



Chapter 14: Indexing

Database System Concepts, 7th Ed.

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use



Chapter 14: Indexing

- Basic Concepts
- Ordered Indices
- B⁺-Tree Index Files
- B-Tree Index Files
- Hashing

<https://www.vertabelo.com/blog/technical-articles/an-introduction-to-mysql-indexes>



Basic Concepts

- Indexing mechanisms used to **speed up access** to desired data.
 - E.g., author catalog in library
- **Search Key** - attribute to set of attributes used to look up records in a file.
- An **index file** consists of records (called **index entries**) of the form

search-key	pointer
------------	---------



What is an index?

■ **Index** = a **data structure** that enable the **user** to **find (= locate)** data items *efficiently (quickly)* using **search keys**

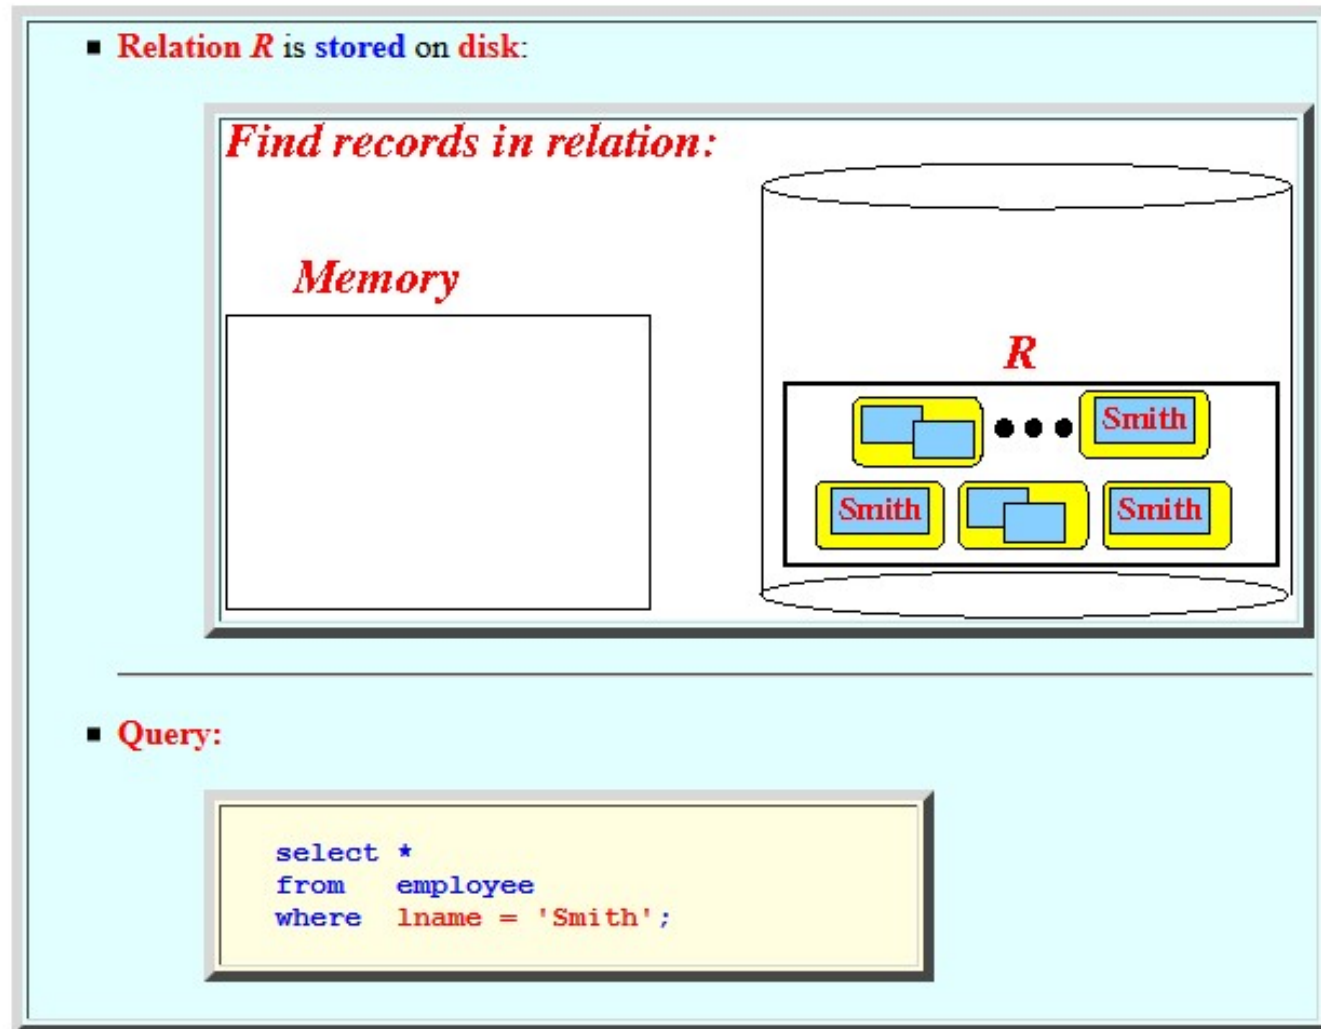
○ Sample of an **index** (found in the **back** of a **text book**):

Index	
A	Dial type 4, 12
About cordless telephones 51	Directory 17
Advanced operation 17	DSL filter 5
Answer an external call during an intercom call 15	E
Answering system operation 27	Edit an entry in the directory 20
B	Edit handset name 11
Basic operation 14	F
Battery 9, 38	FCC, ACTA and IC regulations 53
C	Find handset 16
Call log 22, 37	H
Call waiting 14	Handset display screen messages 36
Chart of characters 18	Handset layout 6
D	I
Date and time 8	Important safety instructions 39
Delete from redial 26	Index 56-57
Delete from the call log 24	Installation 1
Delete from the directory 20	Install handset battery 2
Delete your announcement 32	Intercom call 15
Desk/table bracket installation 4	Internet 4
Dial a number from redial 26	



Example

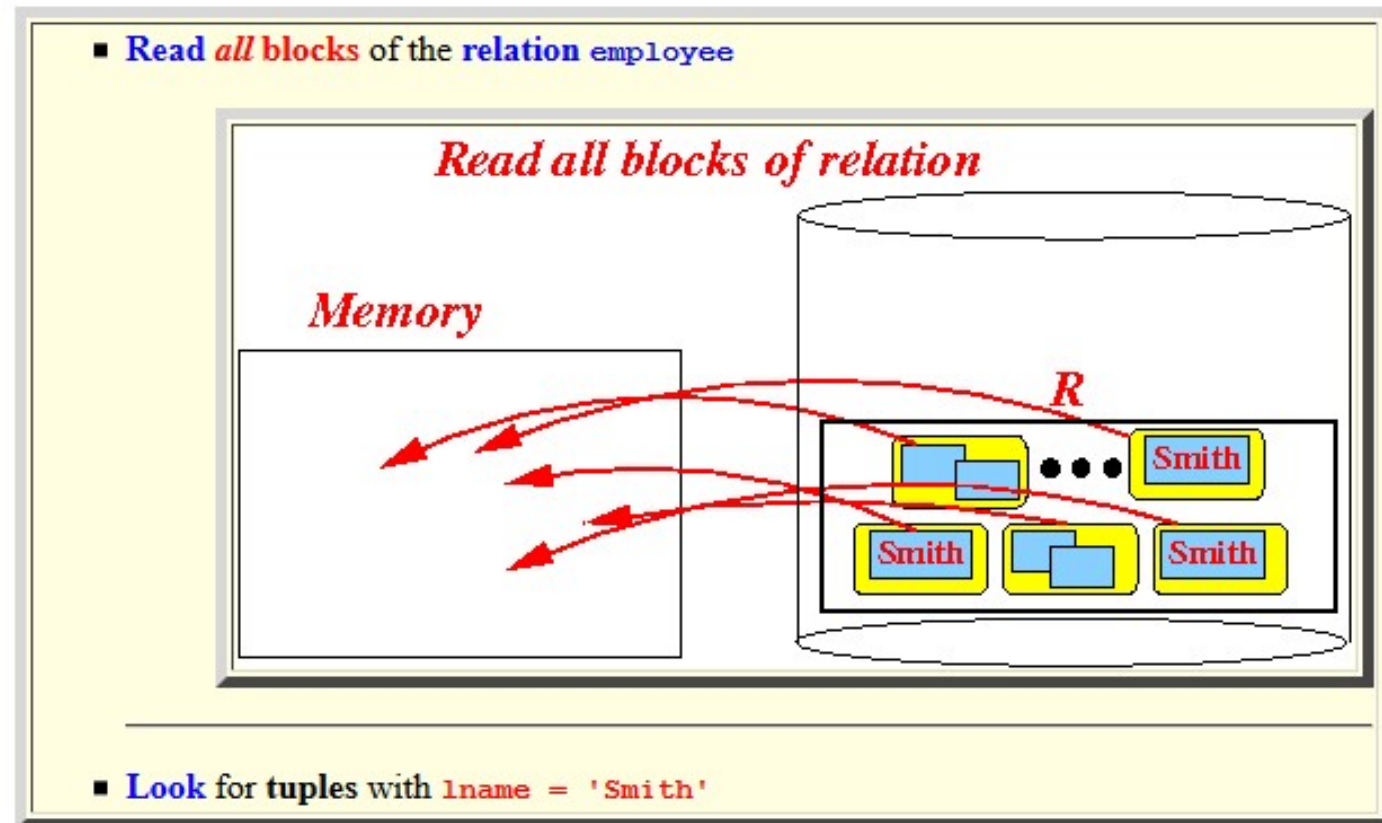
- o Consider the following query on **relation** stored on **disk**:





Example

- o **Without** further information, the **only way** to process the **above query** is:



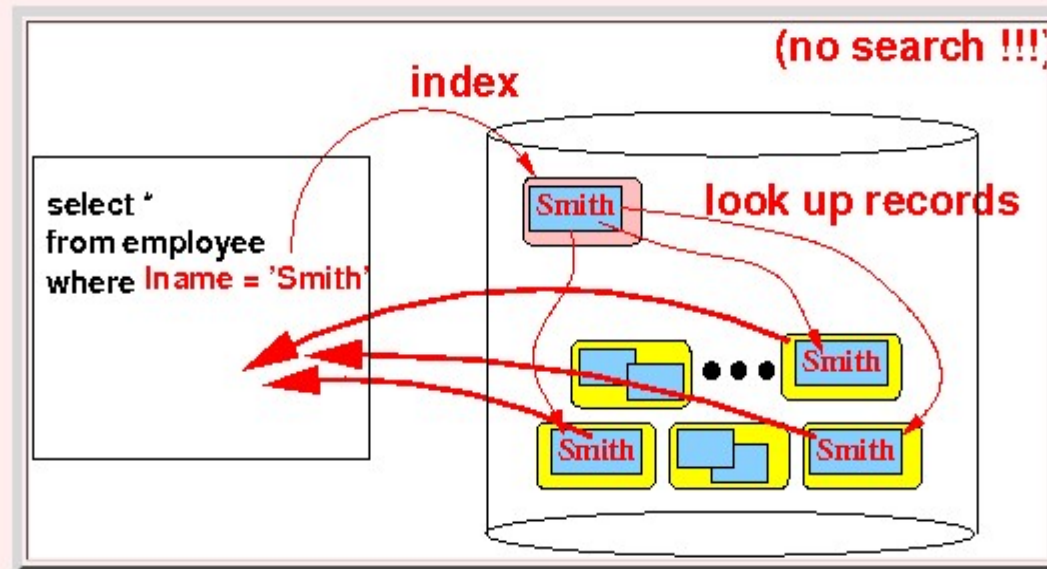


Example

- o **More efficient** processing:

- **Maintain location information** on **specific attributes** (most **useful** ones, e.g., **key attributes**) in the **relation**

Graphically:



