

Big Data Course

Chap 6- Database Types

Electrical Engineering department of
Amirkabir University of technology

Dr. Mohammadreza Pourfard

August 2025

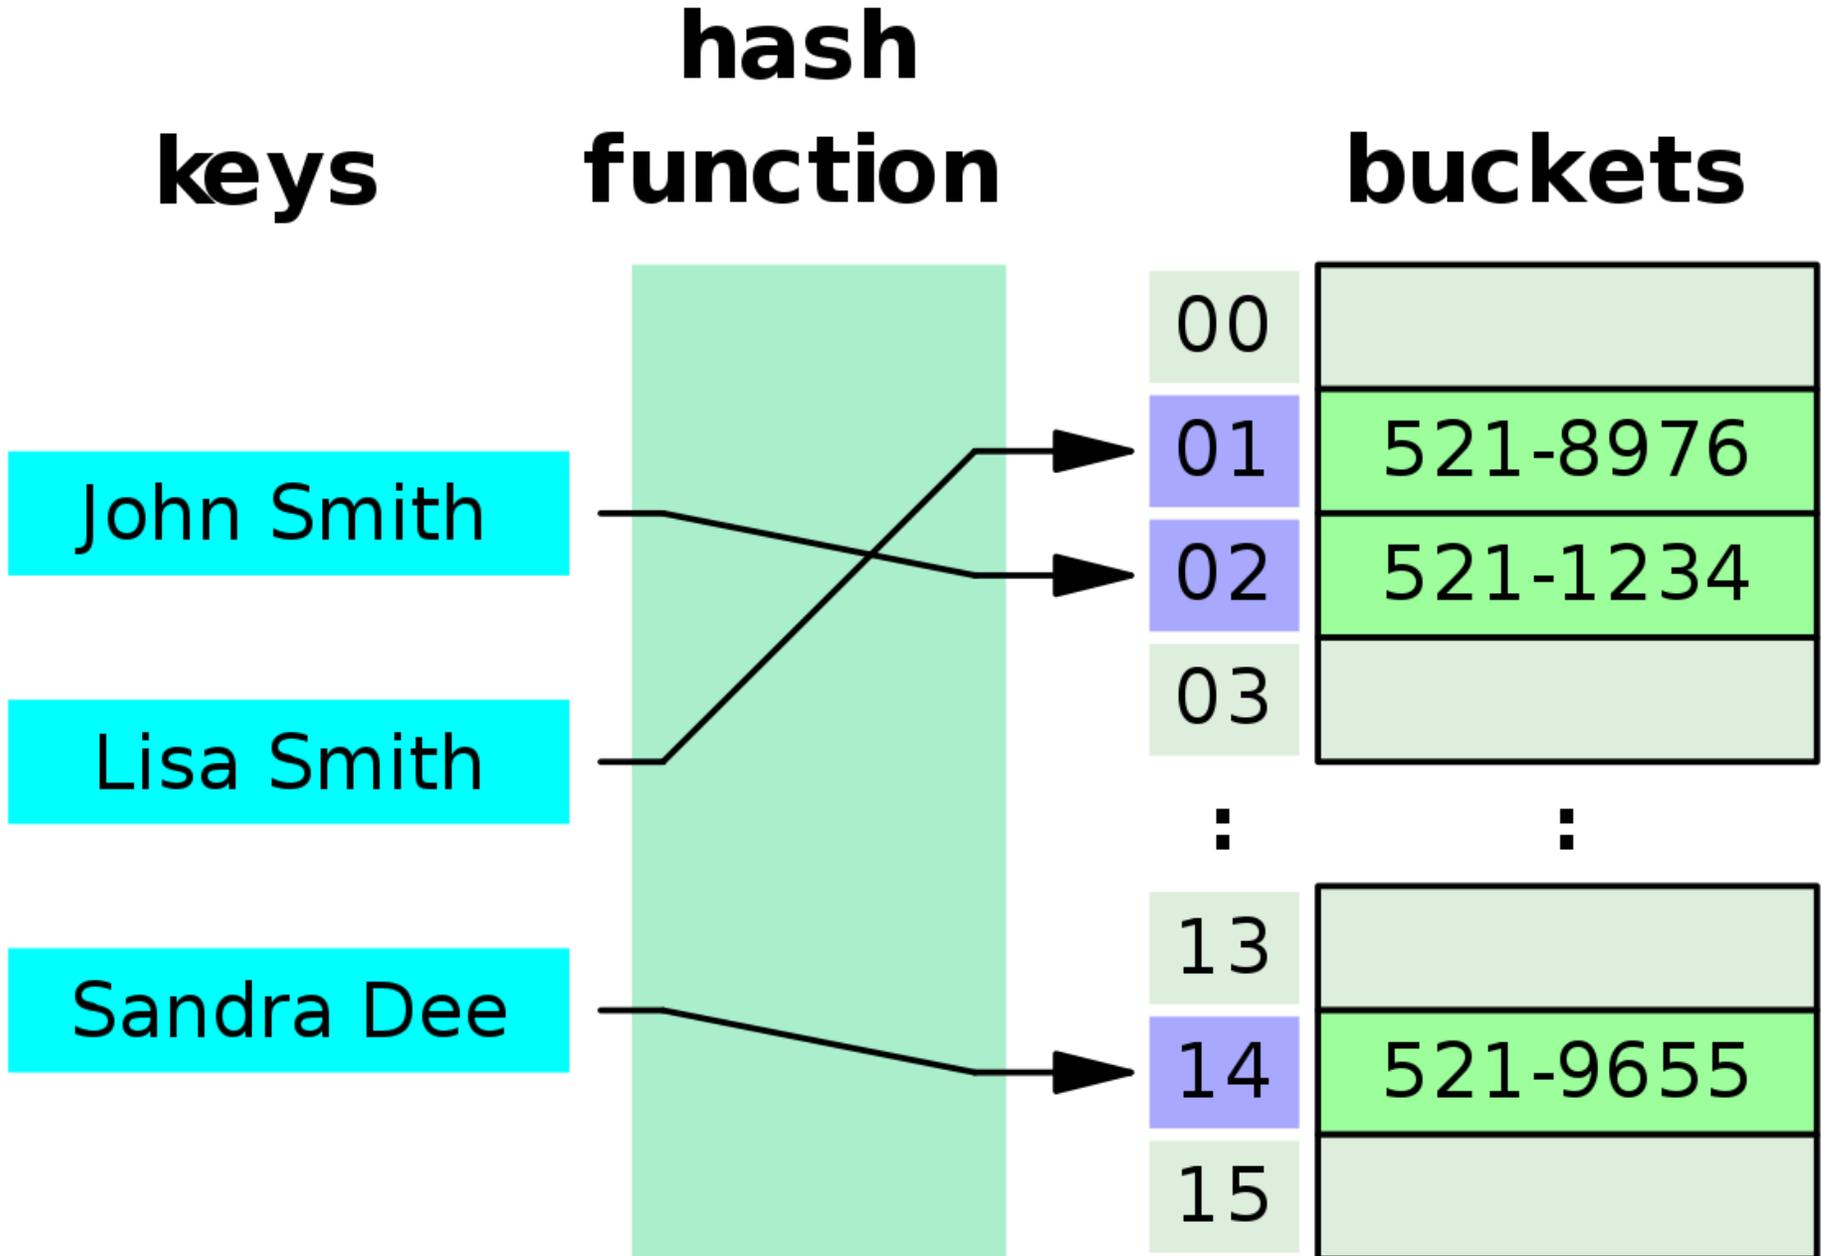
Index

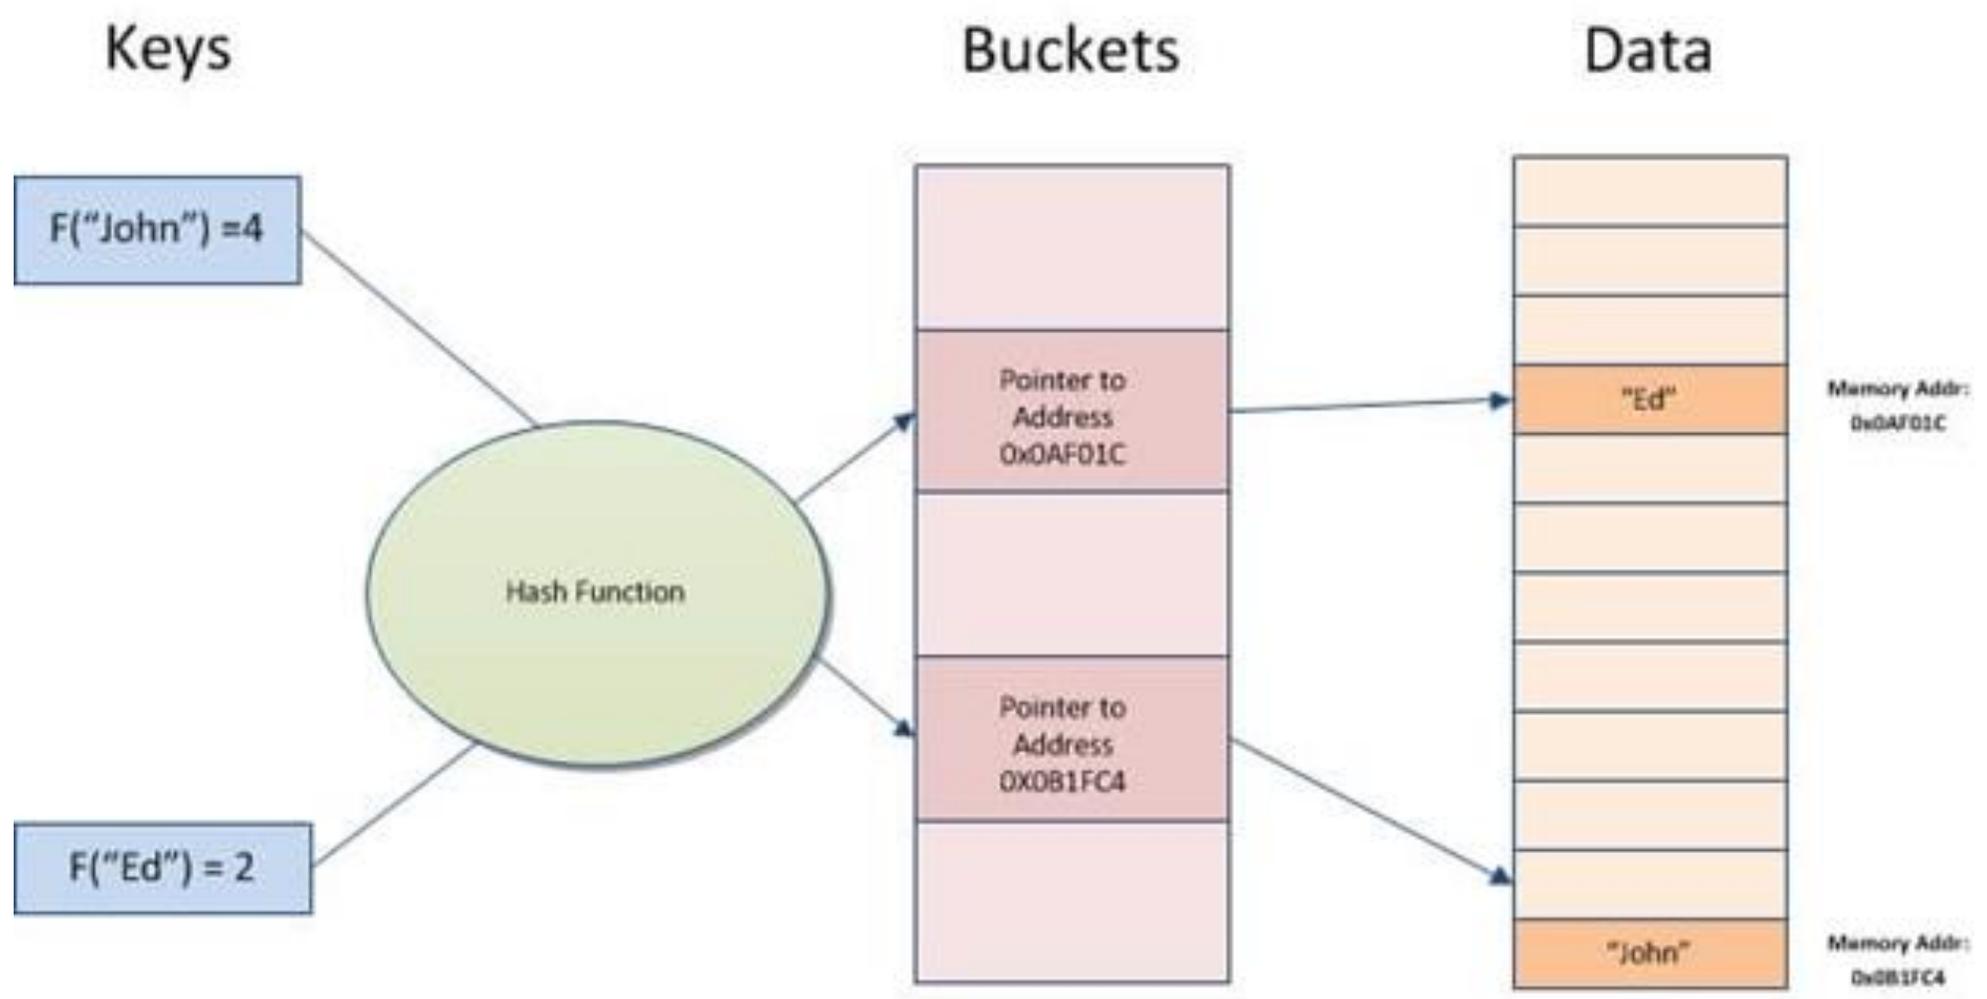

```
CREATE TABLE employees (  
    id INT PRIMARY KEY,  
    first_name VARCHAR(50),  
    last_name VARCHAR(50),  
    email VARCHAR(100),  
    department VARCHAR(50),  
    salary DECIMAL(10, 2)  
);
```

```
CREATE INDEX idx_last_name ON employees (last_name);
```

