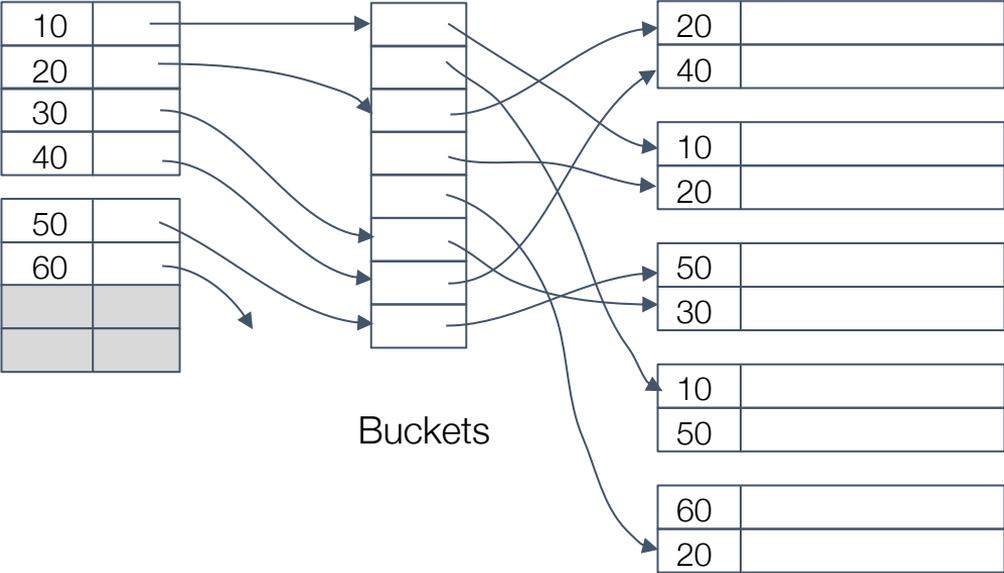
Using a sparse index doesn't make sense

Non-clustered index

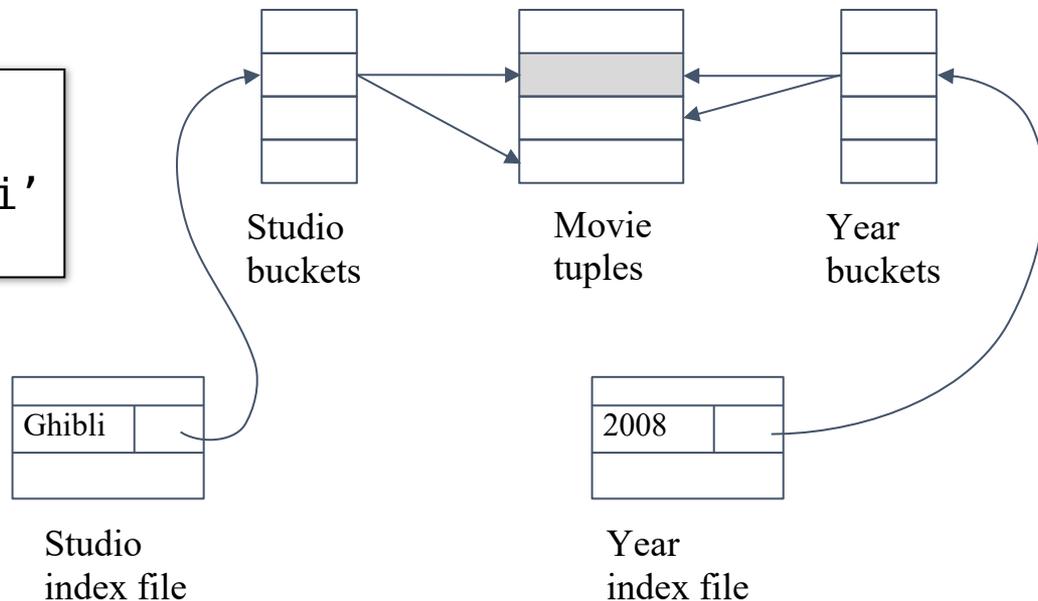As a result, non-clustered indexes are always dense

Non-clustered index

To remove redundant keys in index file, use level of indirection

Buckets

When is indirection useful?