We **only** need to **read** a **subset** of the **data block**



Definitions

- Search key:

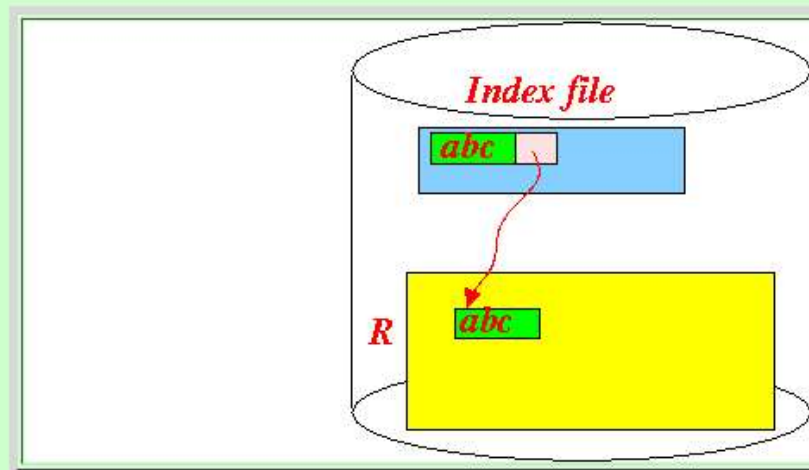
- Search key = field(s) used to create the *additional search information* used to speed up a look up operation

- Index file:

- Index file = a file that store records of the following format:

```
+-----+  
| Search key | block/record pointer |  
+-----+
```

Graphically:



Note:

- The size of an index file is usually *much smaller* than the size of a data file



Note

- The **pointer** is usually a **block pointer**

- So the **index** allow you to locate the **block** that contain the **record** quickly

- The **record** is found by a **search operation** inside the **block** (after the **block** is **read** into **main memory**)

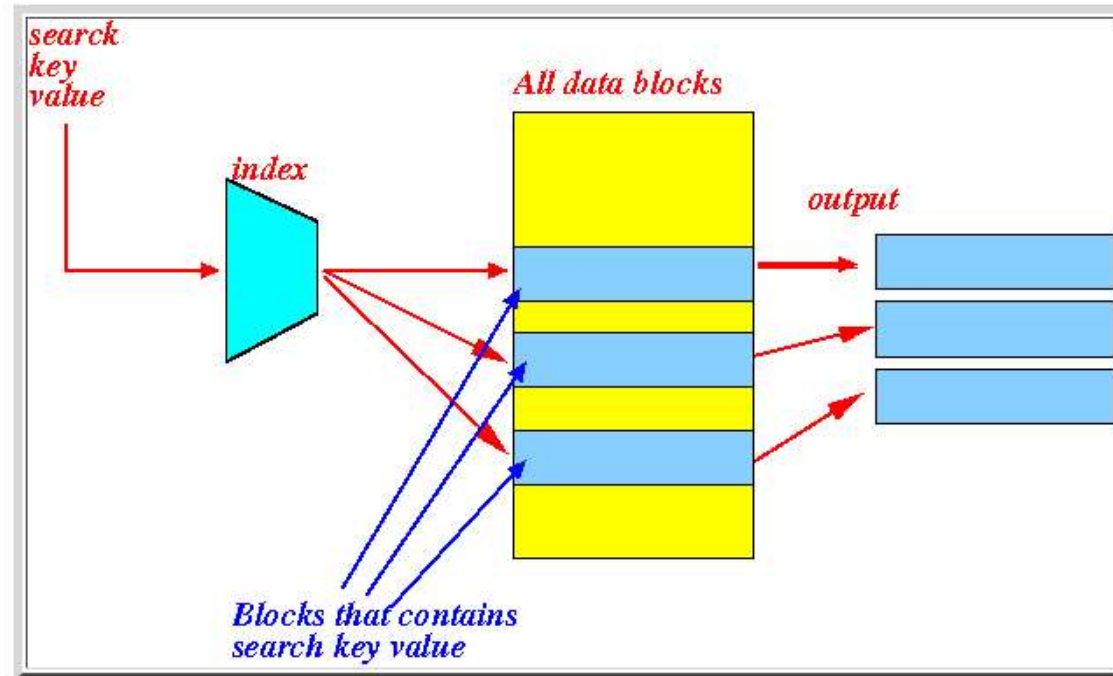
- Because the **search operation** *only* access **main memory**, the **search** is **relatively quick** (because you don't use **disk access** operations in the **search**)



Basic Concepts of Indexing

- Speed up data access

An **index** will **achieve** this effect:



In other words:

- An **index** performs the following **mapping**:

Search key value \longrightarrow List of blocks that contains the search key value



Index File

- An index file is a computer file with an index that allows easy random access to any record given its file key.
- An **index file** consists of records (called **index entries**) of the form

search-key	pointer
------------	---------

- Index files are typically much smaller than the original file
- Two basic kinds of indices:
 - **Ordered indices:** search keys are stored in sorted order
 - **Hash indices:** search keys are distributed uniformly across “buckets” using a “hash function”



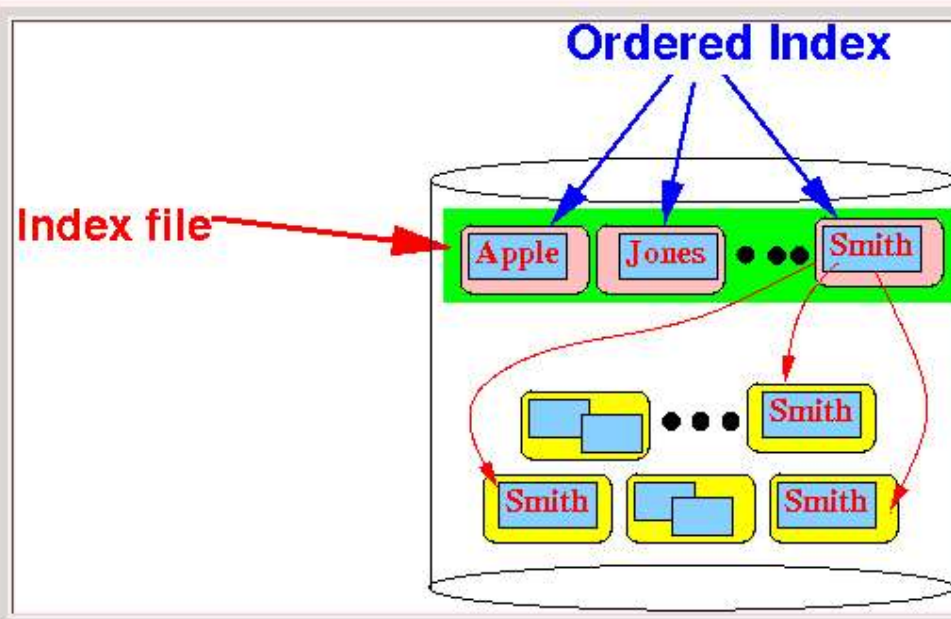
Type of index

- o There are **2 basic** kinds of **indexes**:

- **Ordered indexes**

- The **search keys** are **stored in sorted order** inside the **index file**

Schematically:



We can **find** a **search key** in the **index file** **quickly** using **binary search** !!!

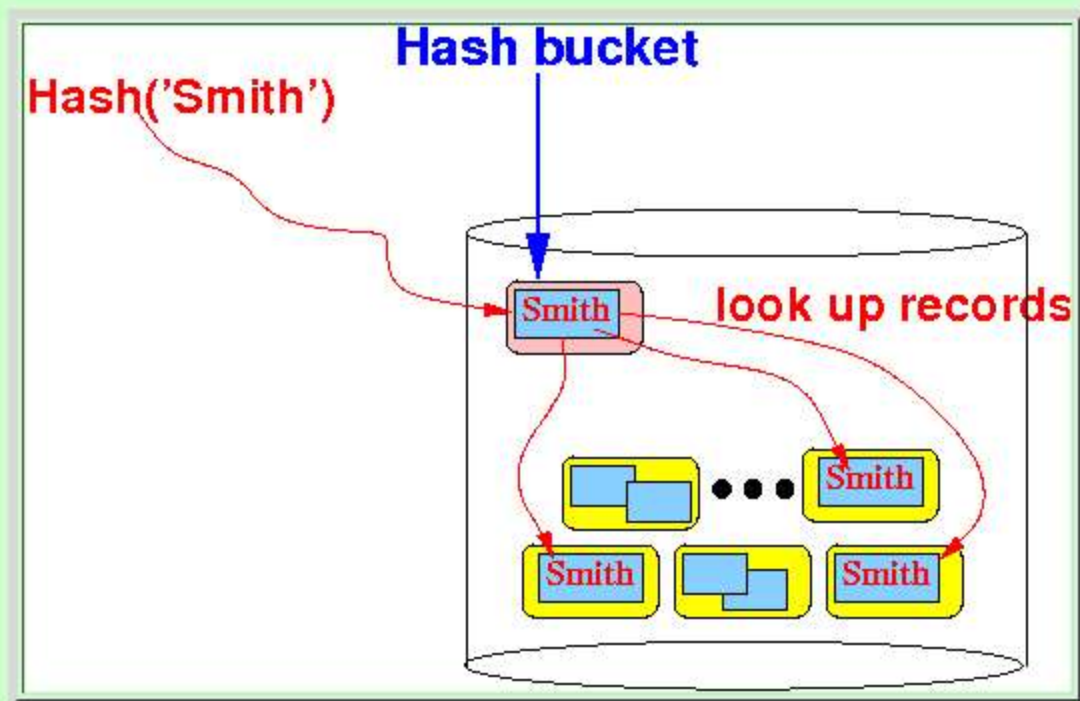


Hash Index

- The **search keys** are stored in **hash buckets**.

- Each **hash bucket** is stored as a **1 or more blocks** in the **index file**

Schematically:



We can **find** the **disk blocks** containing the **search key** *quickly* by using the **hash function**



Index Evaluation Metrics

- Access types supported efficiently. E.g.,
 - records with a specified value in the attribute
 - or records with an attribute value falling in a specified range of values.
- Access time
- Insertion time
- Deletion time
- Space overhead



Ordered Indices

- Index entries are sorted by search key values
 - E.g., author catalog in library.

- The **search keys** are stored in **sorted order** in the **index file**

Schematically:

Index file content:

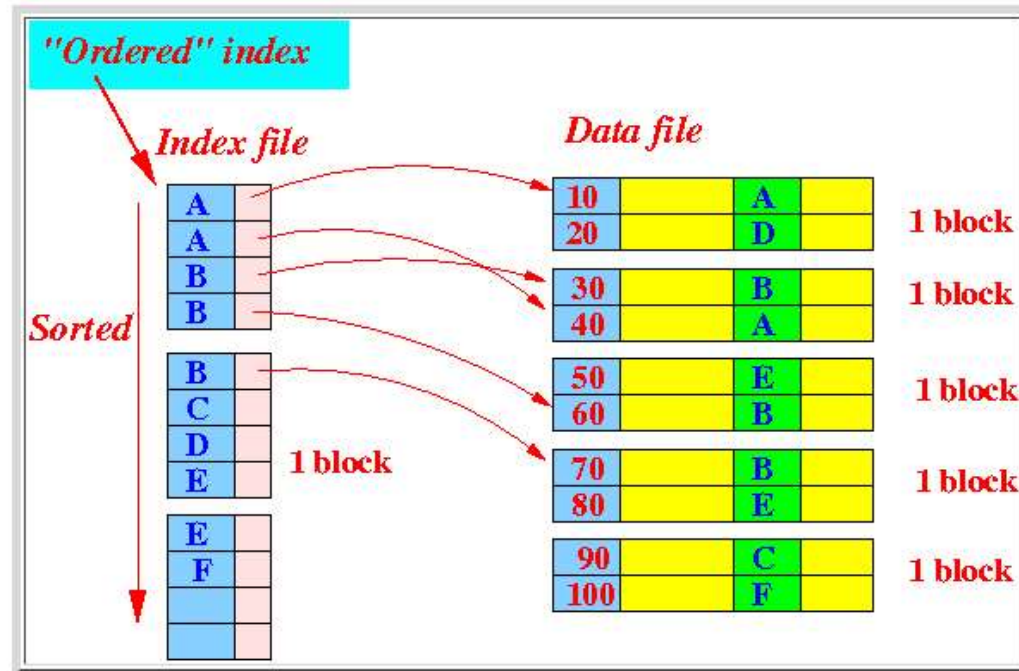
Search key	
cat	p1
dog	p2
elephant	p3
...	