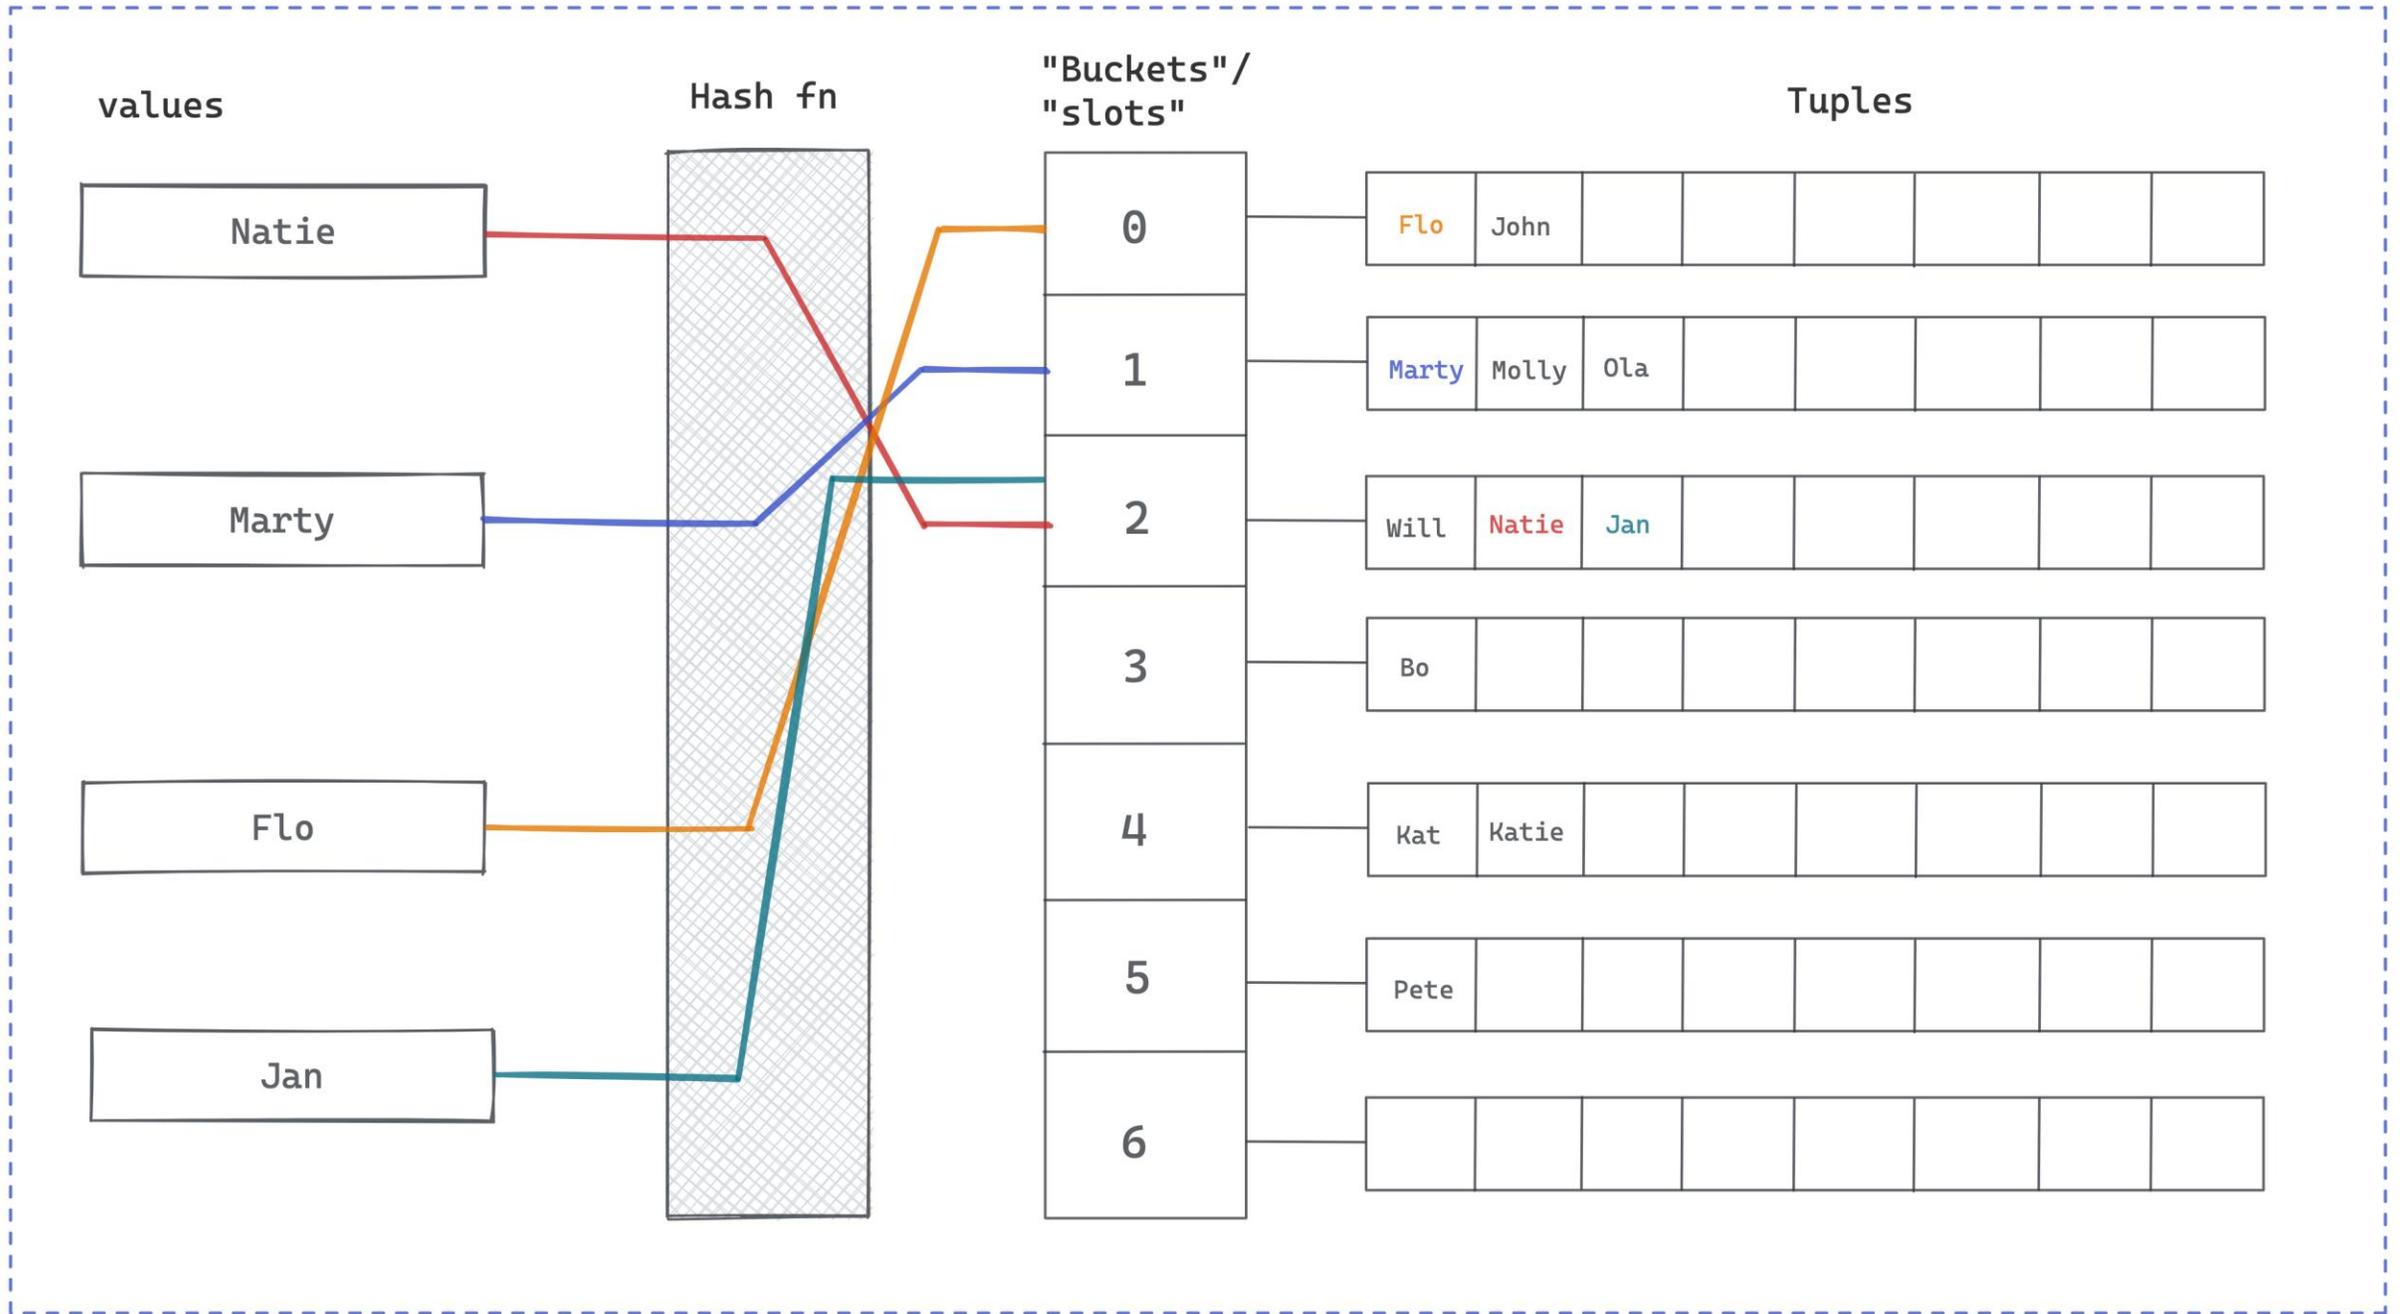
```
SELECT * FROM employees WHERE last_name = 'Smith';
```

```
CREATE INDEX idx_full_name ON employees (first_name, last_name);
```

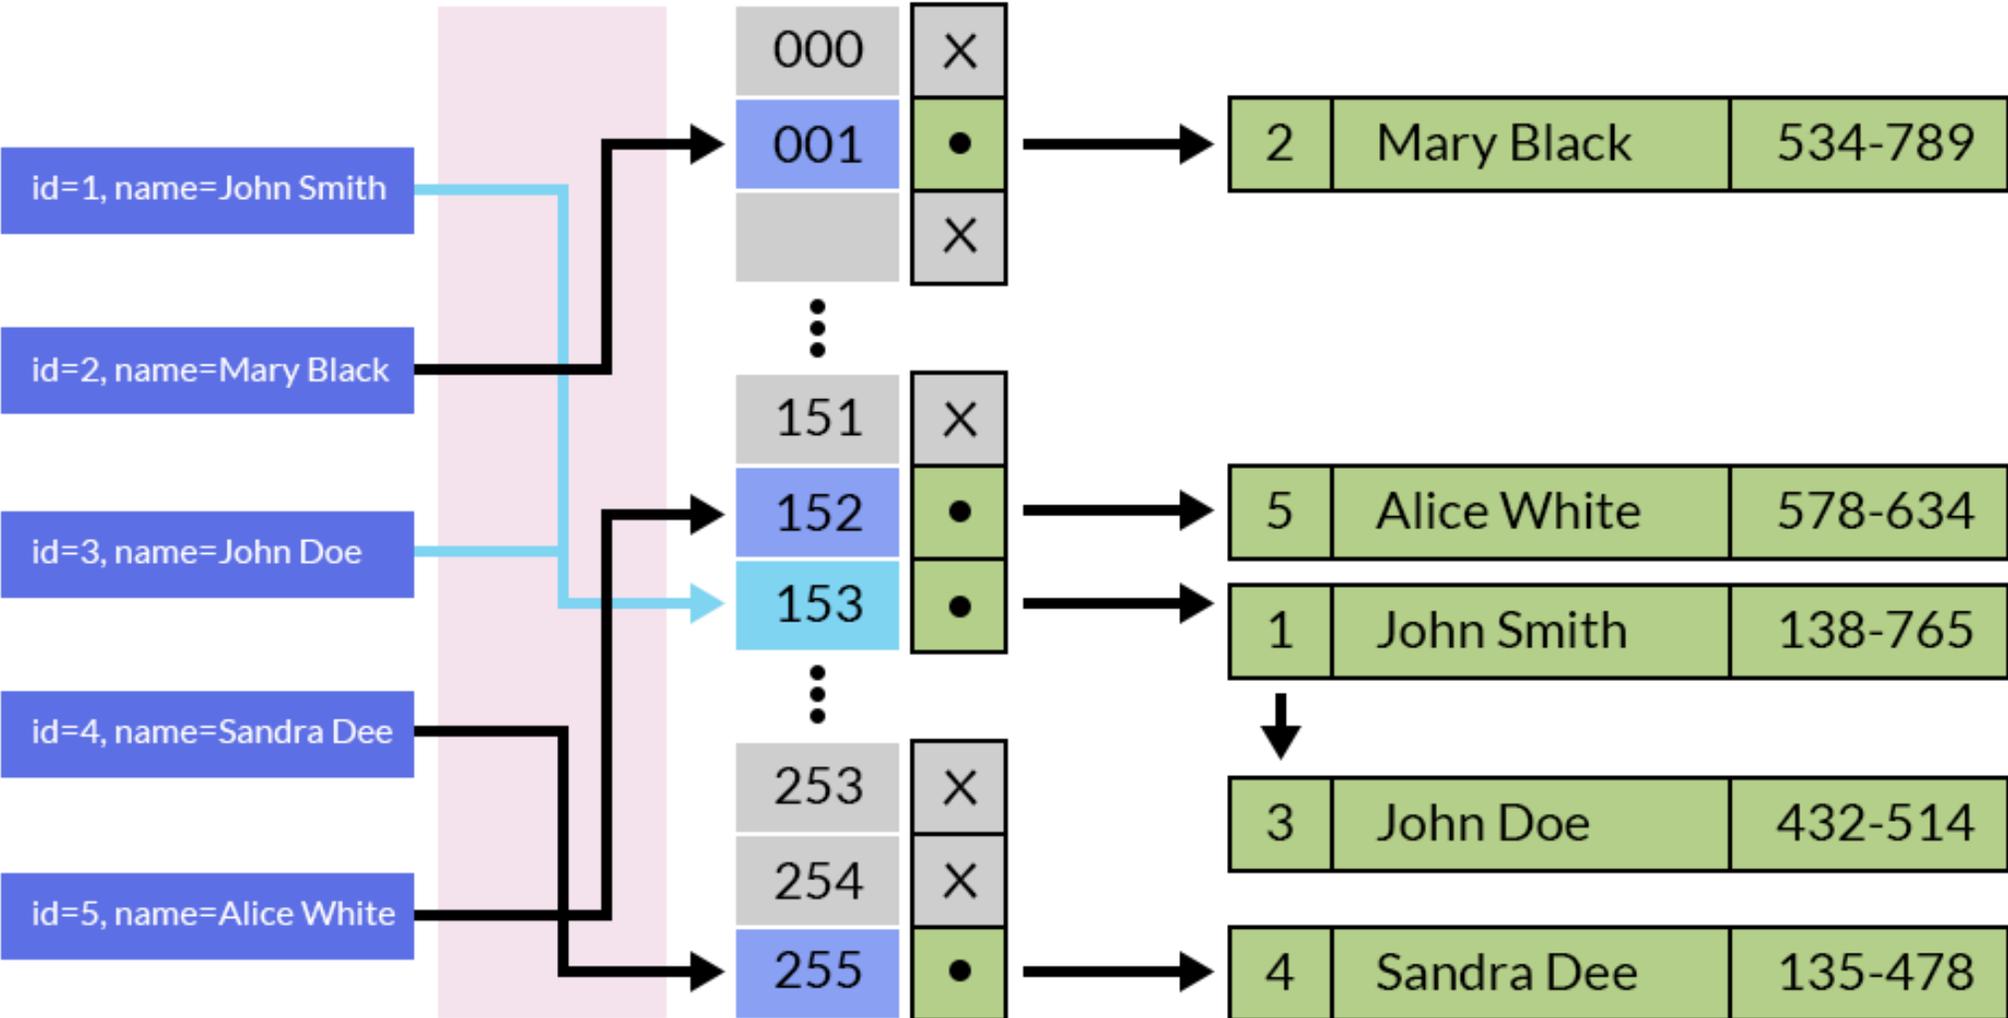




Hash Index



Hash function Buckets



How do Database Indexes Work?

- ◆ **Index Creation:** The database administrator creates an index on a specific column or set of columns.
- ◆ **Index Building:** The database management system builds the index by scanning the table and storing the values of the indexed column(s) along with a pointer to the corresponding data.
- ◆ **Query Execution:** When a query is executed, the database engine checks if an index exists for the requested column(s).
- ◆ **Index Search:** If an index exists, the database searches the index for the requested data, using the pointers to quickly locate the data.
- ◆ **Data Retrieval:** The database retrieves the requested data, using the pointers from the index.