- When a key is larger than a pointer and each key appears twice on average
- Another advantage: use bucket pointers without looking at most of the records

```
SELECT title
FROM Movies
WHERE studioName = 'Ghibli'
    AND year = 2008;
```
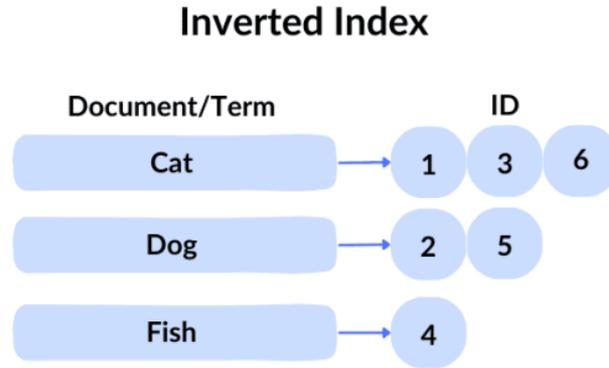
Studio
buckets

Movie
tuples

Year
buckets

Ghibli

2008

Studio
index file

Year
index file

Inverted index
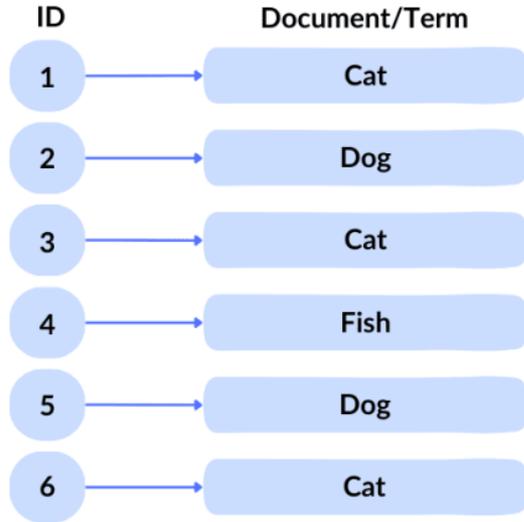
Basically the same structure as the non-clustered index
- e.g., Search for documents containing "cat" or "dog" (or both)

cat				...the cat is fat...
dog				… raining cats and dogs ...
				…the dog is eating ...

Inverted index

Buckets
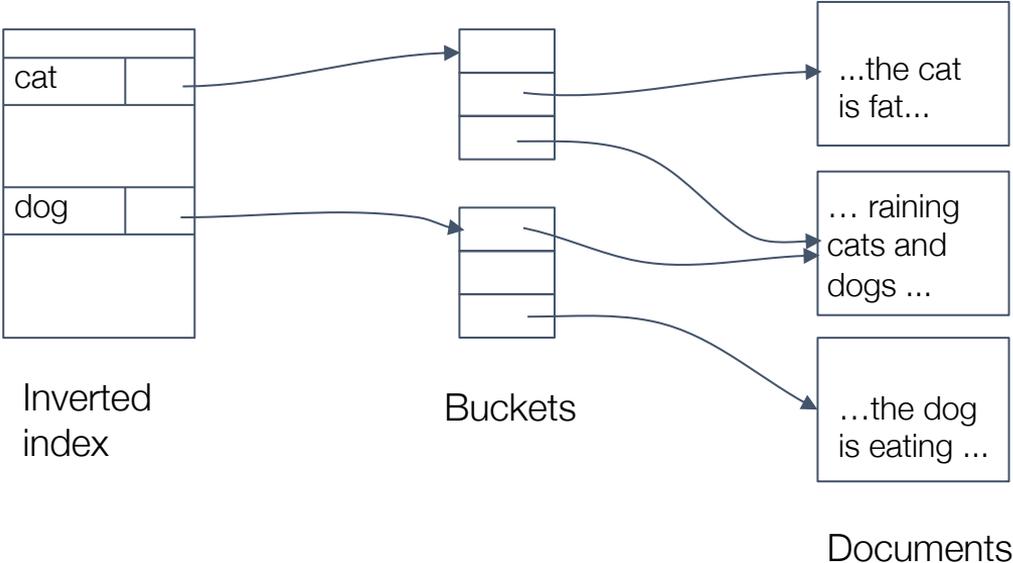
Documents

Inverted Index

Where the name came from

Store more information in inverted index

Can answer more complex queries like:
- Find documents where "dog" and "cat" are within 10 words
- Find documents about dogs that refer to other documents about cats

	Type	Position	
	title	5	
	anchor	3	
	title	11	
	text	20	
	text	50	

cat → (title, anchor rows)

dog → (title, text, text rows)

doc 1
doc 2
doc 3