Ordered Index

- Ordered index = an index file where the index entries are *sorted* (in the order of the search key)

Example: an ordered index



Note:

- Take *sorted* with a grain of salt

We will discuss **B-tree** that store the keys in a **tree structure**

- There is an **ordering** of the keys in a **B-tree**, but the **ordering** is not *sequential*



Ordered Indices

- **Primary index** (clustering index)
 - the index whose search key specifies the sequential order of the file
 - Index-sequential files: files ordered on a primary index
 - The search key of a primary index is usually the primary key

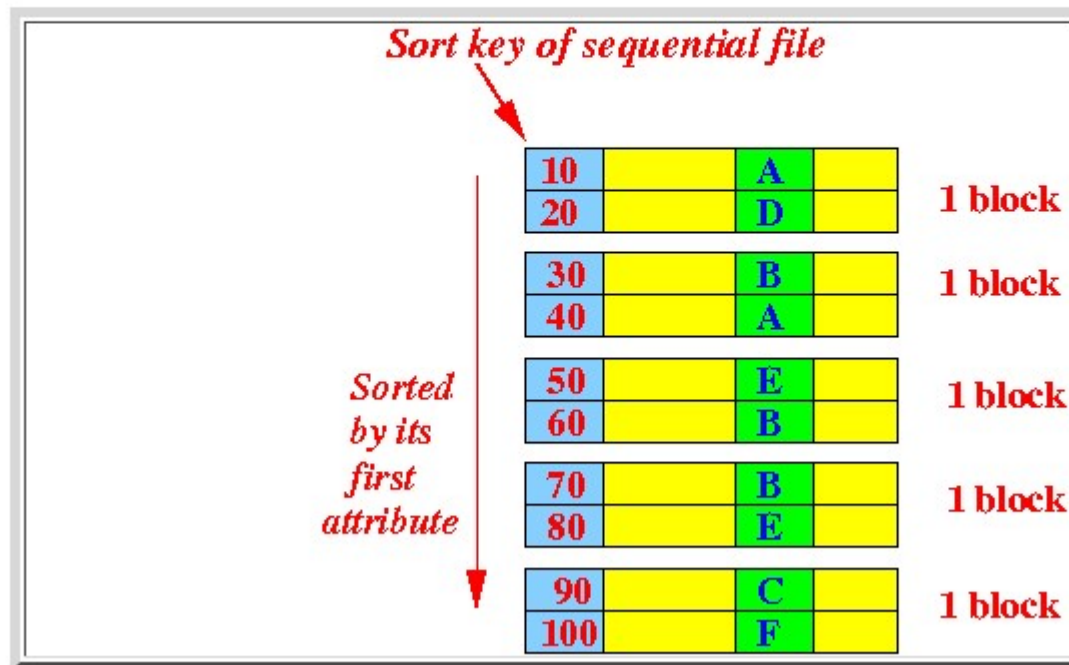
- **Secondary index** (non-clustering index)
 - An index whose search key specifies an order different from the sequential order of the file.



Sequential File (IBM's terminology....)

- **Sequential file** = a **file** whose **records** are **sorted** by some **attribute(s)** (usually its **primary key**)

Example: a **sequential file**



- **Sort key:**

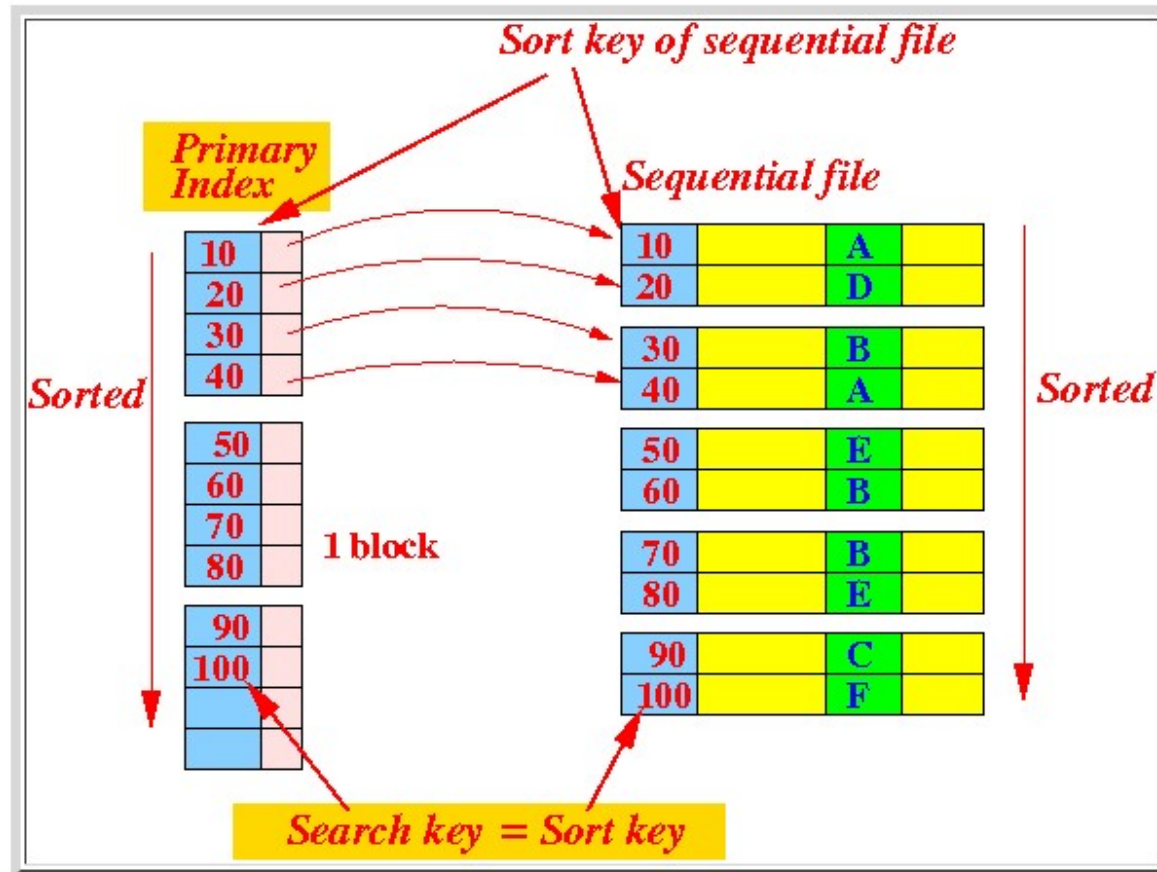
- **Sort key** = **field(s)** whose values are used to **sort/order** the records in a **sequential file**



Primary Index (clustering index)

■ **Primary index** = an **ordered index** whose **search key** is *also* the **sort key** used for the **sequential file**