Benefits of Database Indexes

- ◆ **Faster Query Performance:** Indexes can significantly improve query performance especially for large datasets by reducing the amount of data that needs to be scanned.
- ◆ **Reduced CPU Usage:** By reducing the number of rows that need to be scanned, indexes can decrease CPU usage and optimize resource utilization.
- ◆ **Rapid Data Retrieval:** Indexes enable quick data retrieval for queries that involve equality or range conditions on the indexed columns.
- ◆ **Efficient Sorting:** Indexes can also be used to efficiently sort data based on the indexed columns, eliminating the need for expensive sorting operations.
- ◆ **Better Data Organization:** Indexes can help maintain data organization and structure, making it easier to manage and maintain the database.

Types of Database Indexes

Indexes based on Structure and Key Attributes:

- ◆ **Primary Index:** Automatically created when a primary key constraint is defined on a table. Ensures uniqueness and helps with super-fast lookups using the primary key.
- ◆ **Clustered Index:** Determines the order in which data is physically stored in the table. A clustered index is most useful when we're searching in a range. Only one clustered index can exist per table.
- ◆ **Non-clustered or Secondary Index:** This index does not store data in the order of the index. Instead, it provides a list of virtual pointers or references to the location where the data is actually stored.

Types of Database Indexes

- ◆ **Indexes based on Data Coverage:**
- ◆ **Dense index:** Has an entry for every search key value in the table. Suitable for situations where the data has a small number of distinct search key values or when fast access to individual records is required.
- ◆ **Sparse index:** Has entries only for some of the search key values. Suitable for situations where the data has a large number of distinct search key values.

Specialized Index Types:

- ❖ **Bitmap Index:** Excellent for columns with low cardinality (few distinct values). Common in data warehousing.
- ❖ **Hash Index:** A index that uses a hash function to map values to specific locations. Great for exact match queries.
- ❖ **Filtered Index:** Indexes a subset of rows based on a specific filter condition. Useful to improve query speed on commonly filtered columns.
- ❖ **Covering Index:** Includes all the columns required by a query in the index itself, eliminating the need to access the underlying table data.
- ❖ **Function-based index:** Indexes that are created based on the result of a function or expression applied to one or more columns of a table.
- ❖ **Full-Text Index:** A index designed for full-text search, allowing for efficient searching of text data.
- ❖ **Spatial Index:** Used for indexing geographical data types.

What Data Structure do Indexes use?

- ◇ **B-Tree (Balanced Tree)**
- ◇ **Hash Tables**
- ◇ **Bitmaps**
- ◇ **Burrows–Wheeler (Strings of Characters)**

Burrows–Wheeler transform

- The **Burrows–Wheeler transform (BWT)** rearranges a [character string](#) into runs of similar characters, in a manner that can be reversed to recover the original string.

Burrows–Wheeler transform

Transformation				
1. Input	2. All rotations	3. Sort into lexical order	4. Take the last column	5. Output
<div style="border: 1px solid gray; padding: 10px; width: fit-content; margin: auto;"> \wedgeBANANA\$ </div>	<div style="border: 1px solid gray; padding: 10px; width: fit-content; margin: auto;"> \wedgeBANANA\$ \$$\wedge$BANANA A\$$\wedge$BANAN NA\$$\wedge$BANA ANA\$$\wedge$BAN NANA\$$\wedge$BA ANANA\$$\wedge$B BANANA\$$\wedge$ </div>	<div style="border: 1px solid gray; padding: 10px; width: fit-content; margin: auto;"> ANANA\$$\wedge$B ANA\$$\wedge$BAN A\$$\wedge$BANAN BANANA\$$\wedge$ NANA\$$\wedge$BA NA\$$\wedge$BANA \wedgeBANANA\$ \$$\wedge$BANANA </div>	<div style="border: 1px solid gray; padding: 10px; width: fit-content; margin: auto;"> ANANA\$$\wedge$B ANA\$$\wedge$BAN A\$$\wedge$BANAN BANANA\$$\wedge$ NANA\$$\wedge$BA NA\$$\wedge$BANA \wedgeBANANA\$ \$$\wedge$BANANA </div>	<div style="border: 1px solid gray; padding: 10px; width: fit-content; margin: auto;"> BNN\wedgeAA\$A </div>

Inverse transformation

Input

BNN^AA\$A

Add 1

B
N
N
^
A
A
\$
A

Sort 1

A
A
A
B
N
N
^
\$

Add 2

BA
NA
NA
^B
AN
AN
\$^
A\$

Sort 2

AN
AN
A\$
BA
NA
NA
^B
\$^

Add 5

BANAN
NANA\$
NA\$^B
^BANA
ANANA
ANA\$^
\$^BAN
A\$^BA

Sort 5

ANANA
ANA\$^
A\$^BA
BANAN
NANA\$
NA\$^B
^BANA
\$^BAN

Add 6

BANANA
NANA\$^
NA\$^BA
^BANAN
ANANA\$
ANA\$^B
\$^BANA
A\$^BAN

Sort 6

ANANA\$
ANA\$^B
A\$^BAN
BANANA
NANA\$^
NA\$^BA
^BANAN
\$^BANA

Add 3

BAN
NAN
NA\$
^BA
ANA
ANA
\$^B
A\$^

Sort 3

ANA
ANA
A\$^
BAN
NAN
NA\$
^BA
\$^B

Add 4

BANA
NANA
NA\$^
^BAN
ANAN
ANA\$
\$^BA
A\$^B

Sort 4

ANAN
ANA\$
A\$^B
BANA
NANA
NA\$^
^BAN
\$^BA

Add 7

BANANA\$
NANA\$^B
NA\$^BAN
^BANANA
ANANA\$^
ANA\$^BA
\$^BANAN
A\$^BANA

Sort 7

ANANA\$^
ANA\$^BA
A\$^BANA
BANANA\$
NANA\$^B
NA\$^BAN
^BANANA
\$^BANAN

Add 8

BANANA\$^
NANA\$^BA
NA\$^BANA
^BANANA\$
ANANA\$^B
ANA\$^BAN
\$^BANANA
A\$^BANAN

Sort 8

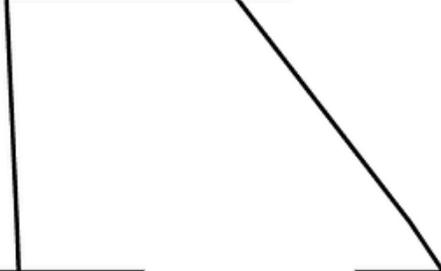
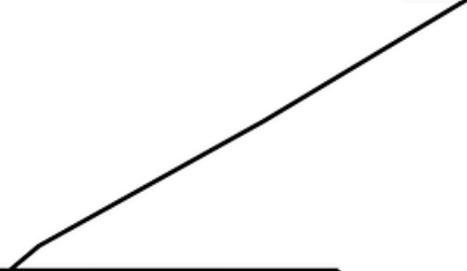
ANANA\$^B
ANA\$^BAN
A\$^BANAN
BANANA\$^
NANA\$^BA
NA\$^BANA
^BANANA\$
\$^BANANA

Output

^BANANA\$

Benefits of B-Tree

- ◆ **Self-Balancing:** B-trees ensure that the 'height' of the tree stays balanced even when inserting or deleting data. This ensures logarithmic time complexity for insertion, deletion, and searching.
- ◆ **Ordered:** B-trees keep the data sorted, making range queries ("find all orders between date X and Y") and inequality comparisons very fast.
- ◆ **Disk-Friendly:** B-trees are designed to work well with disk-based storage. A single node of a B-tree often corresponds to a disk block, minimizing disk access operations.

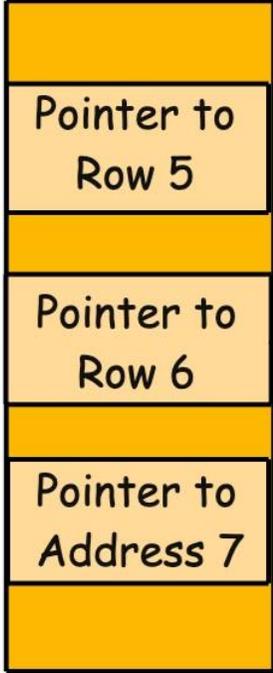


blog.algomaster.io

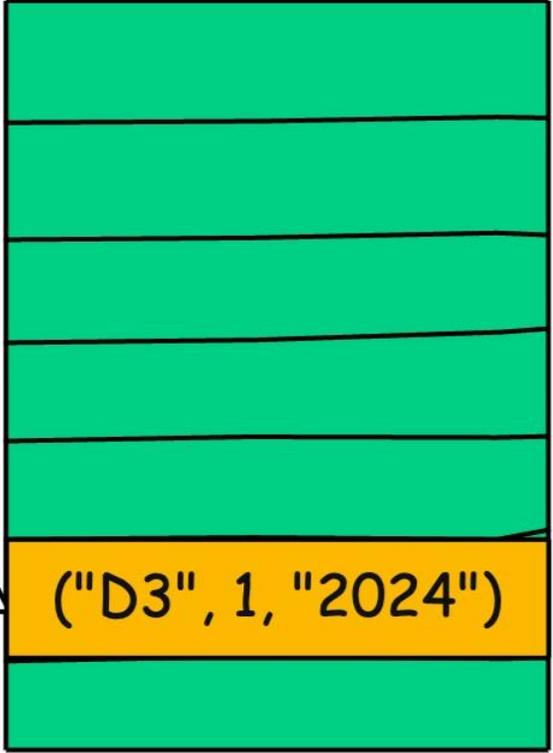
D3
Search Key

Hash Function

$f(\text{"D3"})=6$



Buckets



Data

Row 1
Row 2
Row 3
Row 4
Row 5
Row 6
Row 7

Bitmaps

A	0	1	1	0	1
B	1	0	0	1	1
C	0	1	0	1	0
D	0	0	1	0	0
E	1	0	1	0	0

How to use Database Indexes Smartly?

To get the most out of database indexes, consider these best practices:

- ◆ **Identify Query Patterns:** Analyze the most frequent and critical queries executed against your database to determine which columns to index and which type of index to use.
- ◆ **Index Frequently Used Columns:** Consider indexing columns that are frequently used in WHERE, JOIN, and ORDER BY clauses.
- ◆ **Index Selective Columns:** Indexes are most effective on columns with a good spread of data values (high cardinality). Indexing a gender column might be less beneficial than one with a unique customer_id.
- ◆ **Use Appropriate Index Types:** Choose the right index type for your data and queries.
- ◆ **Consider Composite Indexes:** For queries involving multiple columns, consider creating composite indexes that encompass all relevant columns. This reduces the need for multiple single-column indexes and improves query performance.
- ◆ **Monitor Index Performance:** Regularly monitor index performance, remove unused indexes and adjust your indexing strategy as the database workload evolves.
- ◆ **Avoid Over-Indexing:** Avoid creating too many indexes, as this can lead to increased storage requirements and slower write performance.
 - ◆ Indexes take up extra disk space since they're additional data structures that need to be stored alongside your tables.
 - ◆ Every time you insert, update, or delete data in a table with an index, the index needs to update too. This can slightly slow down write operations.

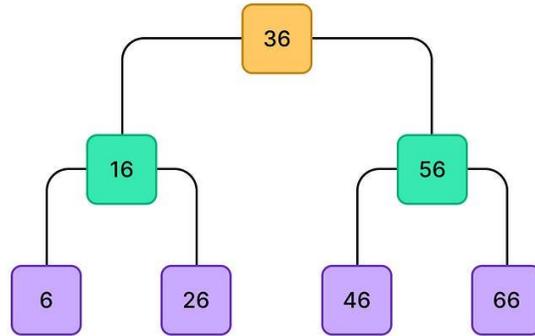
To summarize, indexes are a powerful tool to optimize database query performance.

But remember to choose the right column and index type, monitor performance, and avoid over-indexing to get the most out of them.