Example: a **primary index**



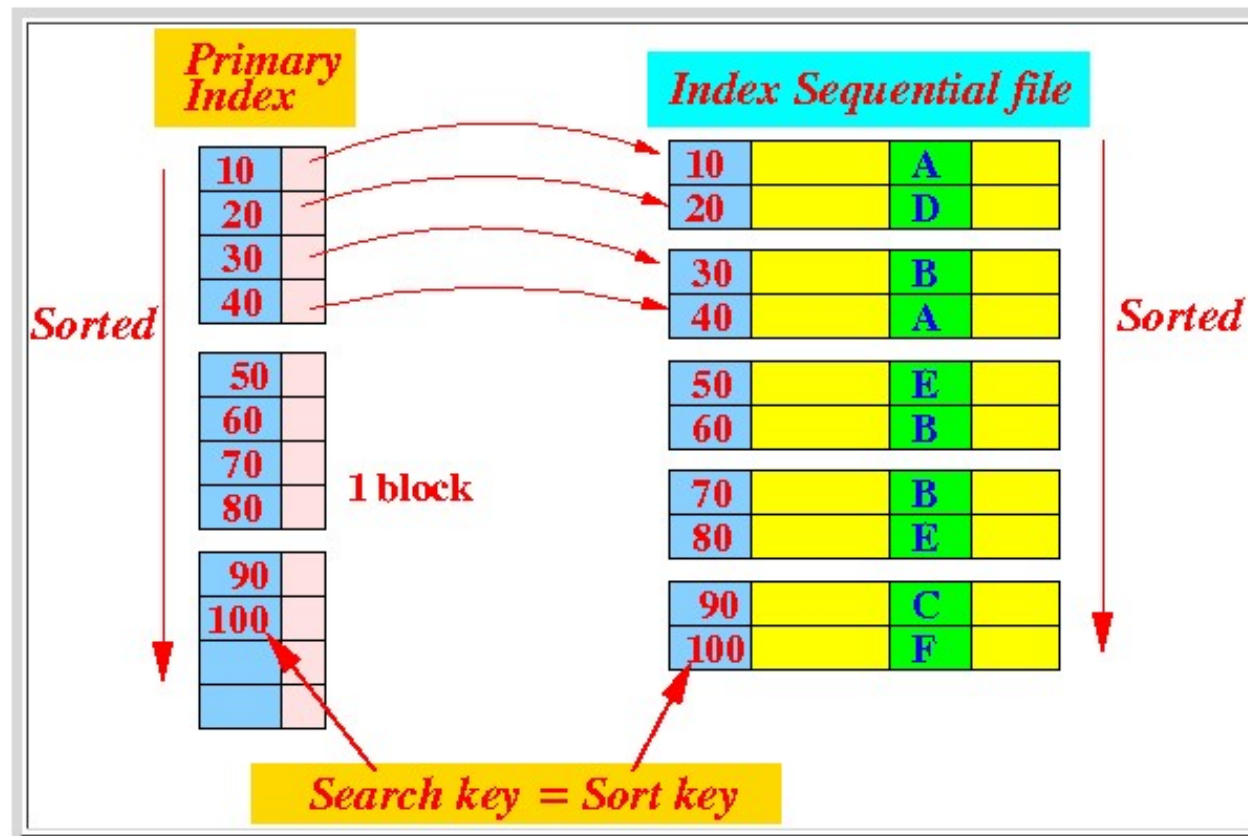


Index File

- o **Index sequential file**

▪ **Index sequential file** = a **sequential file** that has a *primary index*

Example: an **index sequential file**

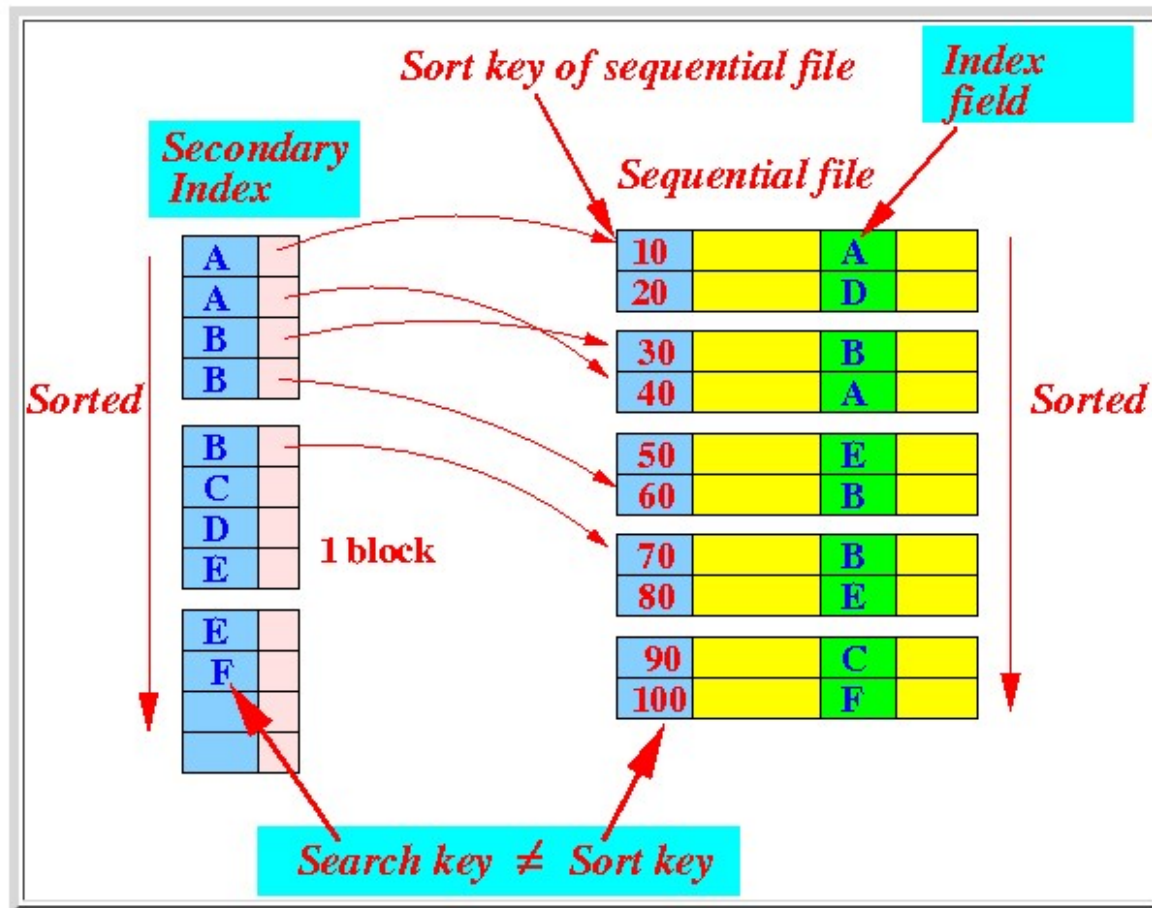




Secondary Index (non-clustering index)

▪ **Secondary index** = an **ordered index** whose **search key** is **NOT** the **sort key** used for the **sequential file**

Example: a **secondary index**





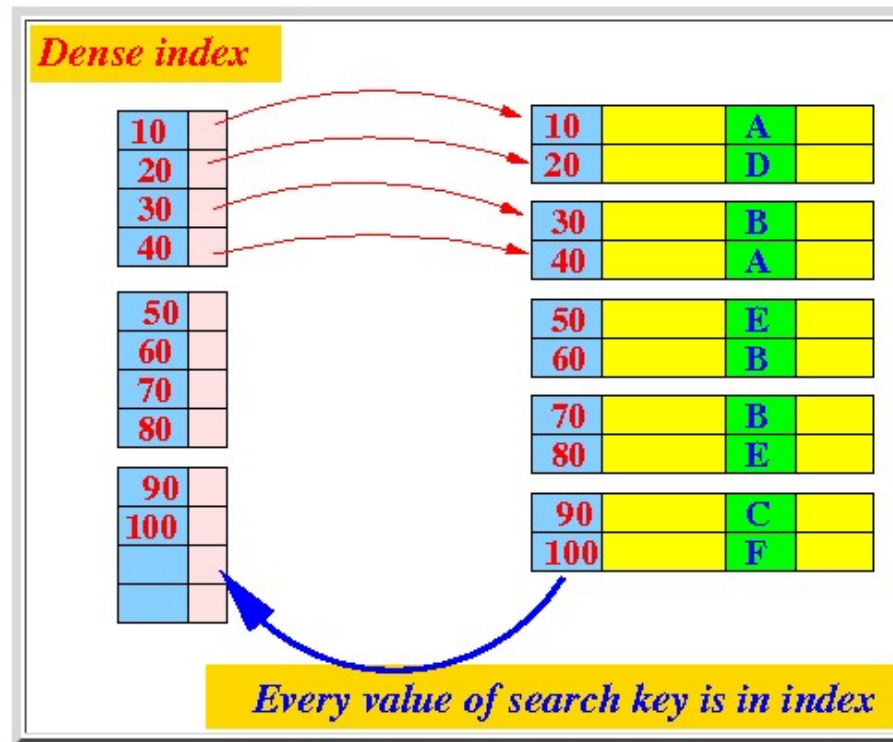
Dense Index

- **Dense index** = the index file contains an **index entry**:

-----+	-----+
Search key	block/record pointer
-----+	-----+

for **every value** of the **search key** in the data file

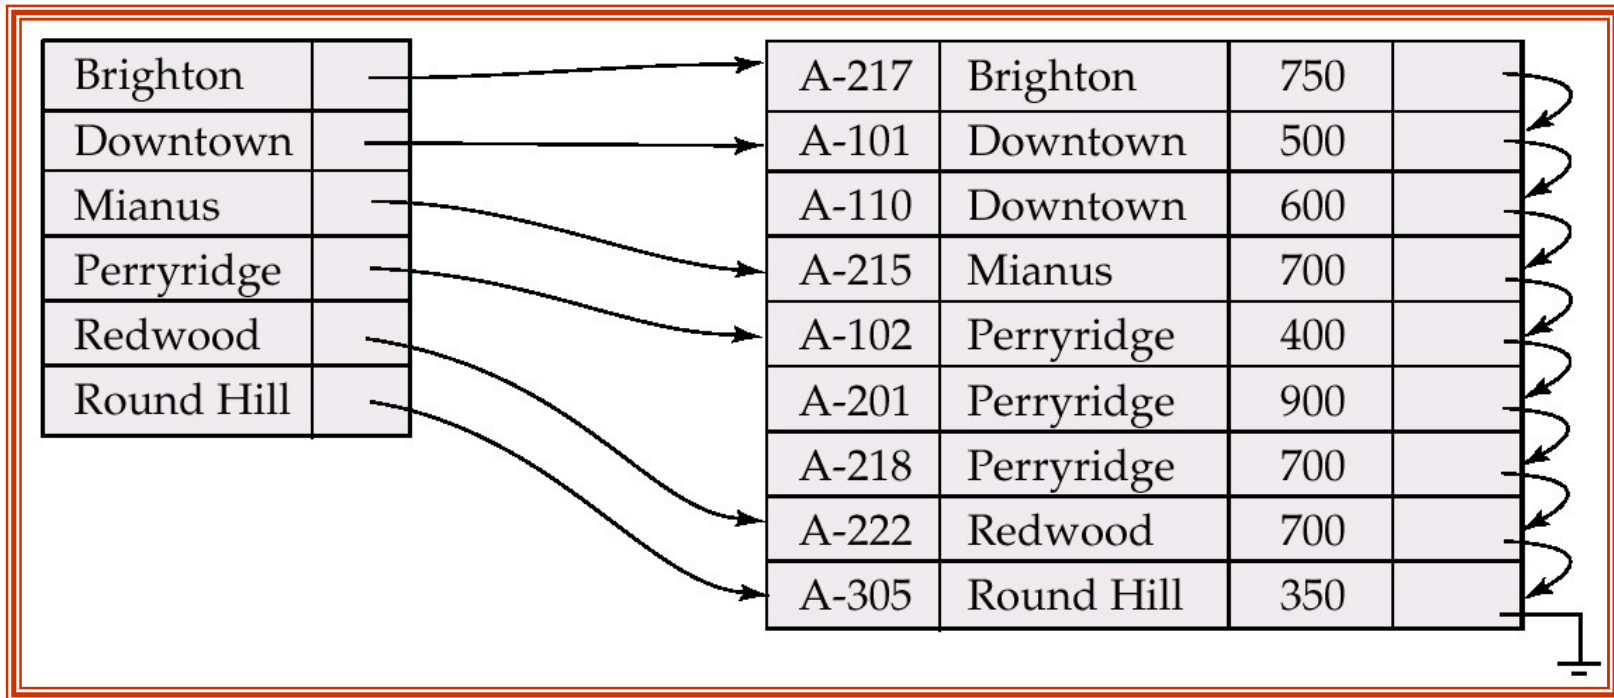
Example: a **dense** index





Dense Index Files

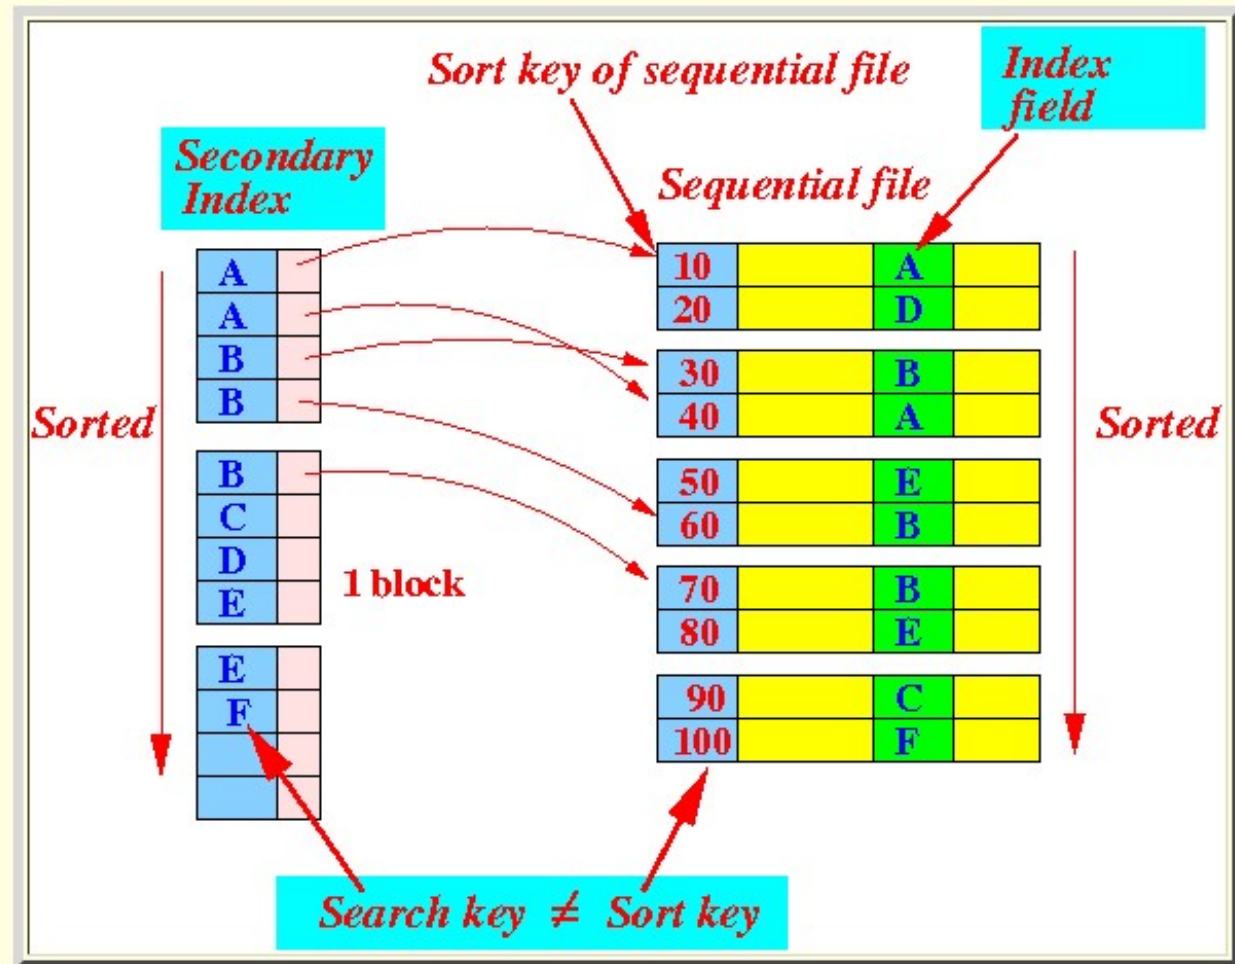
- **Dense index** — Index record exists for every search-key value in the file.