Key Data Structures

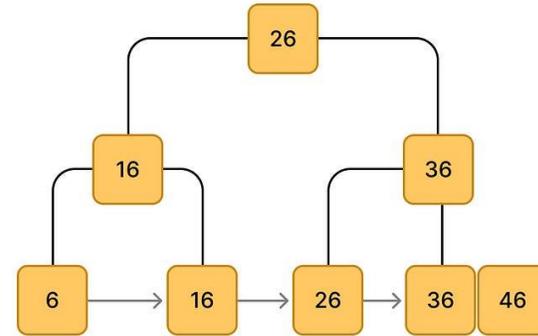
B-Tree Index

Inserting 6, 16, 26, 36, 24, 56, 66

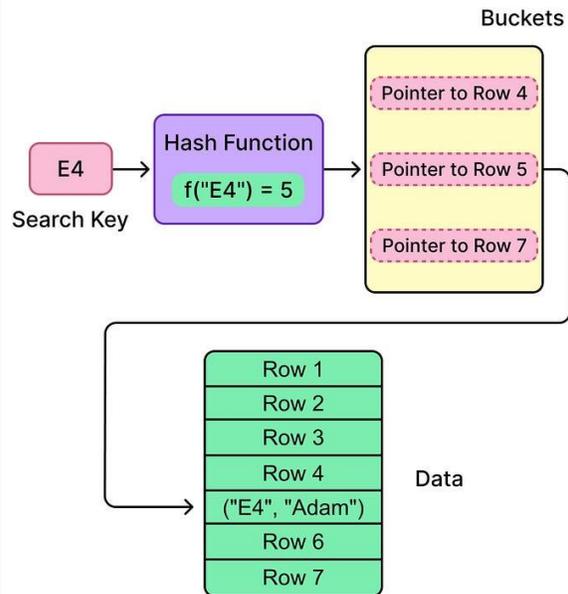


B+ Tree Index

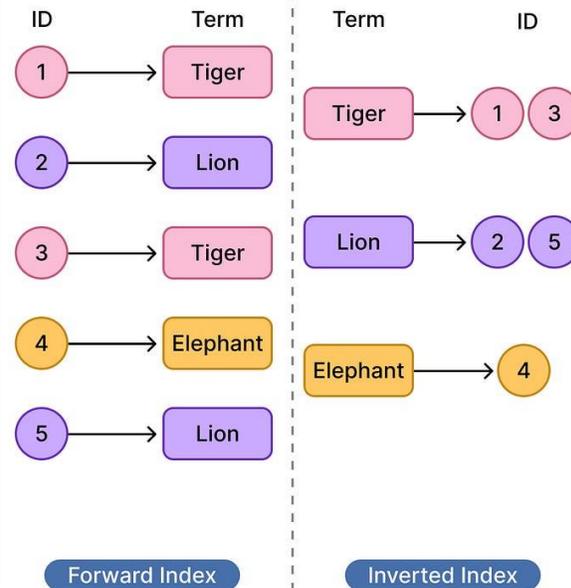
Inserting 6, 16, 26, 36, 46



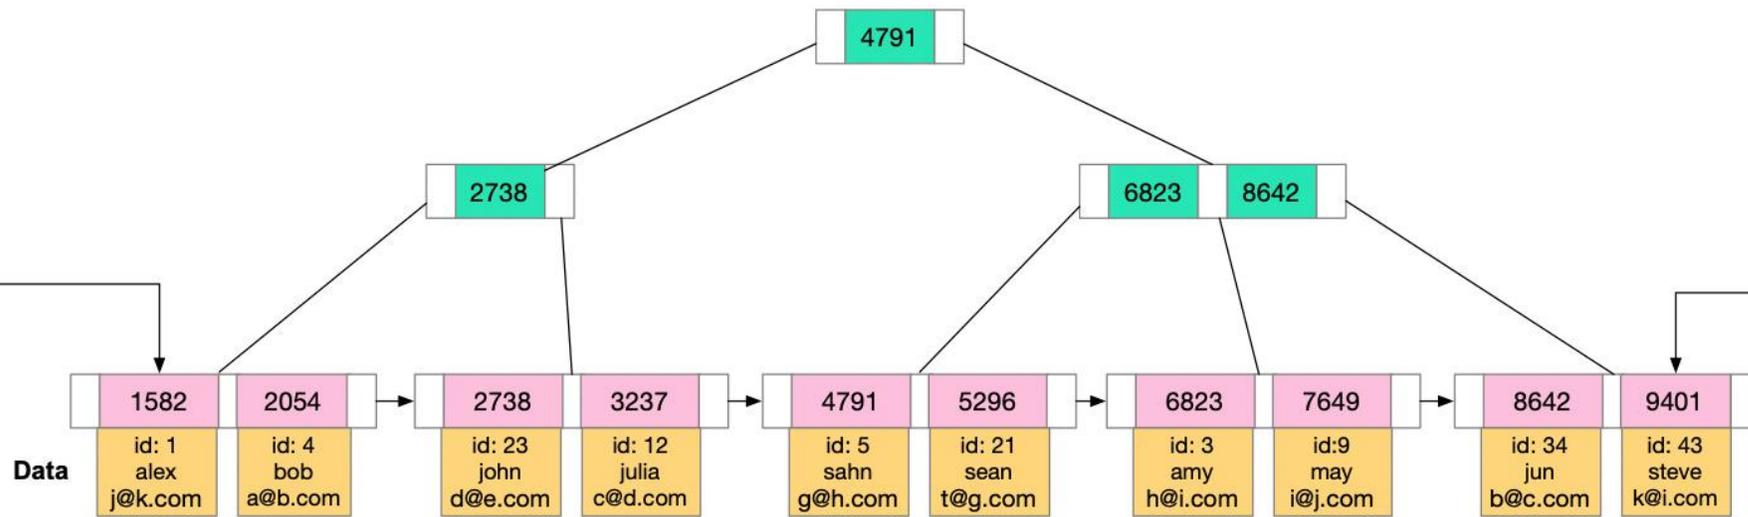
Hash Index



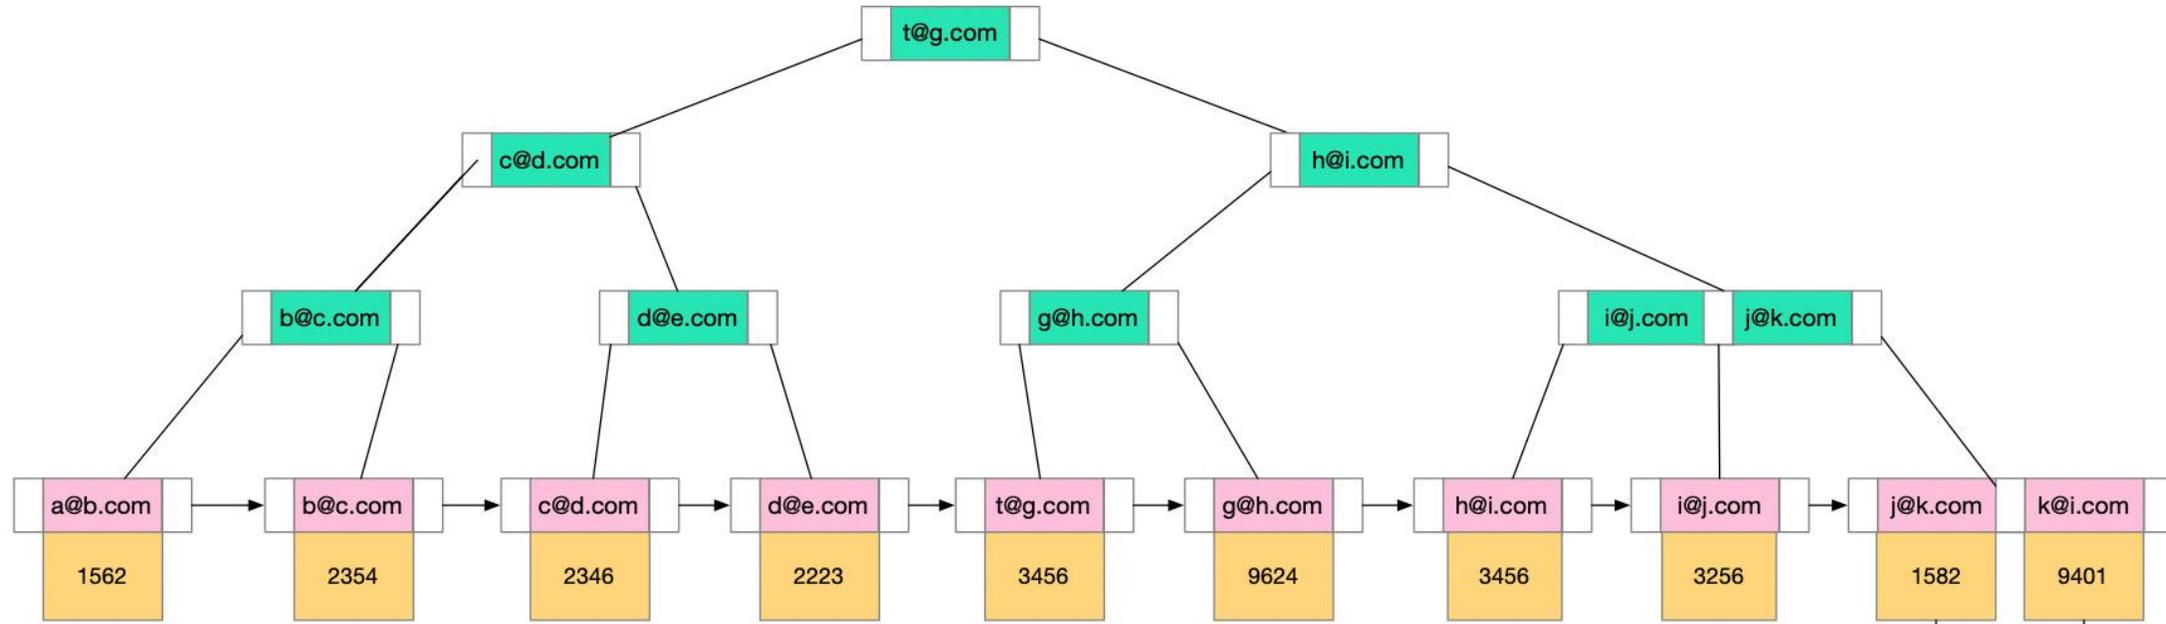
Inverted Index



Clustered index



Secondary Index: email



Databases

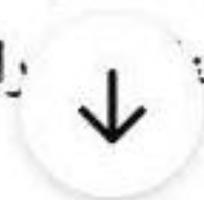
کاربرد	پایگاه داده مناسب	توضیح
سامانه بانکی	PostgreSQL / Oracle SQL	نیاز به تراکنش دقیق و روابط پیچیده
شبکه اجتماعی	MongoDB / Cassandra	داده متنوع و متغیر، مقیاس بالا
اپلیکیشن چت	Redis / DynamoDB	سرعت بالا و real-time
تحلیل کلان داده‌ها	HBase / Cassandra	داده توزیع شده، حجم بسیار زیاد



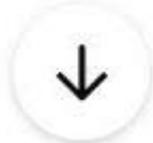
تحلیل کلان داده‌ها HBase / Cassandra داده توزیع شده، حجم بسیار زیاد

فروشگاه آنلاین ترکیبی از MySQL (برای سفارش‌ها) و MongoDB (برای محصولات) سیستم هیبریدی

نوع NoSQL	توضیح	مثال سیستم
Key-Value	ذخیره داده به صورت جفت کلید و مقدار	Redis, DynamoDB
Document-Oriented	ذخیره داده به صورت سند JSON یا BSON	MongoDB, CouchDB
Column-Based	ذخیره داده به صورت ستونها،	Cassandra, HBase
Graph-Based	ذخیره داده به صورت گره و یال، زای شبکه‌ها	Neo4j, Amazon Neptune



مورد استفاده	مناسب تر SQL	مناسب تر NoSQL
سیستم‌های مالی و بانکی	✓	✗
شبکه‌های اجتماعی با داده متغیر	✗	✓
گزارش‌گیری پیچیده (Analytics)	✓	✗
سیستم‌های real-time و مقیاس بزرگ	✗	✓
داده‌های ساختارمند و پایدار	✓	✗
داده‌های نیمه ساخت یافته (JSON, Logs, Sensors)	✗	✓





Bill Gates

Greater Seattle Area | Philanthropy

Summary

Co-chair of the Bill & Melinda Gates Foundation. Chairman, Microsoft Corporation. Voracious reader. Avid traveler. Active blogger.

Experience

Co-chair - Bill & Melinda Gates Foundation
2000 - Present

Co-founder, Chairman - Microsoft
1975 - Present

Education

Harvard University
1973 - 1975

Lakeside School, Seattle

Contact Info

Blog: thegatesnotes.com
Twitter: @BillGates

users table

user_id	first_name	last_name	summary
251	Bill	Gates	Co-chair of ... blogger.
	region_id	industry_id	photo_id
	us:91	131	57817532

regions table

id	region_name
us:7	Greater Boston Area
us:91	Greater Seattle Area

industries table

id	industry_name
43	Financial Services
48	Construction
131	Philanthropy

positions table

id	user_id	job_title	organization
458	251	Co-chair	Bill & Melinda Gates F...
457	251	Co-founder, Chairman	Microsoft

education table

id	user_id	school_name	start	end
807	251	Harvard University	1973	1975
806	251	Lakeside School, Seattle	NULL	NULL

contact_info table

id	user_id	type	url
155	251	blog	http://thegatesnotes.com
156	251	twitter	http://twitter.com/BillGates

Representing a LinkedIn profile as a JSON document

```
{
  "user_id": 251,
  "first_name": "Bill",
  "last_name": "Gates",
  "summary": "Co-chair of the Bill & Melinda Gates... Active blogger.",
  "region_id": "us:91",
  "industry_id": 131,
  "photo_url": "/p/7/000/253/05b/308dd6e.jpg",

  "positions": [
    {"job_title": "Co-chair", "organization": "Bill & Melinda Gates Foundation"},
    {"job_title": "Co-founder, Chairman", "organization": "Microsoft"}
  ],
  "education": [
    {"school_name": "Harvard University", "start": 1973, "end": 1975},
    {"school_name": "Lakeside School, Seattle", "start": null, "end": null}
  ],
  "contact_info": {
    "blog": "http://thegatesnotes.com",
    "twitter": "http://twitter.com/BillGates"
  }
}
```

<http://www.linkedin.com/in/williamhgates>



Bill Gates

Greater Seattle Area | Philanthropy

Summary

Co-chair of the Bill & Melinda Gates Foundation. Chairman, Microsoft Corporation. Voracious reader. Avid traveler. Active blogger.

Experience

Co-chair - Bill & Melinda Gates Foundation
2000 - Present

Co-founder, Chairman - Microsoft
1975 - Present

Education

Harvard University
1973 - 1975

Lakeside School, Seattle

Contact Info

Blog: thegatesnotes.com
Twitter: @BillGates

user_id	first_name	last_name	summary
251	Bill	Gates	Co-chair of ... blogger.

region_id	industry_id	photo_id
us:91	131	57817532

id	region_name
us:7	Greater Boston Area
us:91	Greater Seattle Area

id	industry_name
43	Financial Services
48	Construction
131	Philanthropy

id	user_id	job_title	organization
458	251	Co-chair	Bill & Melinda Gates F...
457	251	Co-founder, Chairman	Microsoft

id	user_id	school_name	start	end
807	251	Harvard University	1973	1975
806	251	Lakeside School, Seattle	NULL	NULL

id	user_id	type	url
155	251	blog	http://thegatesnotes.com
156	251	twitter	http://twitter.com/BillGates



Bill Gates

Greater Seattle Area | Philanthropy

Summary

Co-chair of the Bill & Melinda Gates Foundation. Chairman, Microsoft Corporation. Voracious reader. Avid traveler. Active blogger.

Experience

Co-chair • Bill & Melinda Gates Foundation
2000 – Present

Co-founder, Chairman • Microsoft
1975 – Present

Education

Harvard University
1973 – 1975

Lakeside School, Seattle

Contact Info

Blog: thegatesnotes.com
Twitter: @BillGates

user_id	first_name	last_name	summary
251	Bill	Gates	Co-chair of ... blogger.

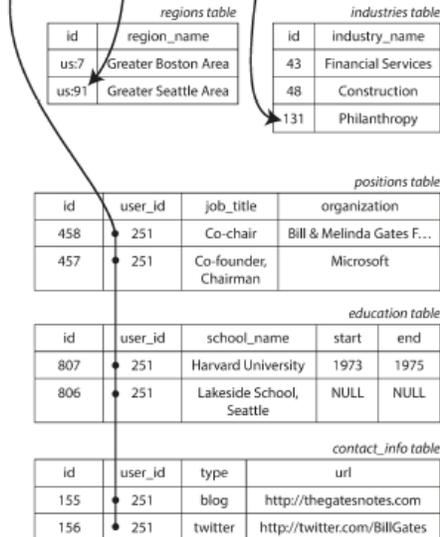
id	region_name
us:7	Greater Boston Area
us:91	Greater Seattle Area

id	industry_name
43	Financial Services
48	Construction
131	Philanthropy

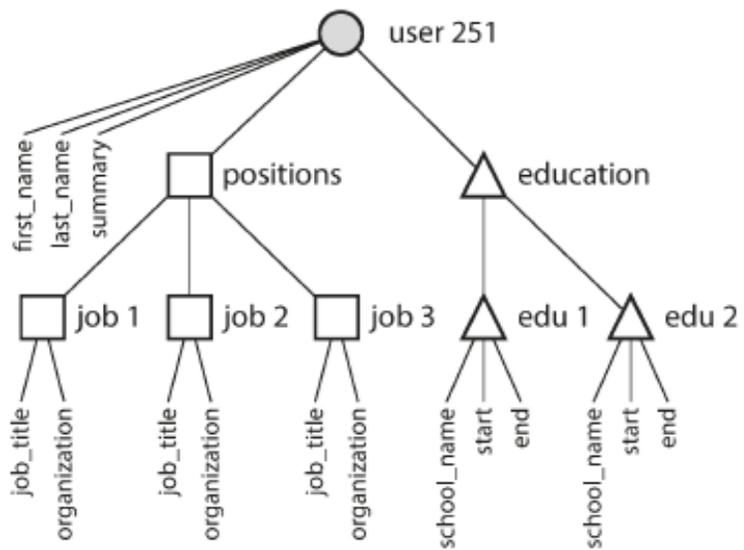
id	user_id	job_title	organization
458	251	Co-chair	Bill & Melinda Gates F...
457	251	Co-founder, Chairman	Microsoft

id	user_id	school_name	start	end
807	251	Harvard University	1973	1975
806	251	Lakeside School, Seattle	NULL	NULL

id	user_id	type	url
155	251	blog	http://thegatesnotes.com
156	251	twitter	http://twitter.com/BillGates



- Consistent style and spelling across profiles
- Avoiding ambiguity (e.g., if there are several cities with the same name)
- Ease of updating—the name is stored in only one place, so it is easy to update across the board if it ever needs to be changed (e.g., change of a city name due to political events)
- Localization support—when the site is translated into other languages, the standardized lists can be localized, so the region and industry can be displayed in the viewer’s language
- Better search—e.g., a search for philanthropists in the state of Washington can match this profile, because the list of regions can encode the fact that Seattle is in Washington (which is not apparent from the string "Greater Seattle Area")



One-to-many relationships forming a tree structure.

<http://www.linkedin.com/in/williamhgates>



Bill Gates

Greater Seattle Area | Philanthropy

Summary

Co-chair of the Bill & Melinda Gates Foundation. Chairman, Microsoft Corporation. Voracious reader. Avid traveler. Active blogger.

Experience

Co-chair • Bill & Melinda Gates Foundation
2000 – Present

Co-founder, Chairman • Microsoft
1975 – Present

Education

Harvard University
1973 – 1975

Lakeside School, Seattle

Contact Info

Blog: thegatesnotes.com
Twitter: @BillGates

user_id	first_name	last_name	summary
251	Bill	Gates	Co-chair of ... blogger.
	region_id	industry_id	photo_id
	us:91	131	57817532

id	region_name
us:7	Greater Boston Area
us:91	Greater Seattle Area

id	industry_name
43	Financial Services
48	Construction
131	Philanthropy

id	user_id	job_title	organization
458	251	Co-chair	Bill & Melinda Gates F...
457	251	Co-founder, Chairman	Microsoft

id	user_id	school_name	start	end
807	251	Harvard University	1973	1975
806	251	Lakeside School, Seattle	NULL	NULL

id	user_id	type	url
155	251	blog	http://thegatesnotes.com
156	251	twitter	http://twitter.com/BillGates



Experience

Co-chair

Bill & Melinda Gates Foundation

2000 – Present (13 years)

Co-founder, Chairman

Microsoft

1975 – Present



Education

Harvard University

1973 – 1975

Microsoft

Come as you are. Do what you love. At Microsoft we help people and businesses throughout the world realize their full potential. We make this simple mission come to life every day through our ... [More »](#)

Co. Size: 10,001+ employees

Website: <http://www.microsoft.com/>

HQ: Greater Seattle Area

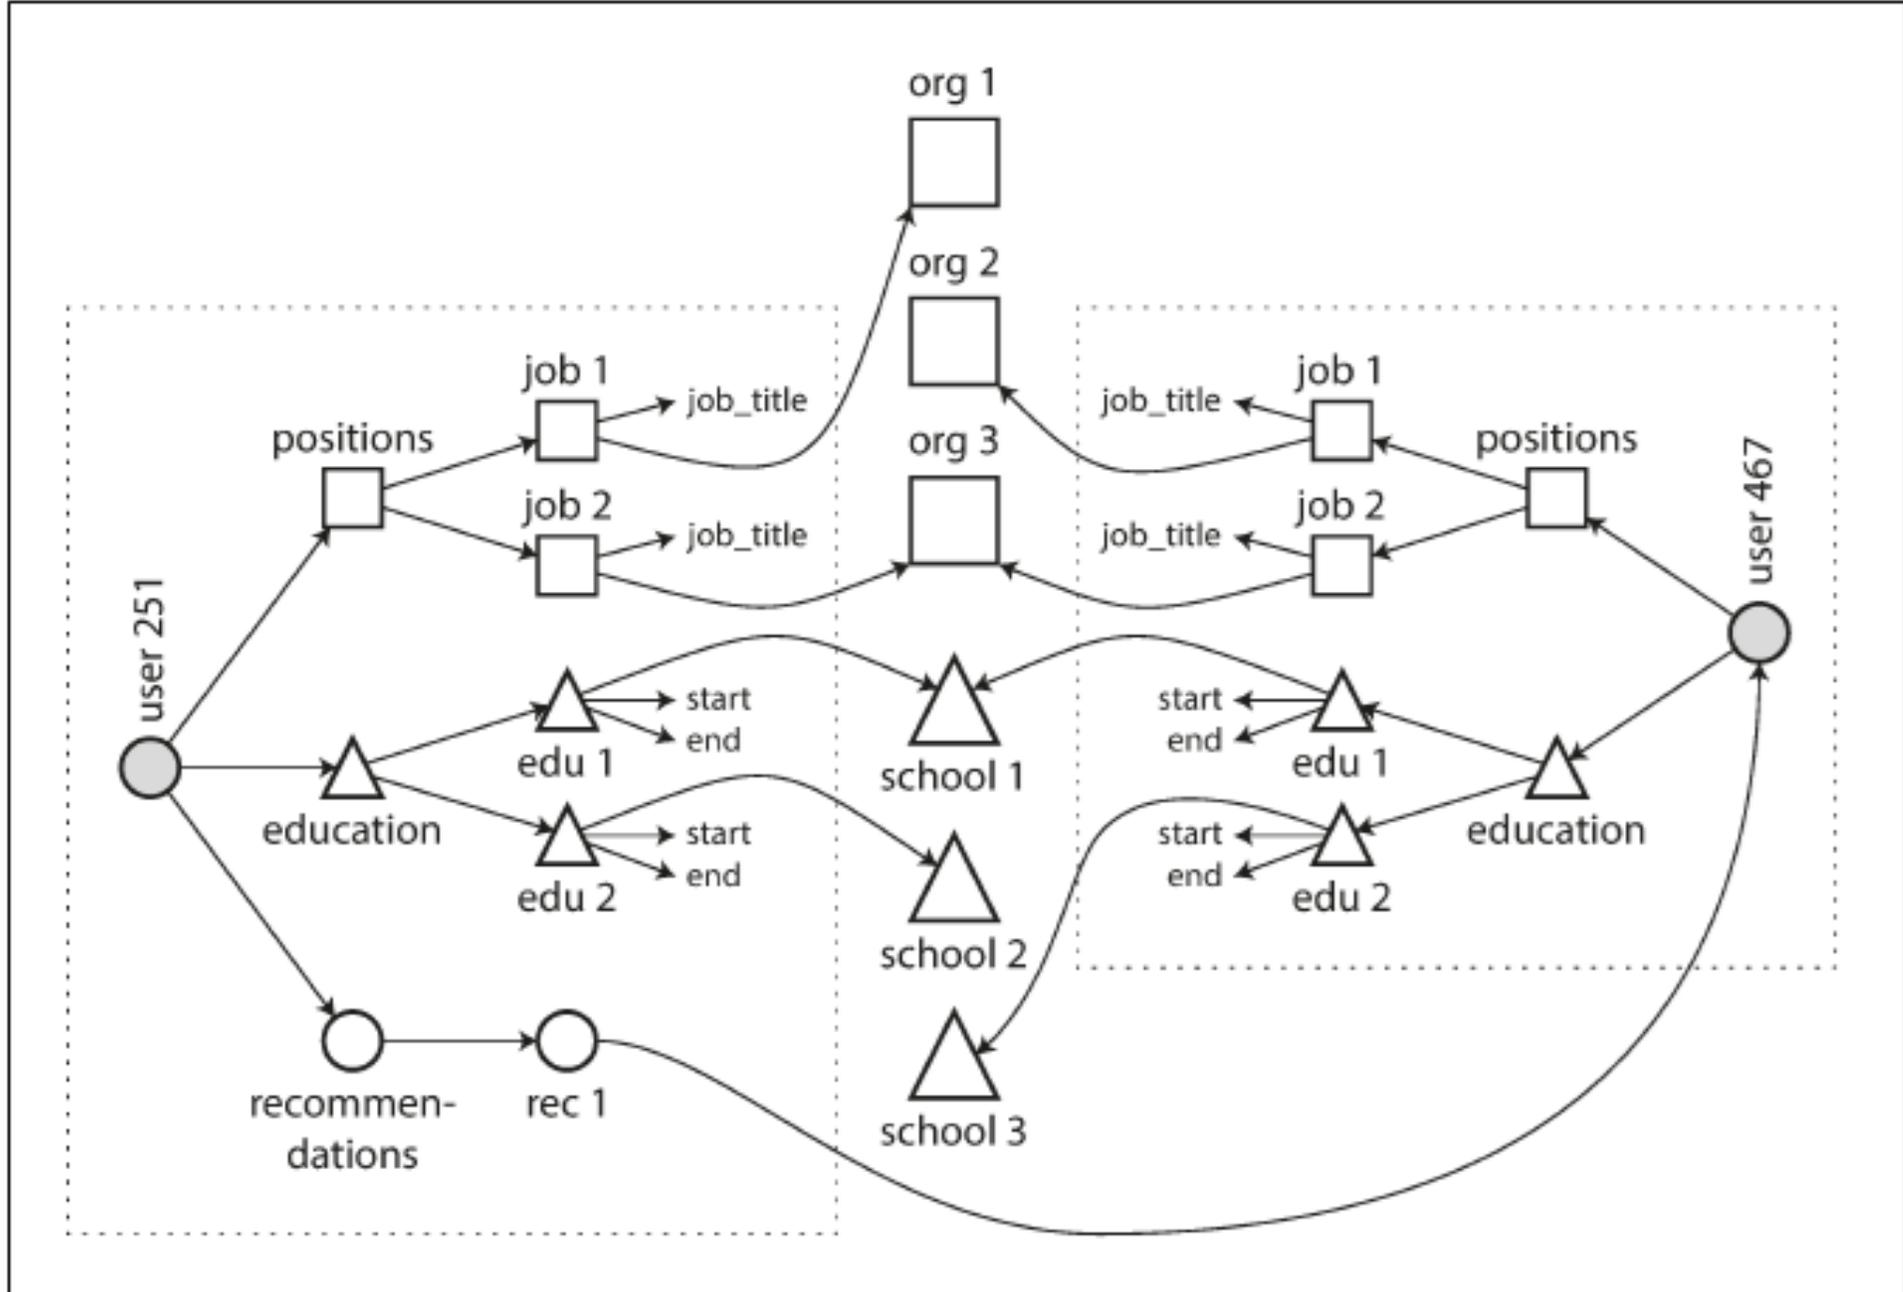
Industry: Computer Software

[Follow company](#) | [Careers](#)

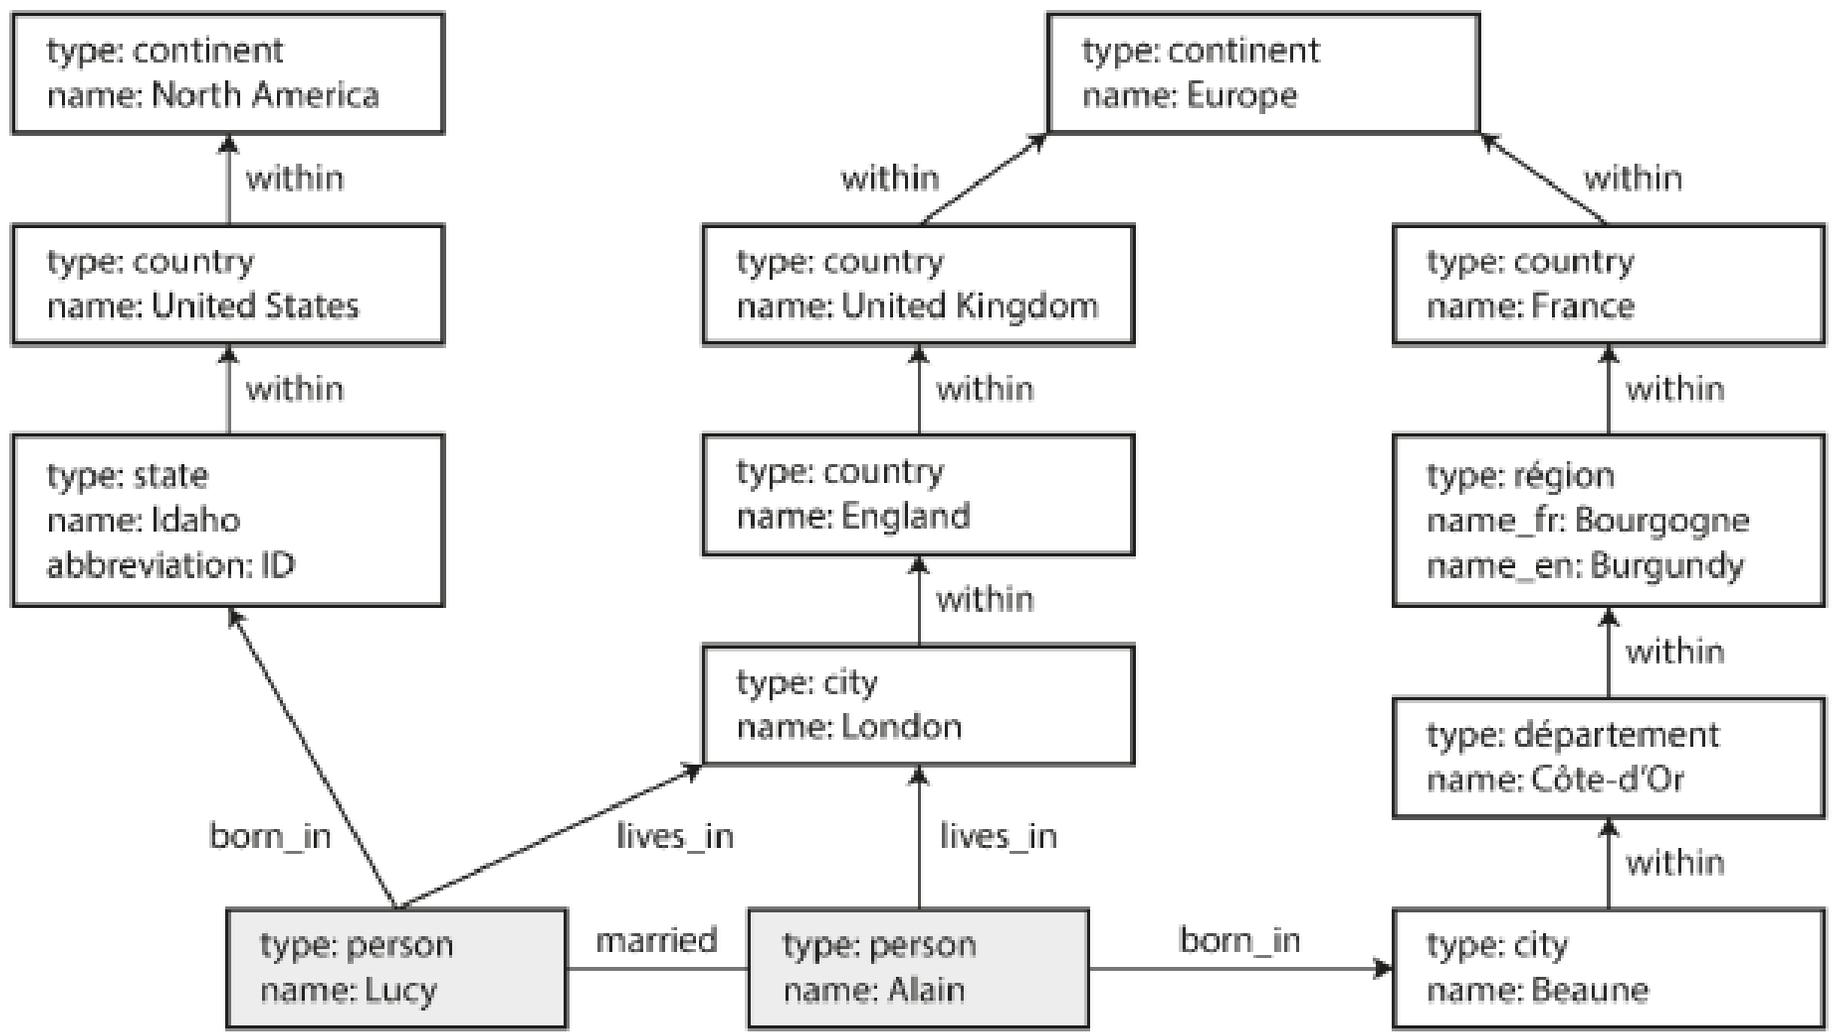


The company name is not just a string, but a link to a company entity.