Note

- A *secondary index* is *always* a *dense index*:





Sparse Index Files

- **Sparse Index:** contains index records for only some search-key values.
 - Applicable when records are sequentially ordered on search-key
- To locate a record with search-key value K we:
 - Find index record with largest search-key value K
 - Search file sequentially starting at the record to which the index record points
- Less space and less maintenance overhead for insertions and deletions.
- Slower than dense index for locating records.
- Good tradeoff: sparse index with an index entry for every block in file, corresponding to least search-key value in the block.




Sparse Index

- **Sparse index** = the **index file's** search entries

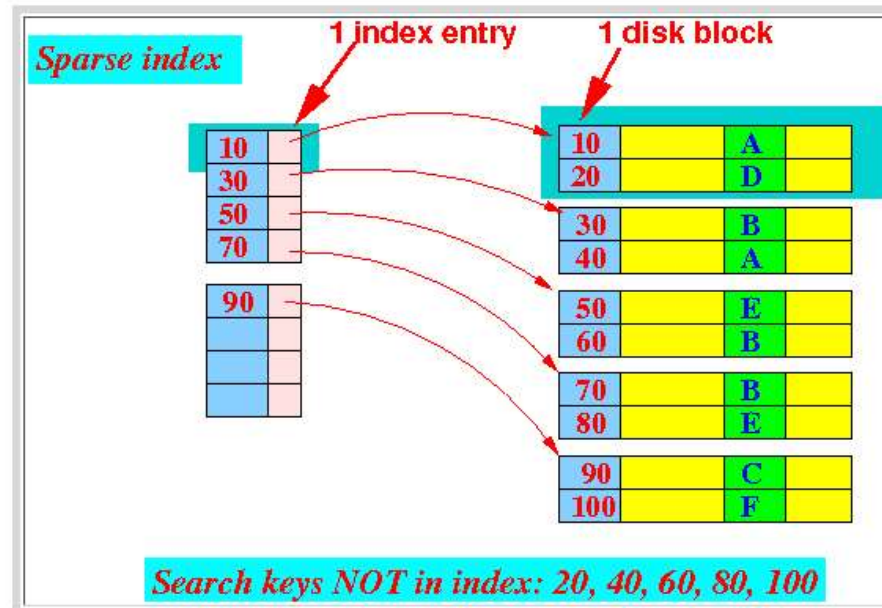


does **not** contain **every value** of the **search key** of the **data file**

- The **search key** of a **sparse index** must be the **sort key** of the **data file** !!!
- A **sparse index** contains:

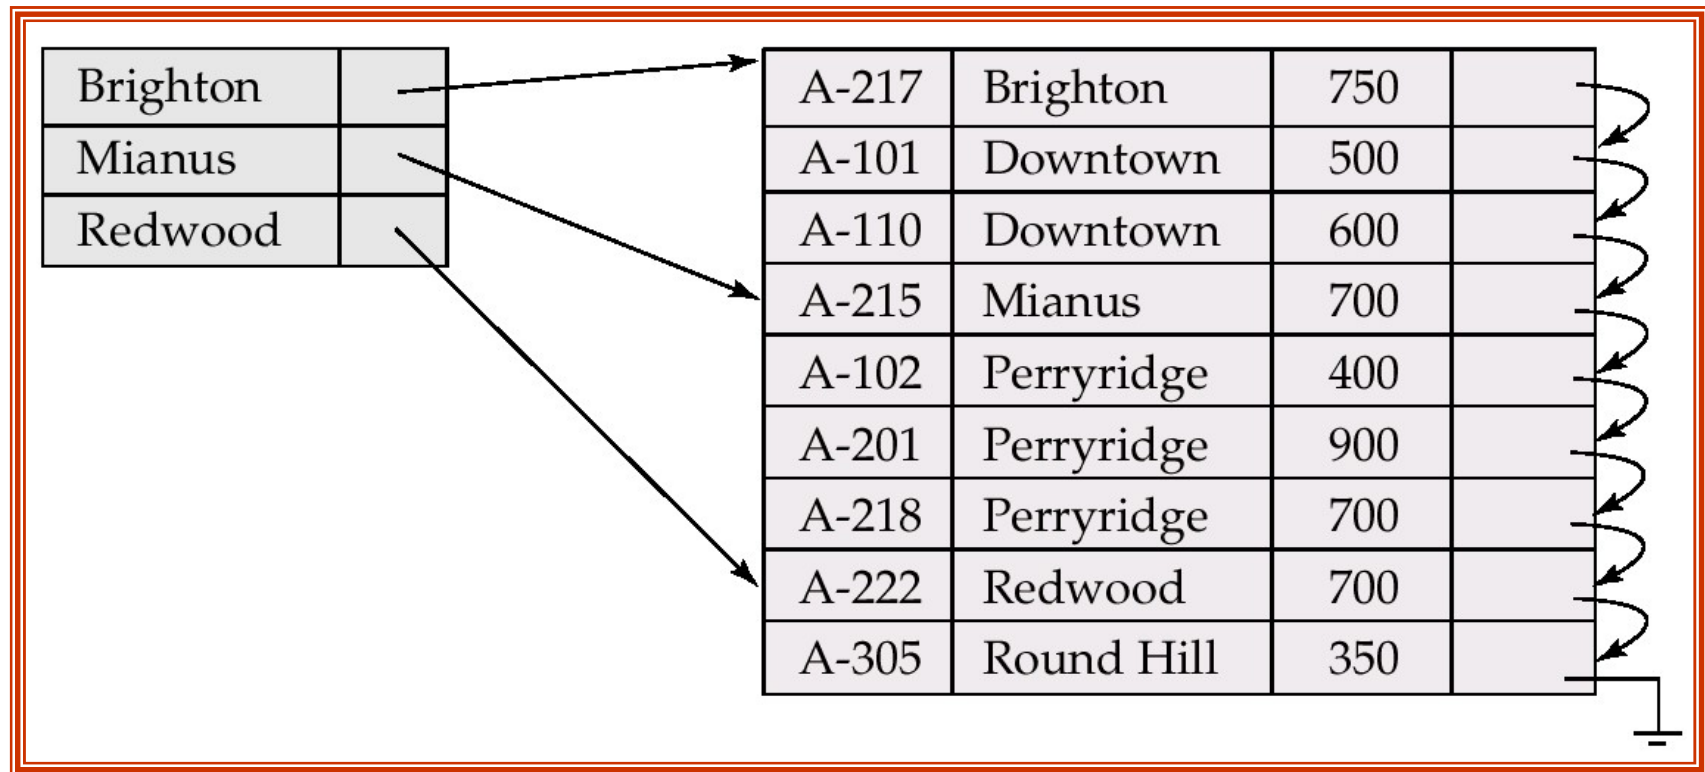


- o **Organization** of a **sparse index**





Example of Sparse Index Files



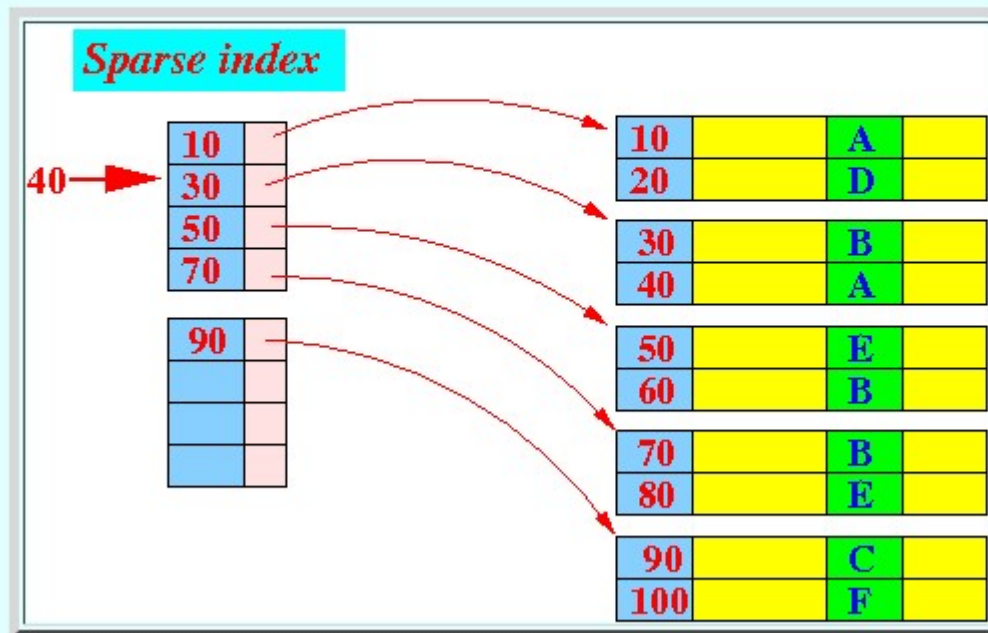


How to use a sparse index

Example:

- Look up the record with search key = 40

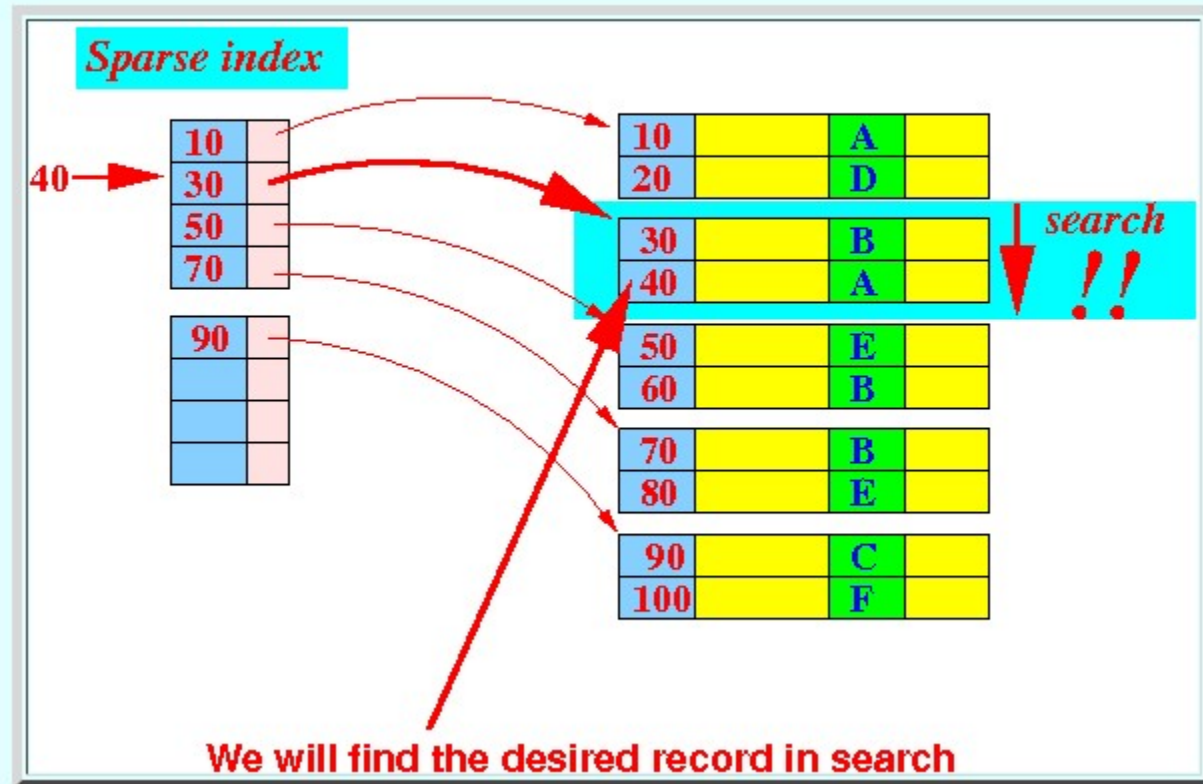
- Find the search key ≤ 40 :