Screenshot of linkedin.com.



Extending résumés with many-to-many relationships.



Example of graph-structured data (boxes represent vertices, arrows represent edges).

Property Graphs

In the property graph model, each vertex consists of:

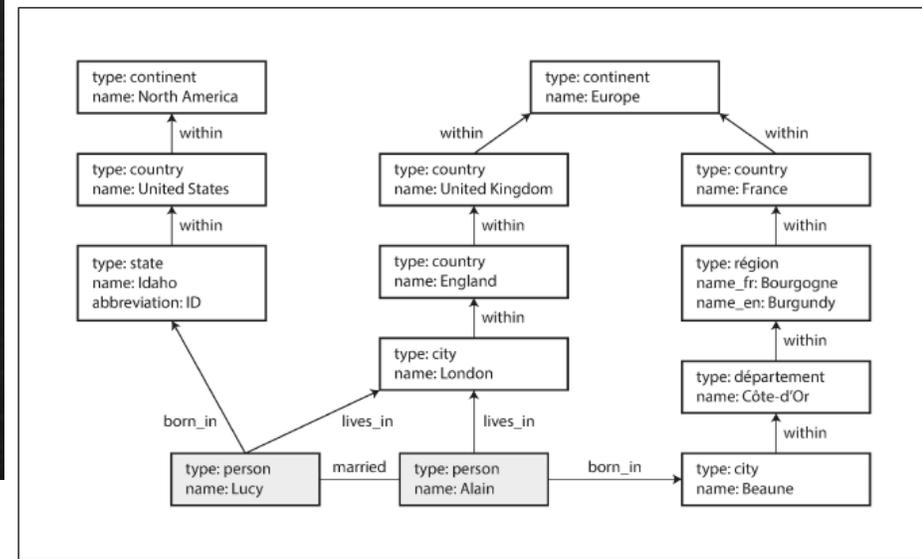
- A unique identifier
- A set of outgoing edges
- A set of incoming edges
- A collection of properties (key-value pairs)

Each edge consists of:

- A unique identifier
- The vertex at which the edge starts (the *tail vertex*)

Representing a property graph using a relational schema

```
CREATE TABLE vertices (  
  vertex_id integer PRIMARY KEY,  
  properties json  
);  
  
CREATE TABLE edges (  
  edge_id integer PRIMARY KEY,  
  tail_vertex integer REFERENCES vertices (vertex_id),  
  head_vertex integer REFERENCES vertices (vertex_id),  
  label text,  
  properties json  
);  
  
CREATE INDEX edges_tails ON edges (tail_vertex);  
CREATE INDEX edges_heads ON edges (head_vertex);
```

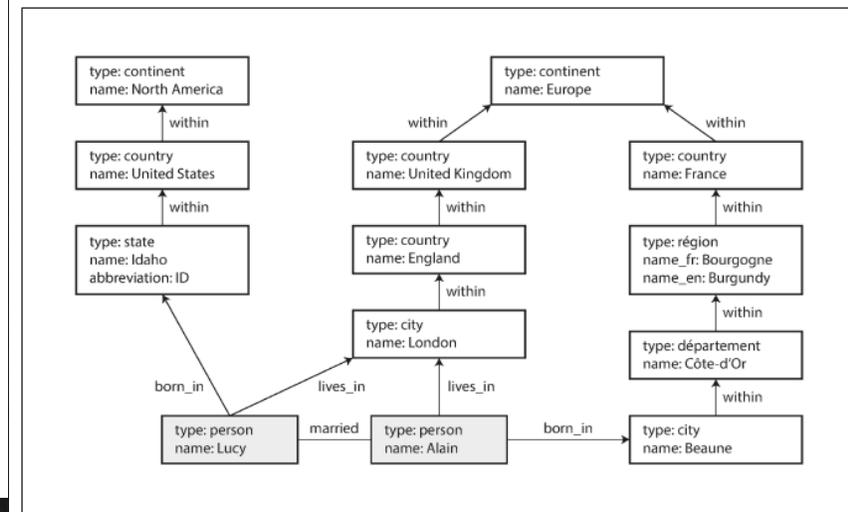


Example of graph-structured data (boxes represent vertices, arrows represent edges).

A subset of the data in Figure 2-5, represented as a Cypher query

CREATE

```
(NAmerica:Location {name:'North America', type:'continent'}),  
(USA:Location {name:'United States', type:'country' }),  
(Idaho:Location {name:'Idaho', type:'state' }),  
(Lucy:Person {name:'Lucy' }),  
(Idaho) -[:WITHIN]-> (USA) -[:WITHIN]-> (NAmerica),  
(Lucy) -[:BORN_IN]-> (Idaho)
```



Example of graph-structured data (boxes represent vertices, arrows represent edges).

Cypher query to find people who emigrated from the US to Europe

MATCH

```
(person) -[:BORN_IN]-> () -[:WITHIN*0..]-> (us:Location {name:'United States'}),  
(person) -[:LIVES_IN]-> () -[:WITHIN*0..]-> (eu:Location {name:'Europe'})
```

RETURN person.name

• خلاصه:

ویژگی	توضیح
ساختار داده	جفت کلید-مقدار
نوع پایگاه داده	NoSQL
کاربرد اصلی	ذخیره‌سازی سریع داده‌های عظیم
مثال‌ها	Redis، DynamoDB، Riak، LevelDB
کاربرد در Big Data	خروجی MapReduce، کش داده، مدیریت session، و داده‌های real-time

• در حوزه‌ی Big Data چه کاربردی دارد؟

در سیستم‌های کلان‌داده، حجم بسیار بالایی از اطلاعات باید به سرعت خوانده و نوشته شود.

مدل **Key-Value** برای این منظور عالی است، چون:

- جستجو با کلید بسیار سریع است ($O(1)$ یا نزدیک به آن)
- ساختار داده‌ها ساده و قابل توزیع بین چندین سرور است
- برای ذخیره‌سازی افقی (horizontal scaling) بسیار مناسب است

در حوزه **Big Data** (کلان‌داده)، مفهوم **Key-Value**

(کلید-مقدار) یکی از بنیادی‌ترین و ساده‌ترین مدل‌های ذخیره‌سازی داده است. در این مدل، داده‌ها به صورت جفت‌های کلید (**Key**) و مقدار (**Value**) ذخیره و بازیابی می‌شوند.

• تعریف ساده:

هر کلید (**Key**) مانند یک شناسه یکتا است که به یک مقدار (**Value**) اشاره می‌کند.

شما می‌توانید با داشتن کلید، به سرعت مقدار مربوطه را پیدا کنید – بدون نیاز به جستجو در کل پایگاه داده.

MapReduce

- در MapReduce، خروجی هر مرحله معمولاً به صورت جفت‌های key-value است.
مثلاً:

 Copy code

```
"word" → count  
"apple" → 57  
"banana" → 34
```

2. پایگاه داده‌های Key-Value برای مدیریت داده‌های عظیم و سریع:

مثال‌ها:

- ▶ Redis
- ▶ Amazon DynamoDB
- ▶ Cassandra (اما wide-column در اصل) (key-value مبتنی بر)
- ▶ Riak
- ▶ LevelDB

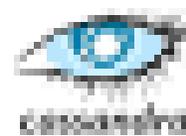
B-Tree vs LSM-Tree

DATABASE STORAGE ENGINES

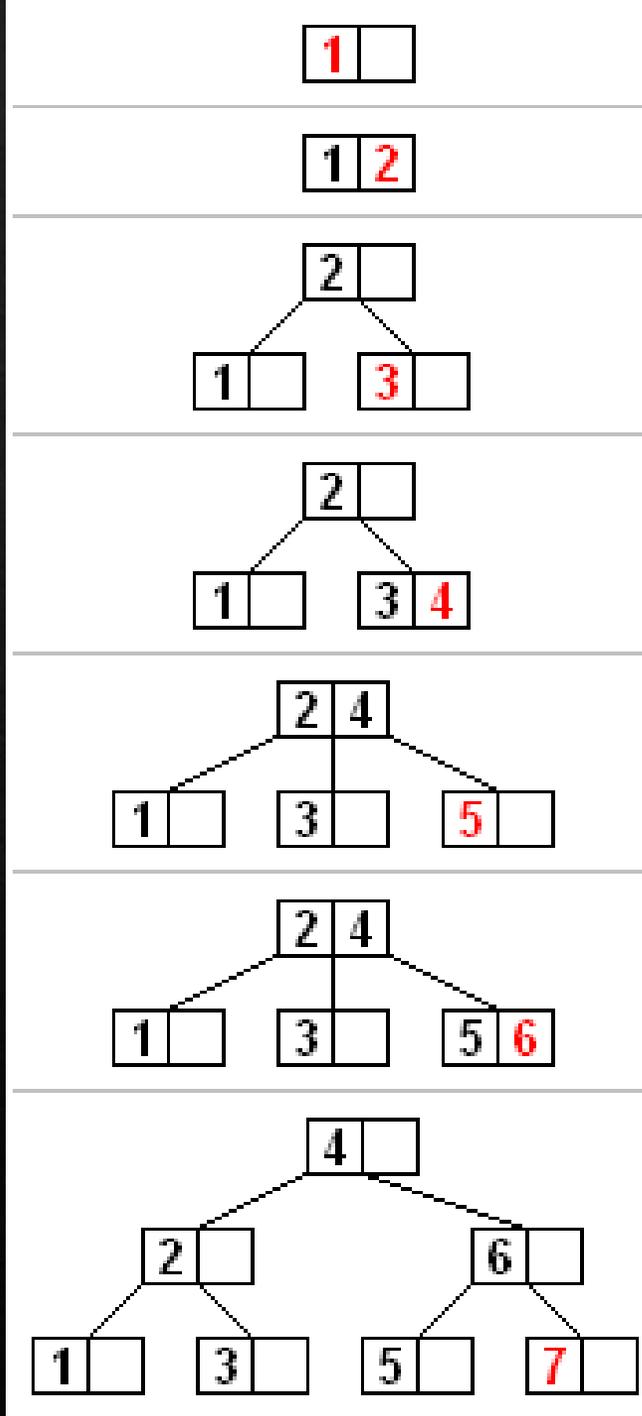
B-TREE



LSM TREE

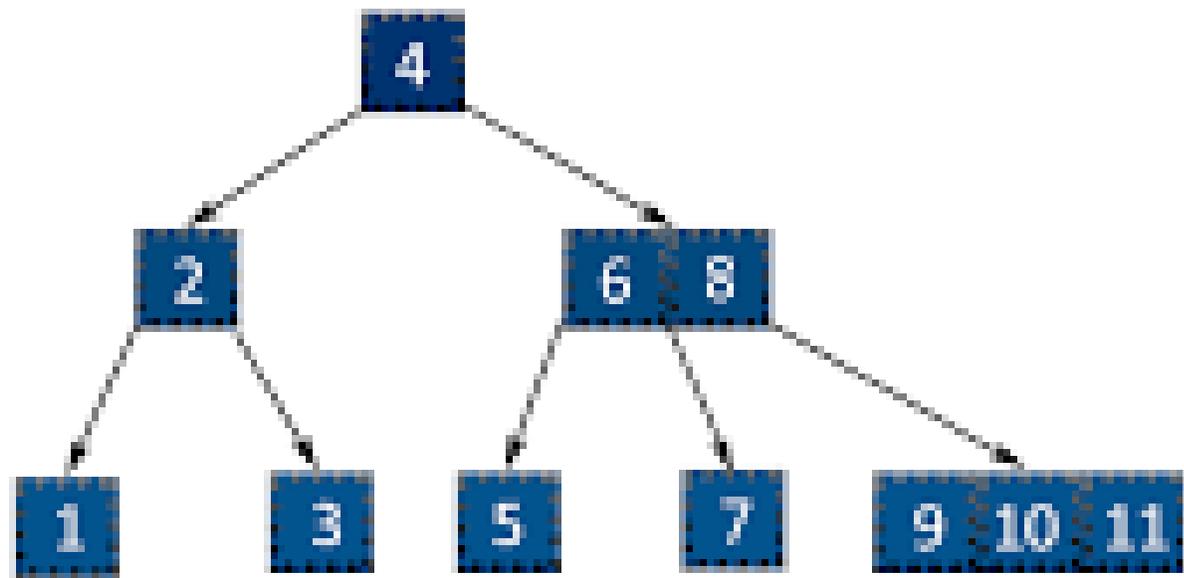


A B-tree insertion example with each iteration. The nodes of this B-tree have at most 3 children (Knuth order 3).

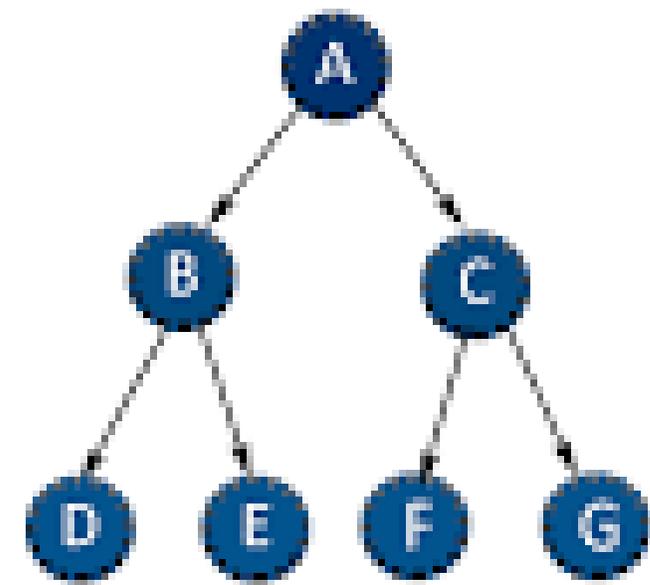


In computer science, a B-tree is a self-balancing tree data structure that maintains sorted data and allows searches, sequential access, insertions, and deletions in logarithmic time. The B-tree generalizes the binary search tree, allowing nodes to have more than two children.

B-TREE



BINARY TREE



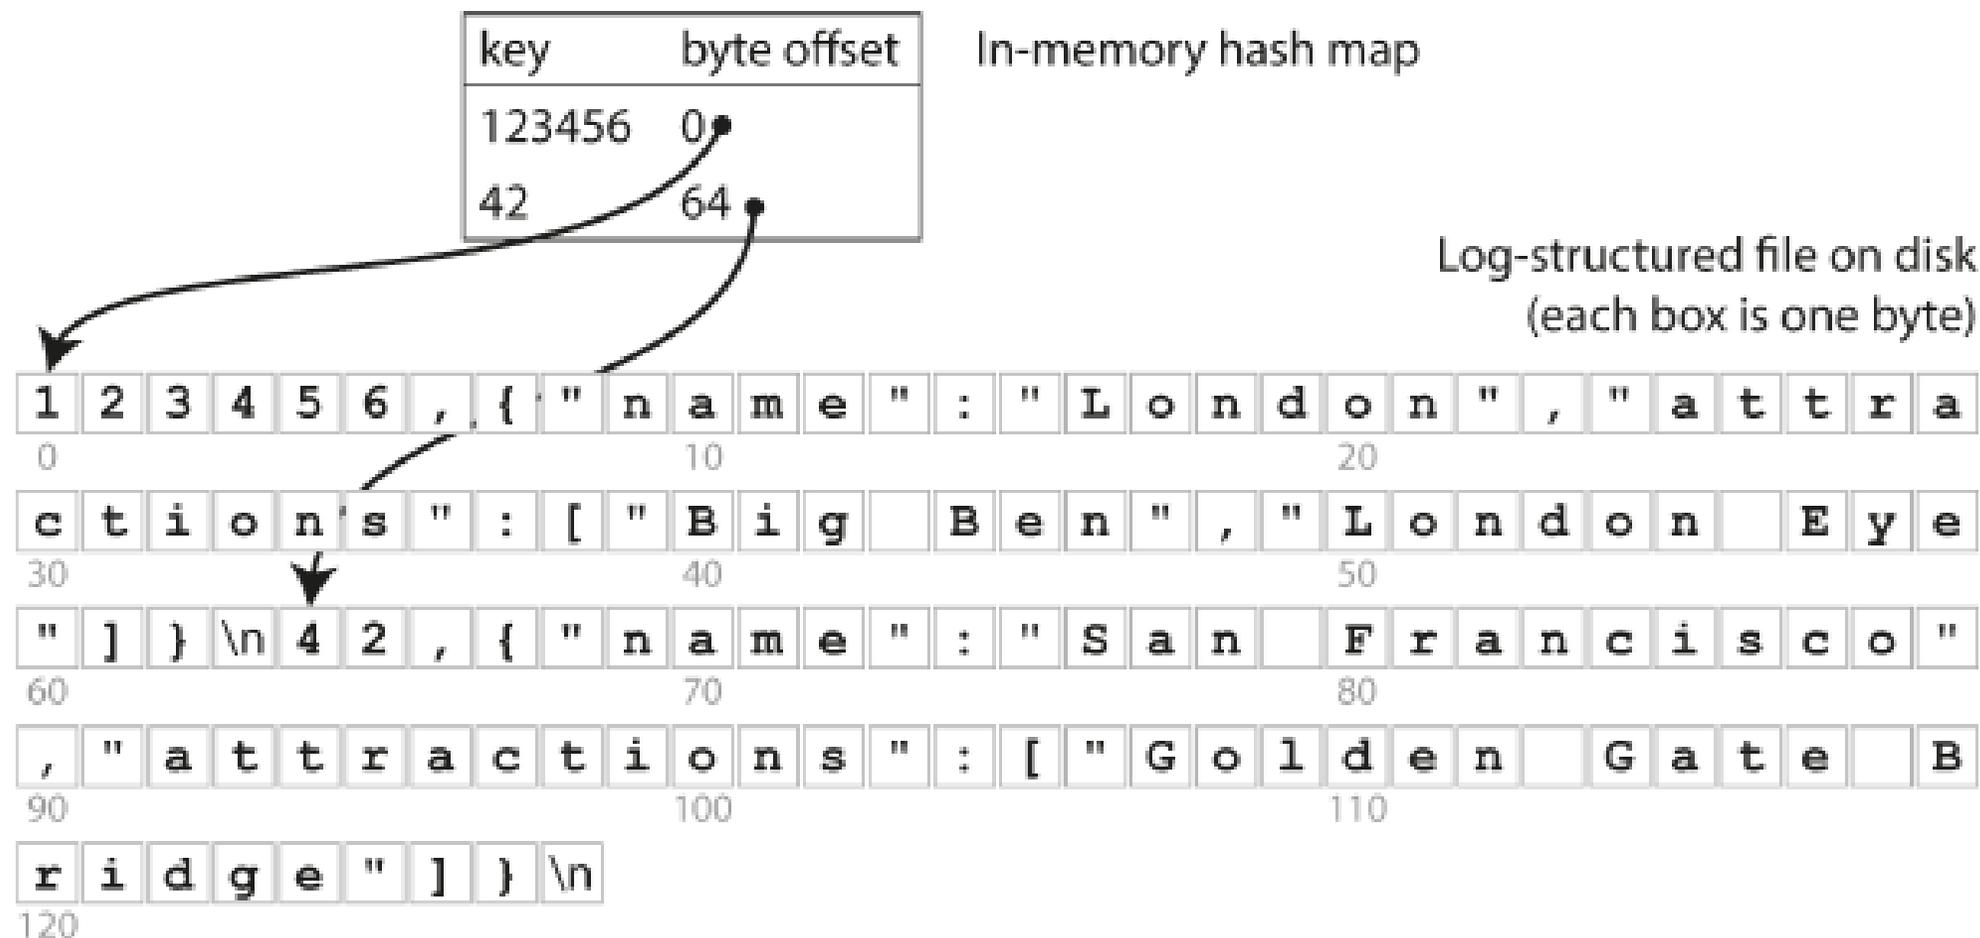


Figure 3-1. Storing a log of key-value pairs in a CSV-like format, indexed with an in-memory hash map.

Data file segment

mew: 1078	purr: 2103	purr: 2104	mew: 1079	mew: 1080	mew: 1081
purr: 2105	purr: 2106	purr: 2107	yawn: 511	purr: 2108	mew: 1082



Compaction process

Compacted segment

yawn: 511	mew: 1082	purr: 2108
-----------	-----------	------------

Figure 3-2. Compaction of a key-value update log (counting the number of times each cat video was played), retaining only the most recent value for each key.

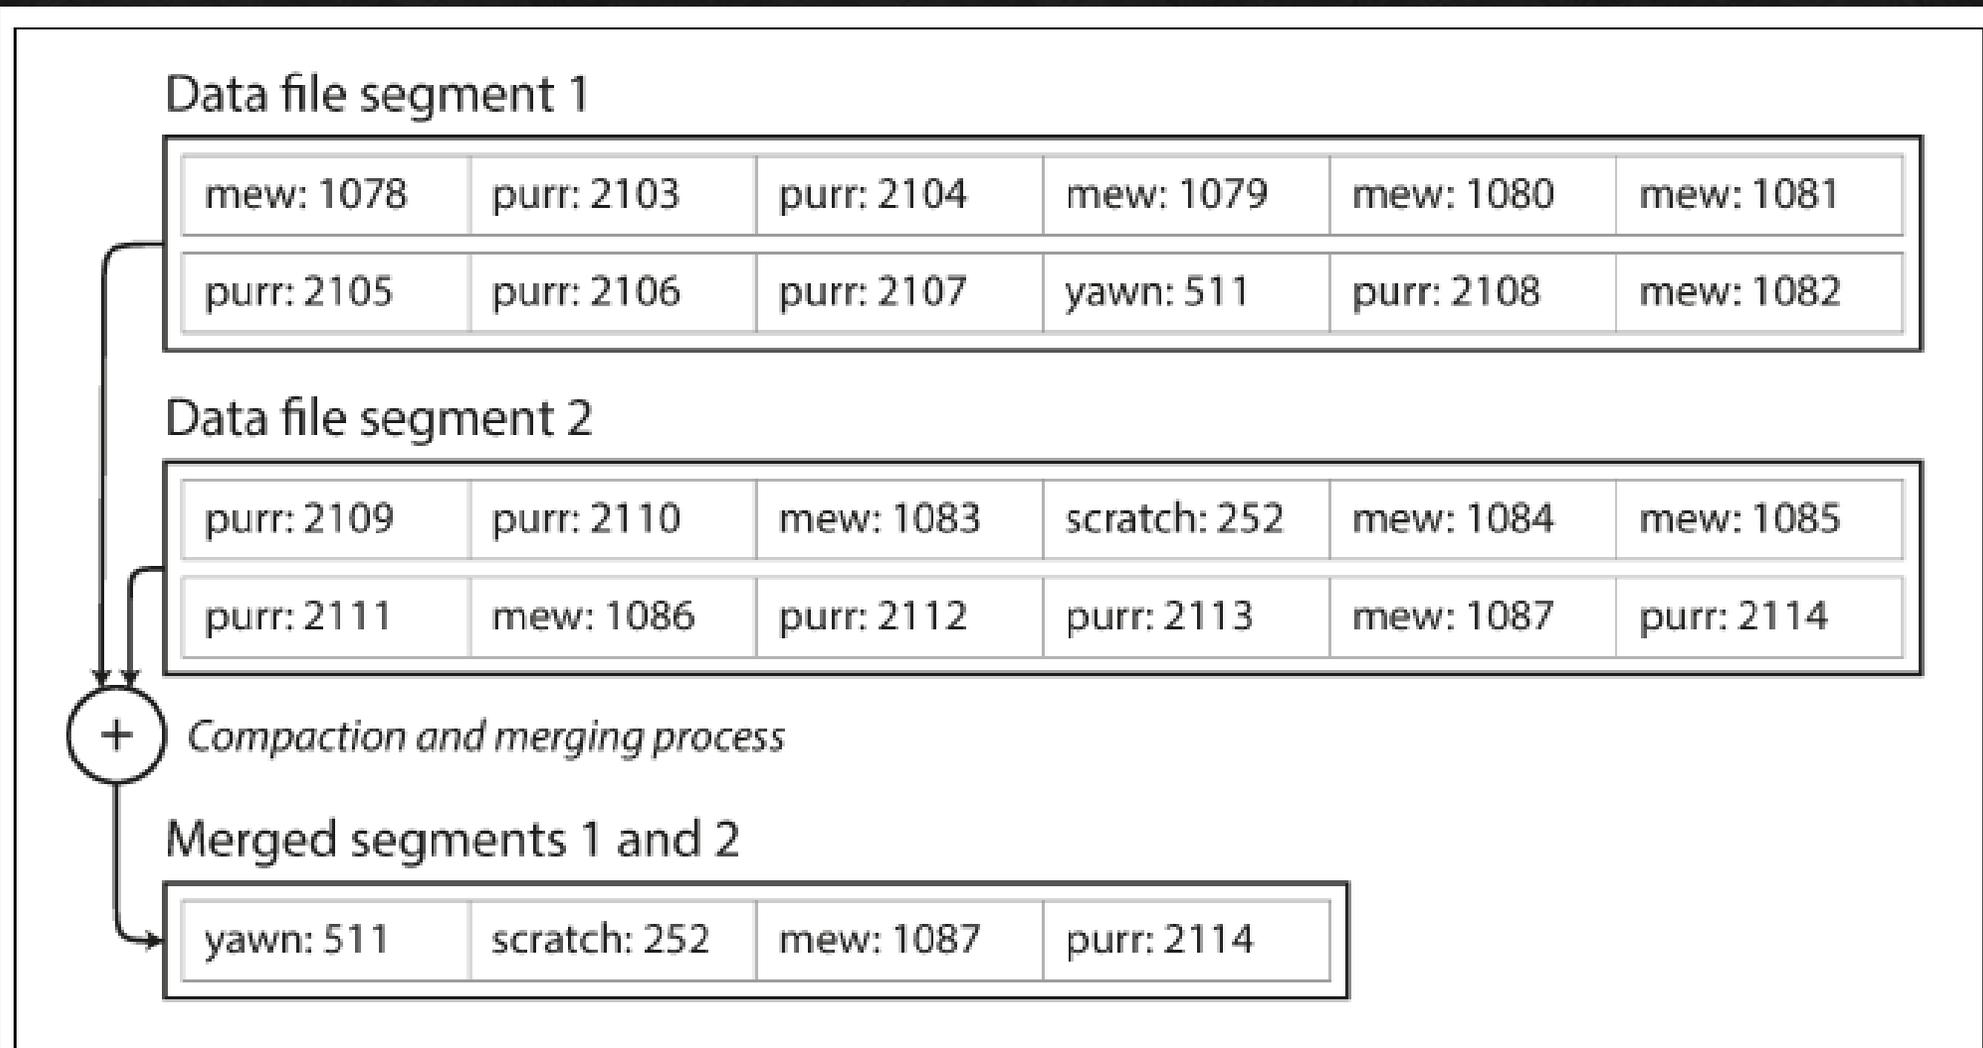


Figure 3-3. Performing compaction and segment merging simultaneously.