How to use a sparse index

- **Search** in the **block** for **search key 40**:



- **Only primary indexes can be a sparse index**

(Using the **technique** above: the **sparse index** stores the **first key** in each **data block**)



Examples

Dense Index:

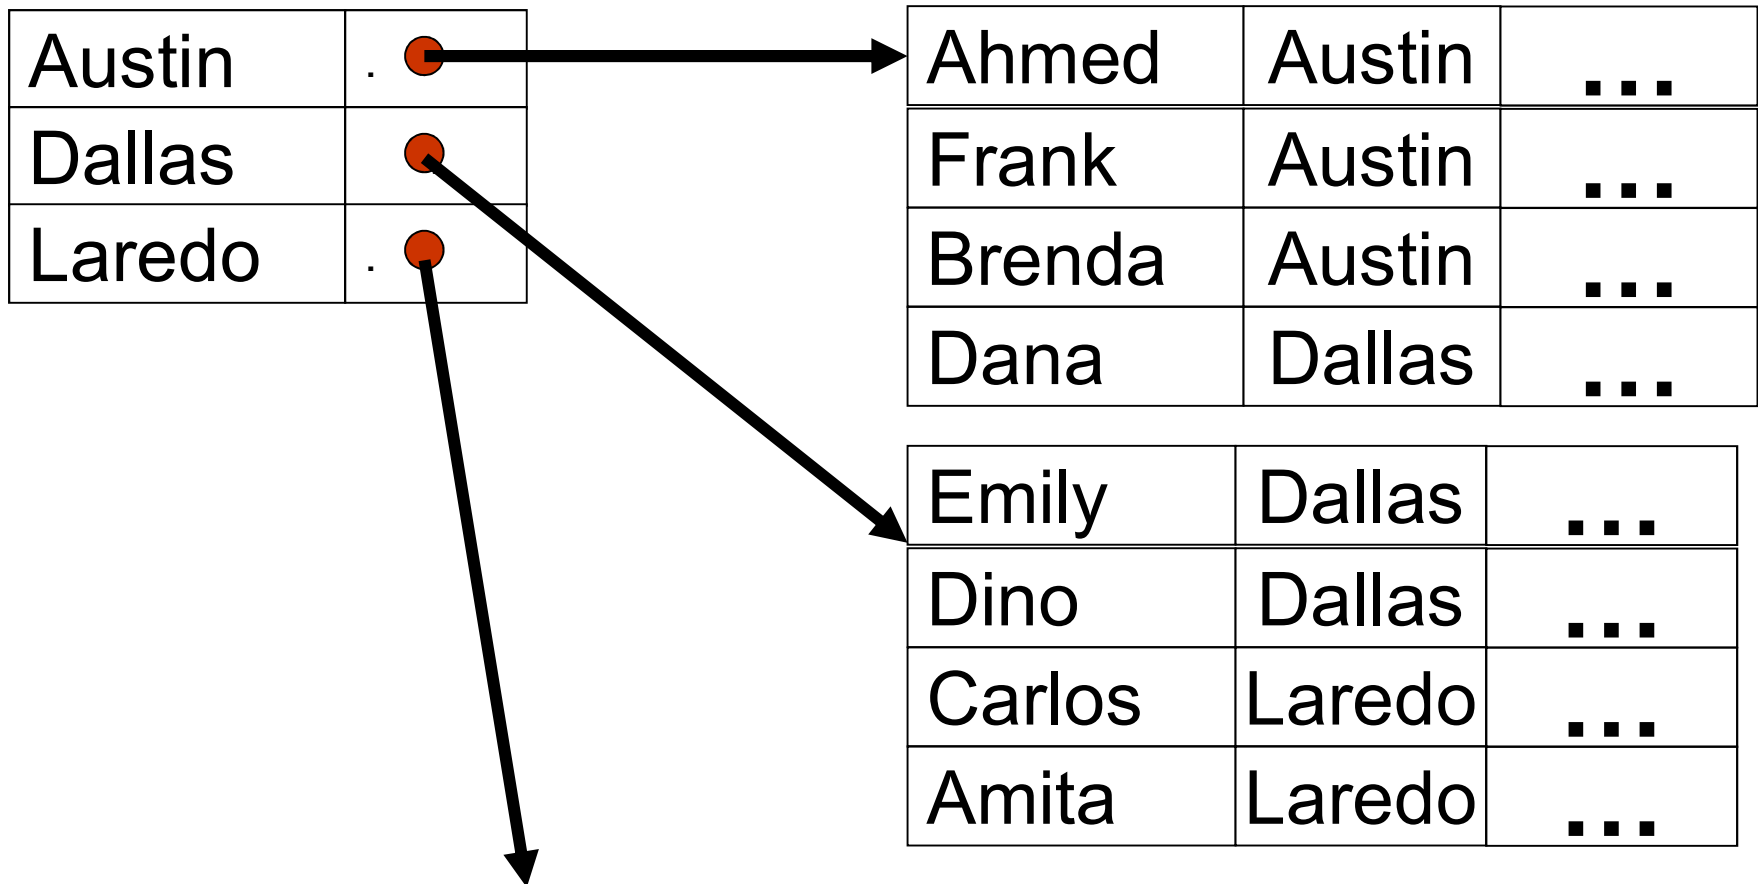
China	•	China	Beijing	3,705,386
Canada	•	Canada	Ottawa	3,855,081
Russia	•	Russia	Moscow	6,592,735
USA	•	USA	Washington	3,718,691

Sparse Index:

China	•	China	Beijing	3,705,386
Russia	•	Canada	Ottawa	3,855,081
USA	•	Russia	Moscow	6,592,735
	•	USA	Washington	3,718,691

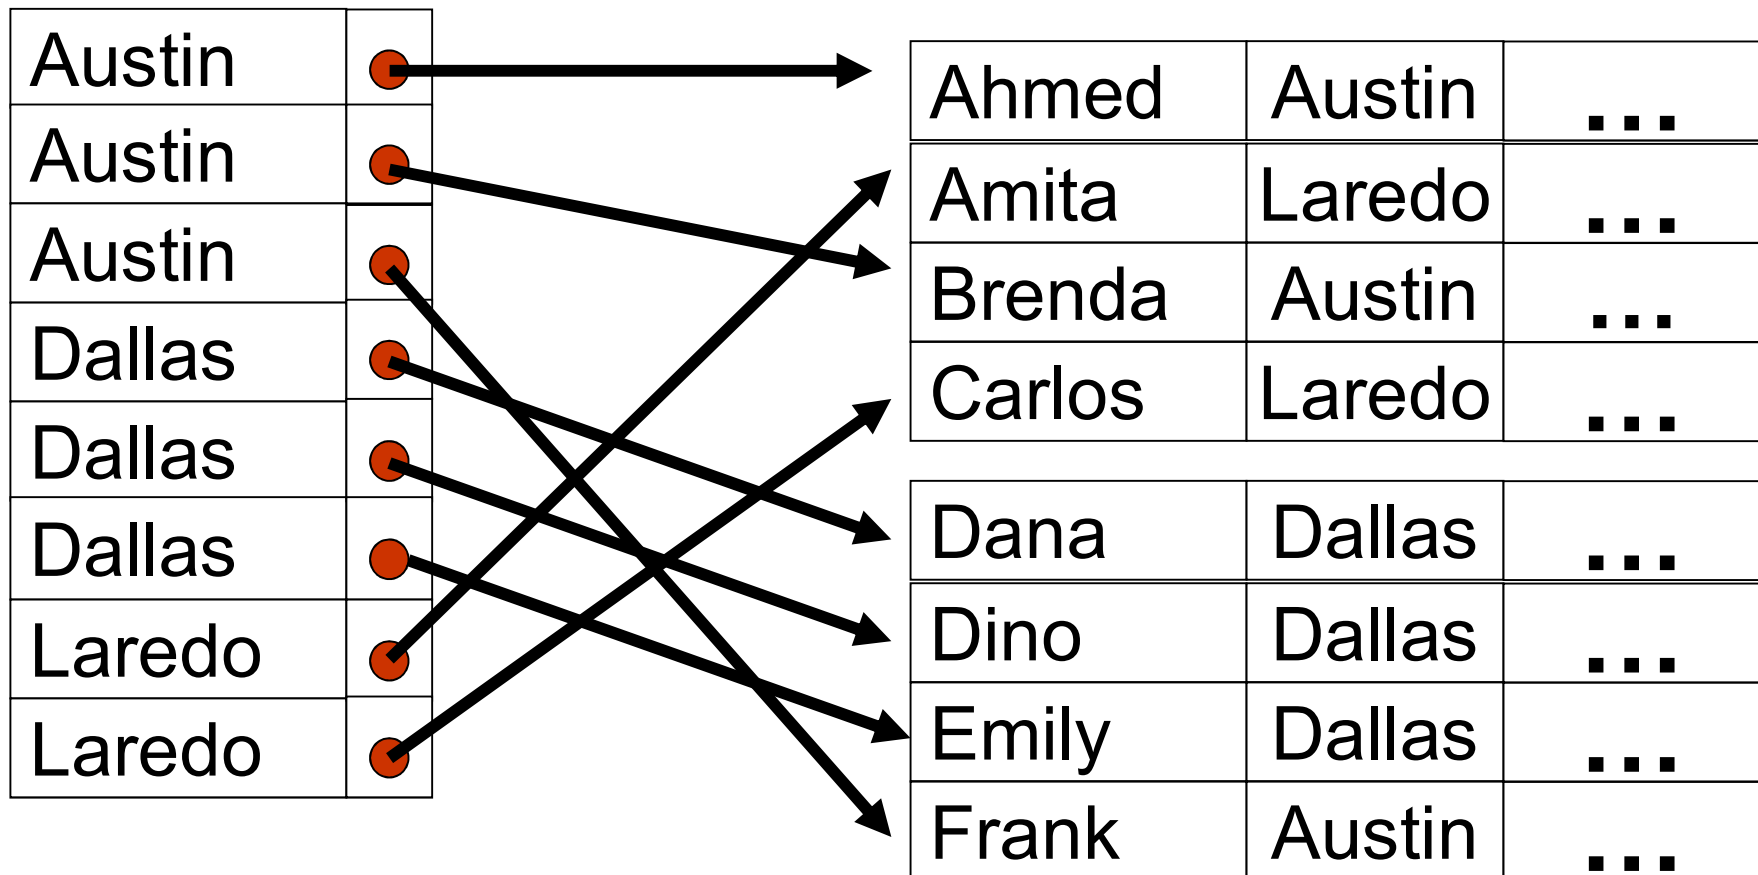


Sparse clustering index





Dense clustering index





Indexing the index

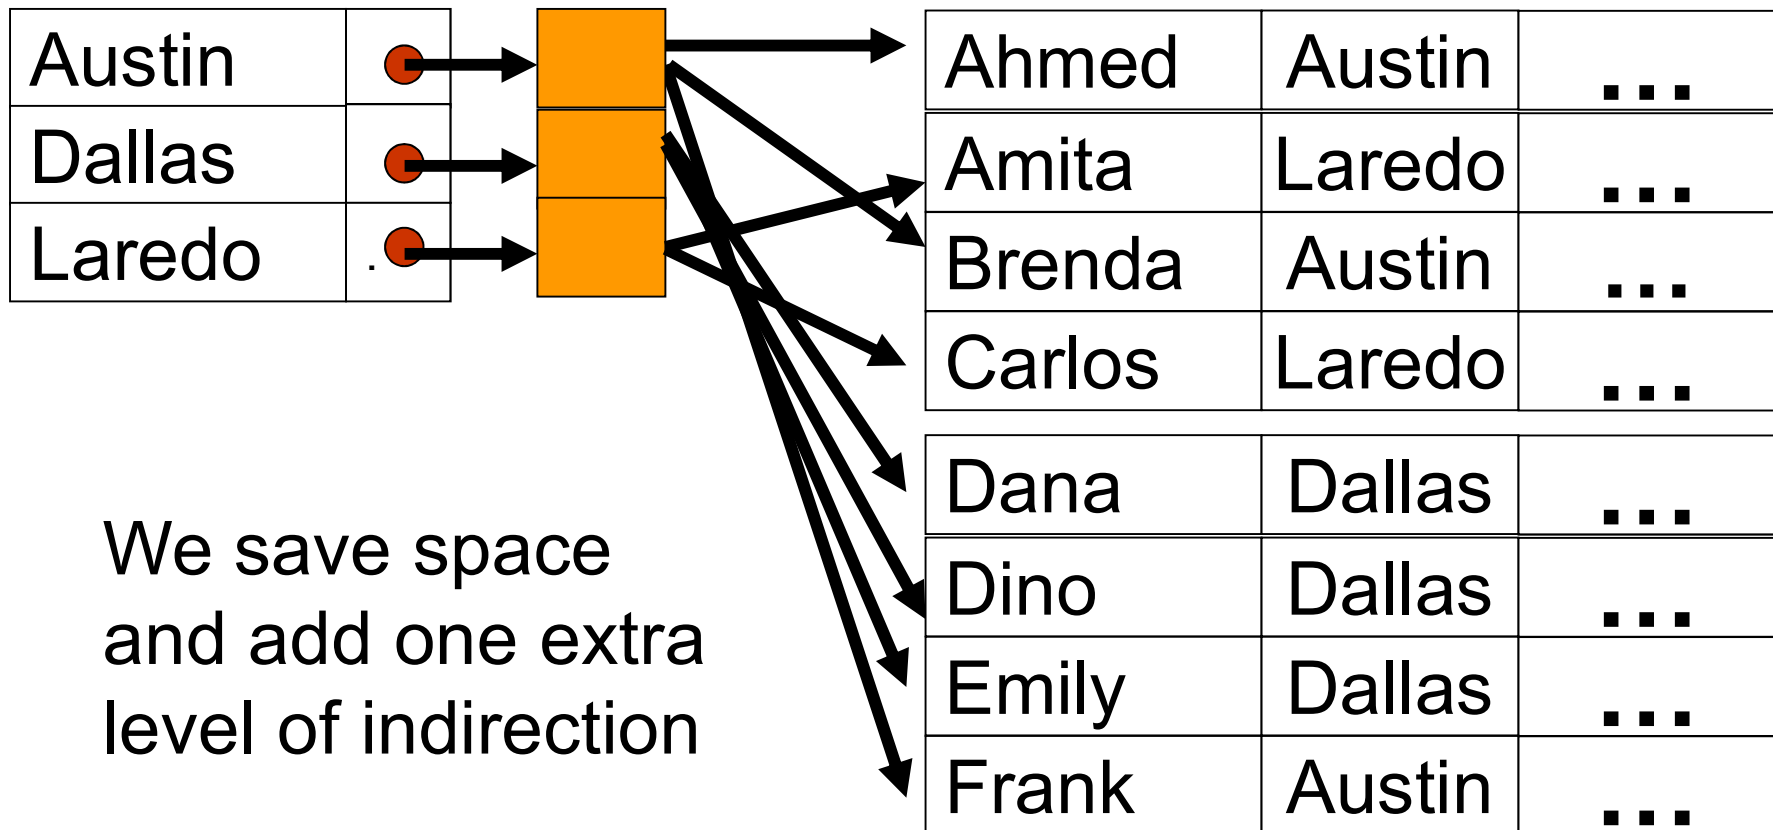
- When index is very large, it makes sense to index the index
 - Two-level or three-level index
 - Index at top level is called *master index*
 - Normally a *sparse index*

- "We can solve any problem by introducing an extra level of indirection, except of course for the problem of too many indirections."

David John Wheeler




Another realization





Two levels





Multilevel Index

o **Fact:**

- An **index file** is *also* a **data file**
(It contains **data** !!!)

Therefore:

- We can **create** an (primary) **index** on an **index file** !!!!

- **Multi-level** index....

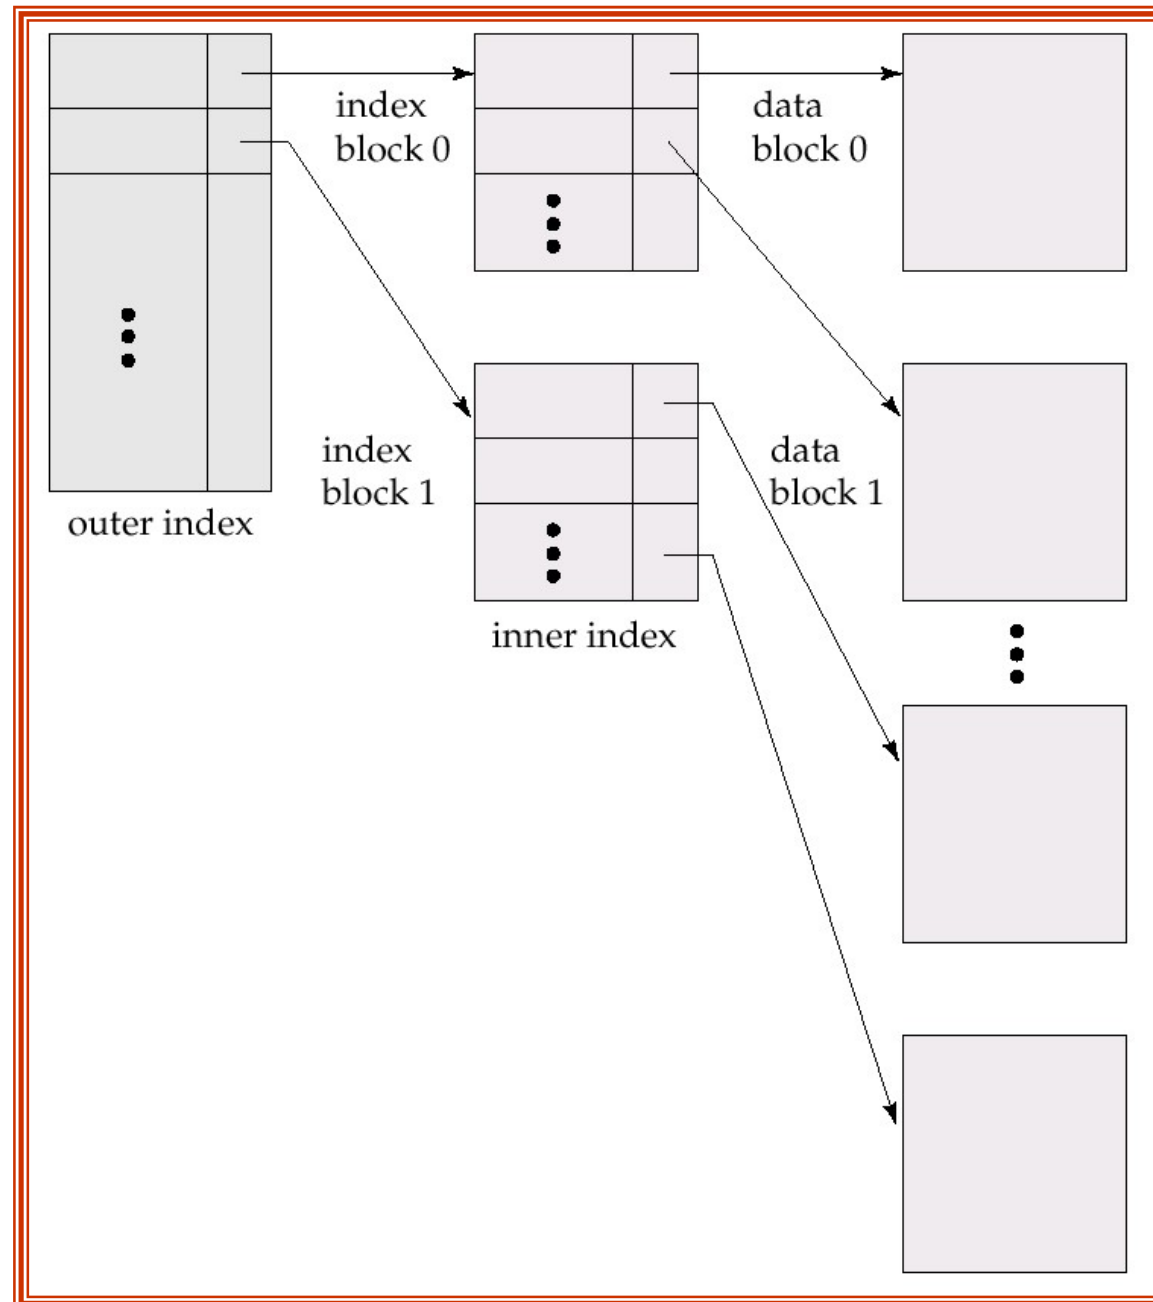


Multilevel Index

- If primary index does not fit in memory, access becomes expensive.
- To reduce number of disk accesses to index records, treat primary index kept on disk as a sequential file and construct a sparse index on it.
 - outer index – a sparse index of primary index
 - inner index – the primary index file
- If even outer index is too large to fit in main memory, yet another level of index can be created, and so on.
- Indices at all levels must be updated on insertion or deletion from the file.

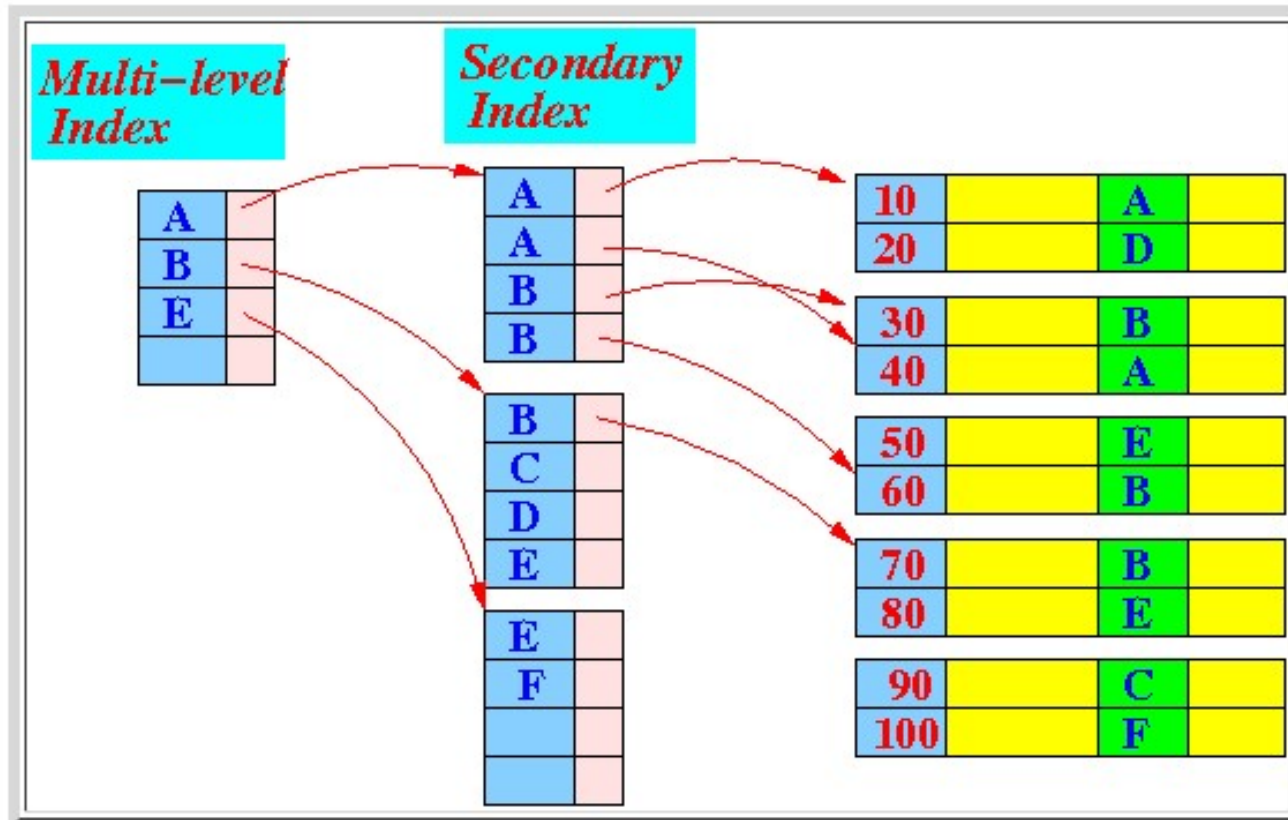


Multilevel Index (Cont.)





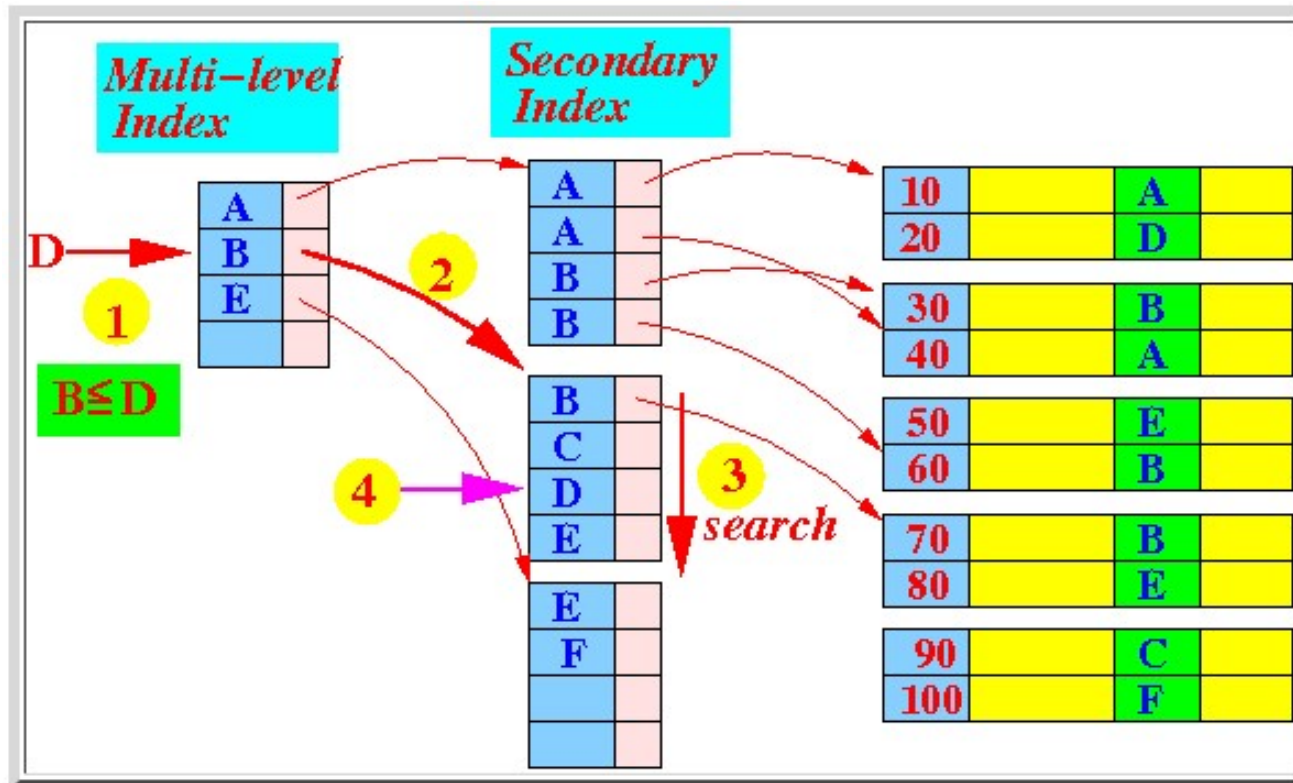
Multilevel Index Example





Multilevel Index

□ How to find the search key = D



o Note:

■ We will soon discuss a *dynamic multi-level* index data structure:

■ B-tree / B⁺-tree



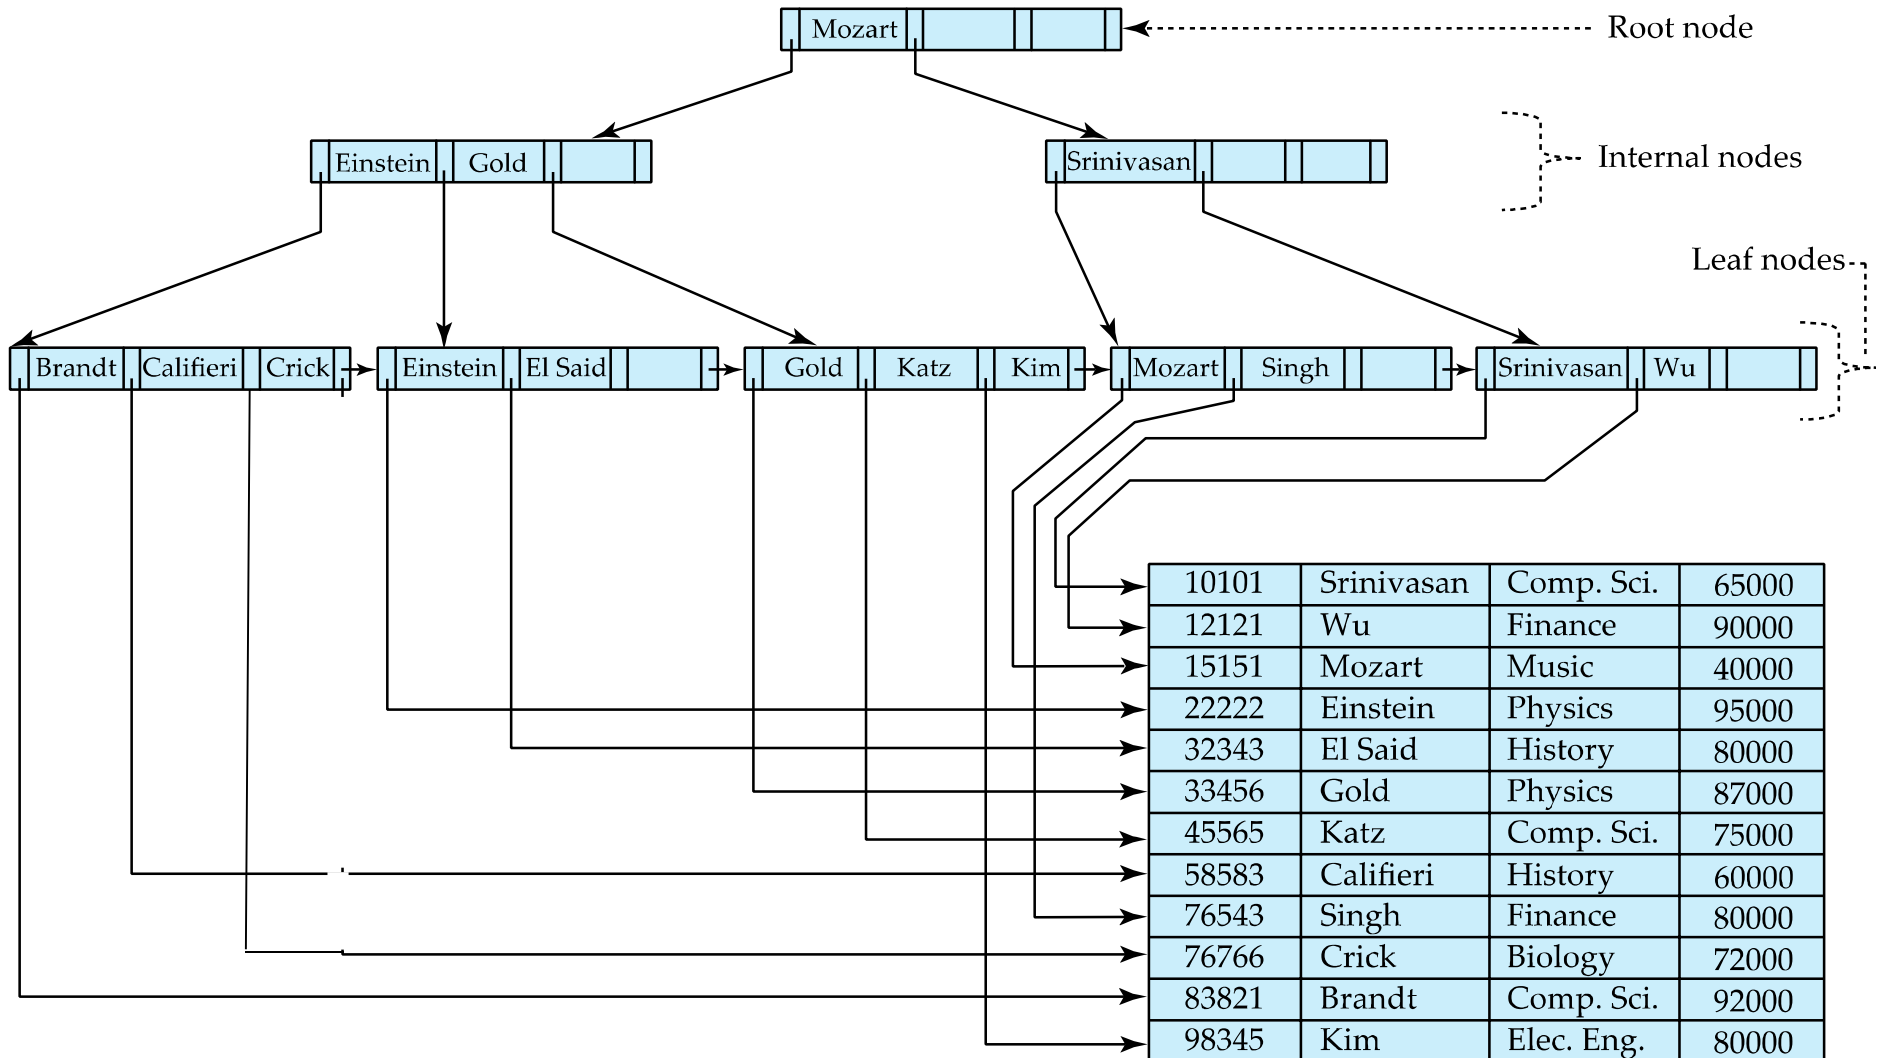
B-tree and B+-tree

Please watch the video:

<https://www.youtube.com/watch?v=aZjYr87r1b8>



Example of B⁺-Tree





Creation of Indices

- **E.g.**
 - create index** *takes_pk* **on** *takes* (*ID, course_ID, year, semester, section*)
 - drop index** *takes_pk*
- Most database systems allow specification of type of index, and clustering.
- Indices on primary key created automatically by all databases
 - Why?
- Some database also create indices on foreign key attributes
 - Why might such an index be useful for this query:
 - *takes* ⋈ $\sigma_{name='Shankar'}(student)$
- Indices can greatly speed up lookups, but impose cost on updates
 - Index tuning assistants/wizards supported on several databases to help choose indices, based on query and update workload



Index Definition in SQL

- Create an index

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

E.g.: **create index** *b-index* **on** *branch(branch_name)*

- Use **create unique index** to indirectly specify and enforce the condition that the search key is a candidate key is a candidate key.
 - Not really required if SQL **unique** integrity constraint is supported
- To drop an index

drop index <index-name>
- Most database systems allow specification of type of index, and clustering.



End of Chapter 14