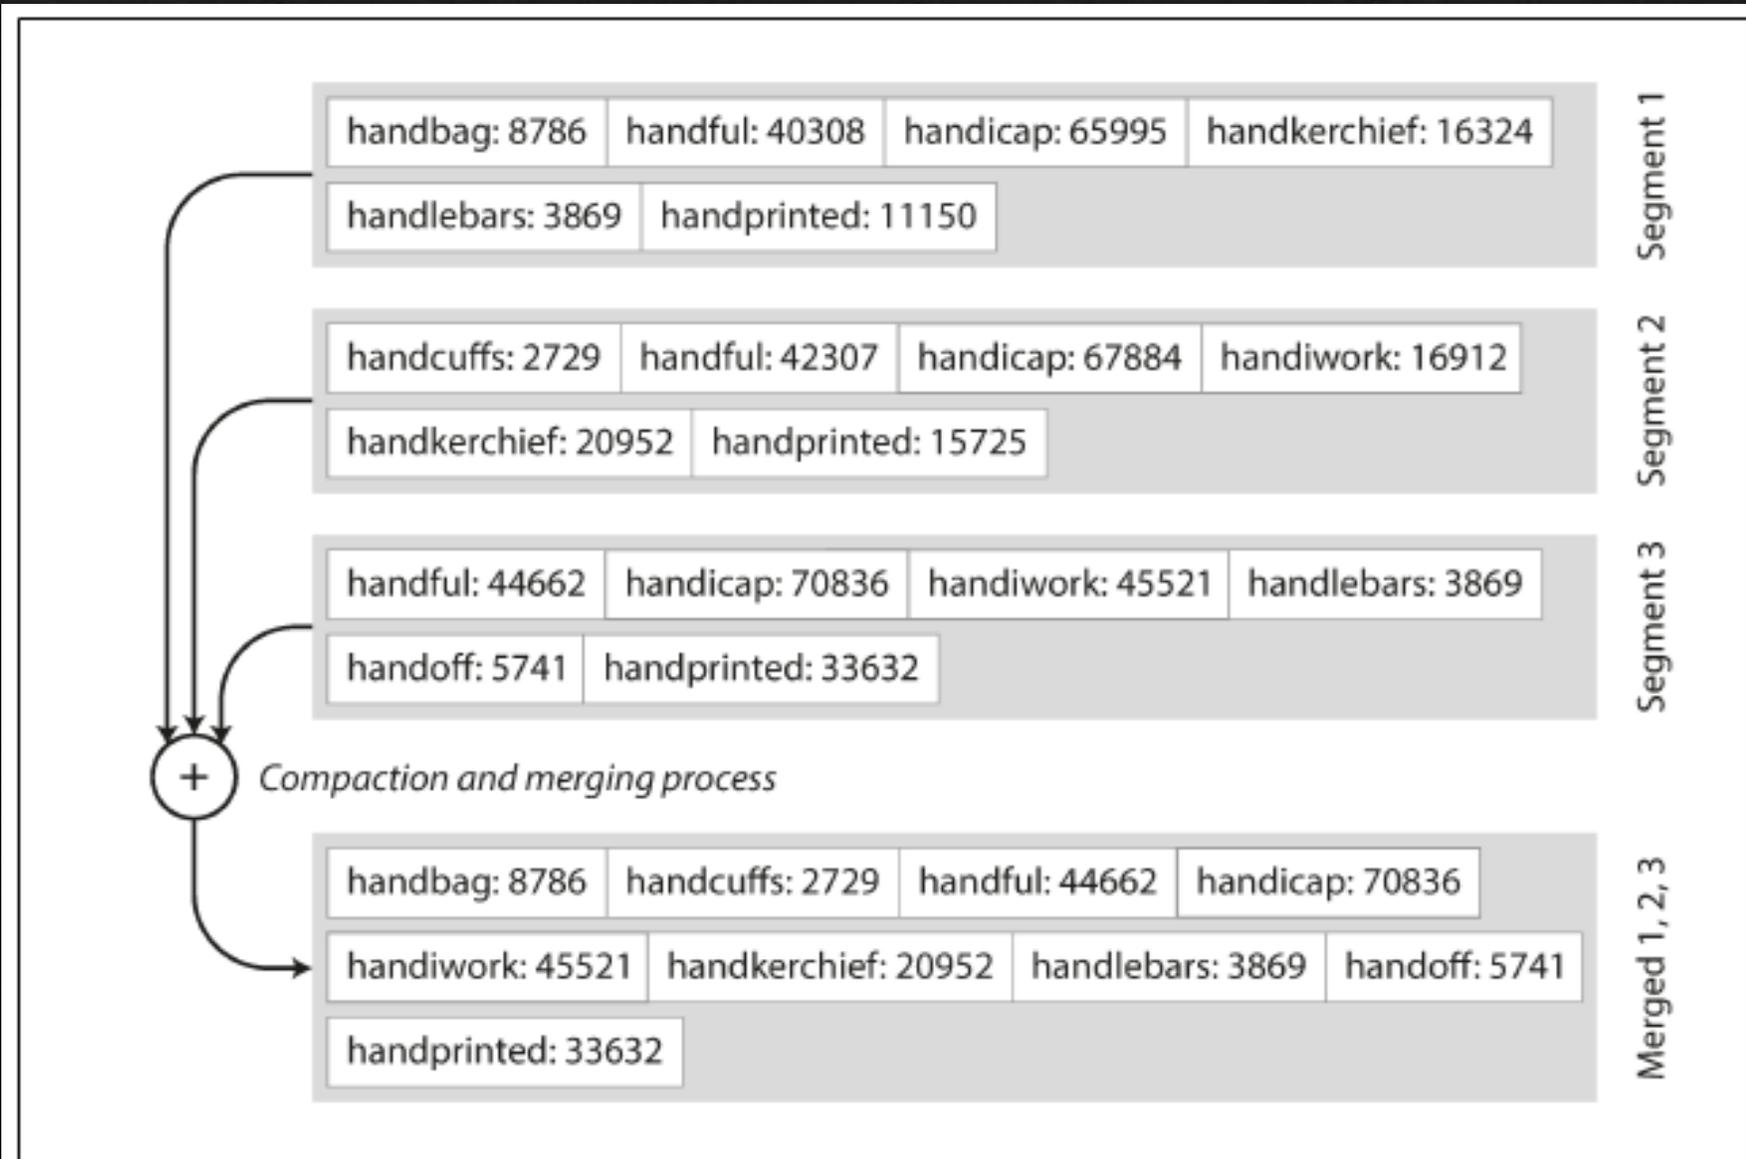


Figure 3-4. Merging several SSTable segments, retaining only the most recent value for each key.

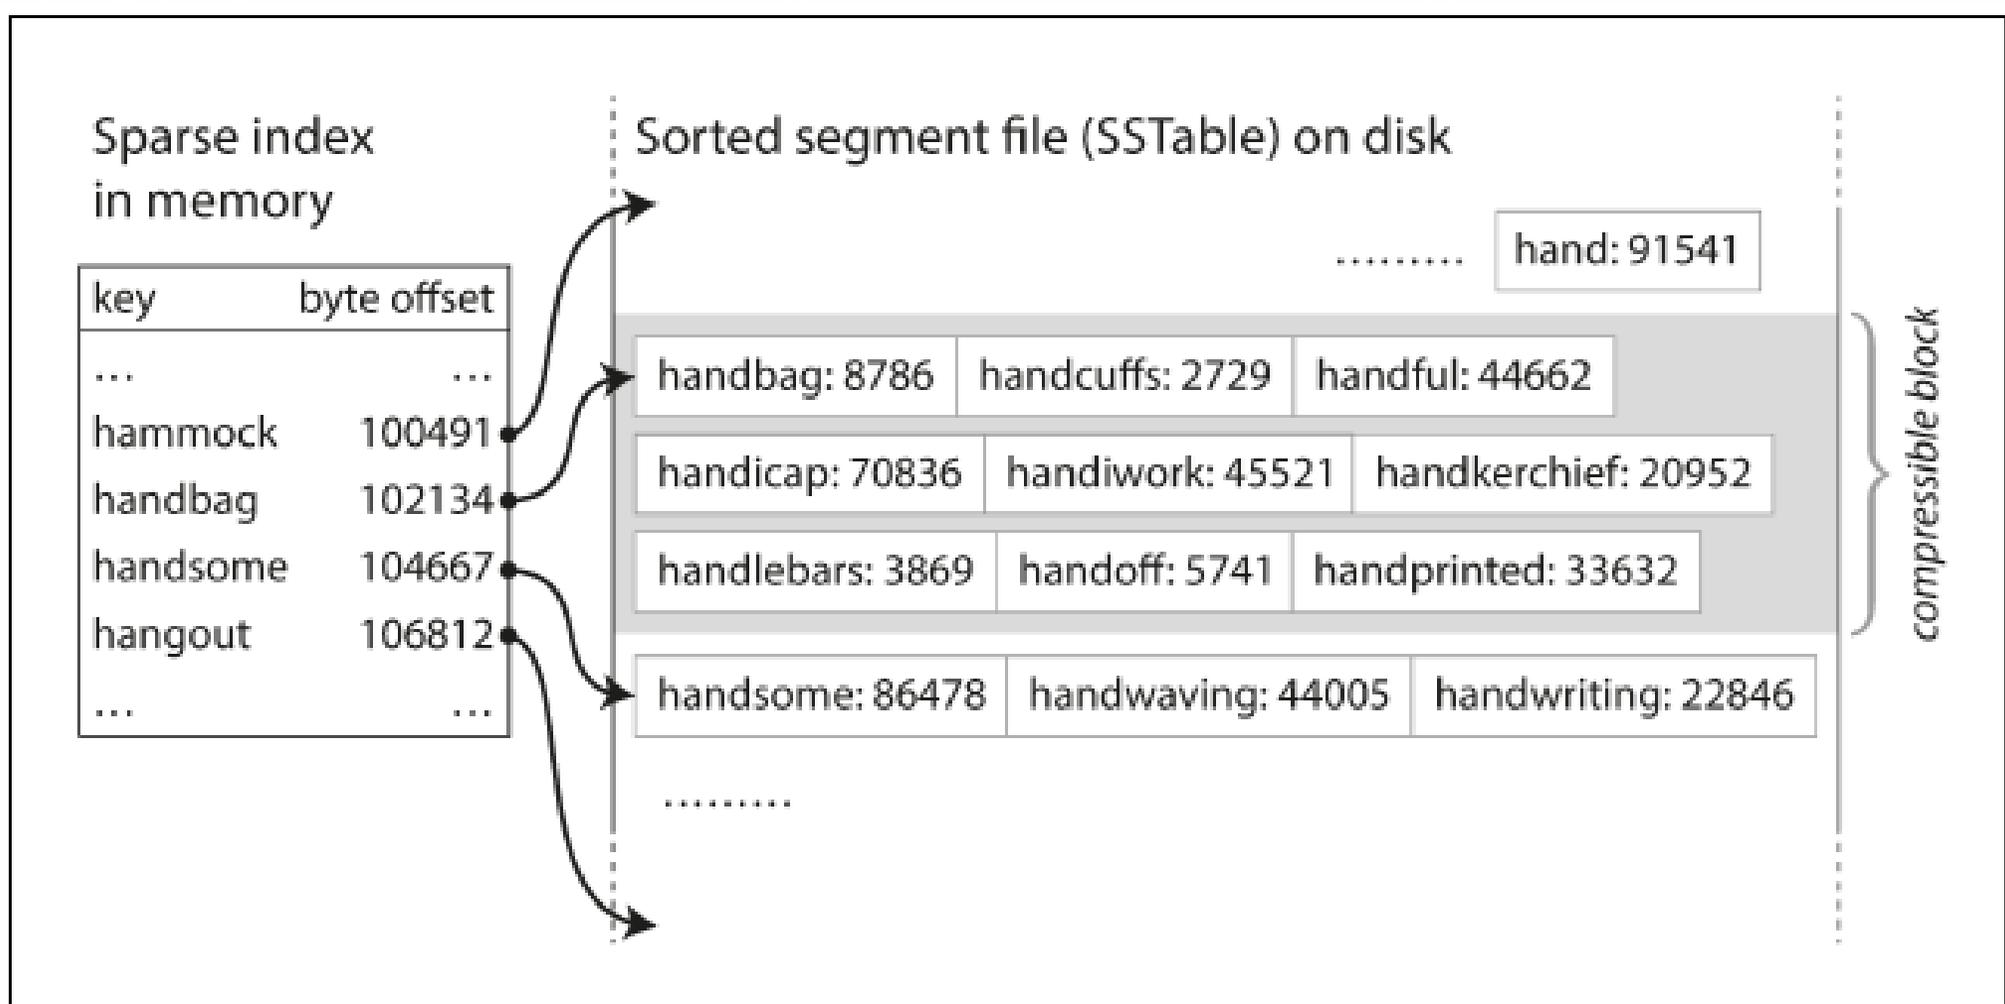
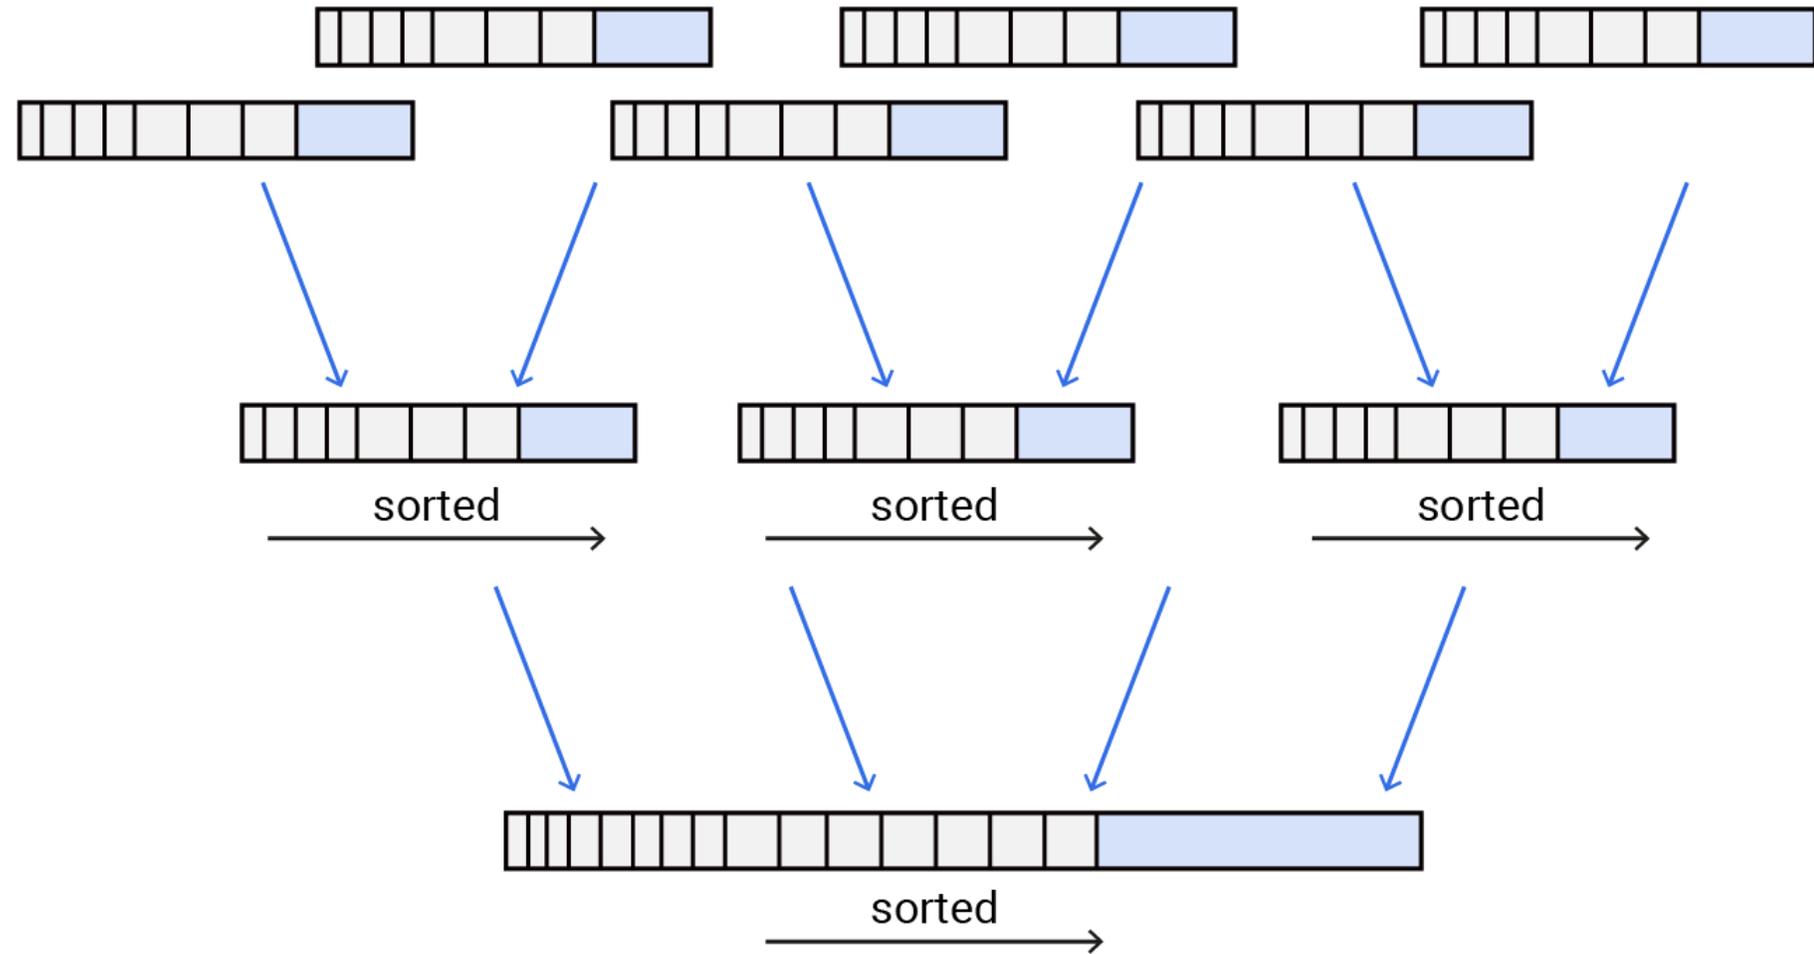


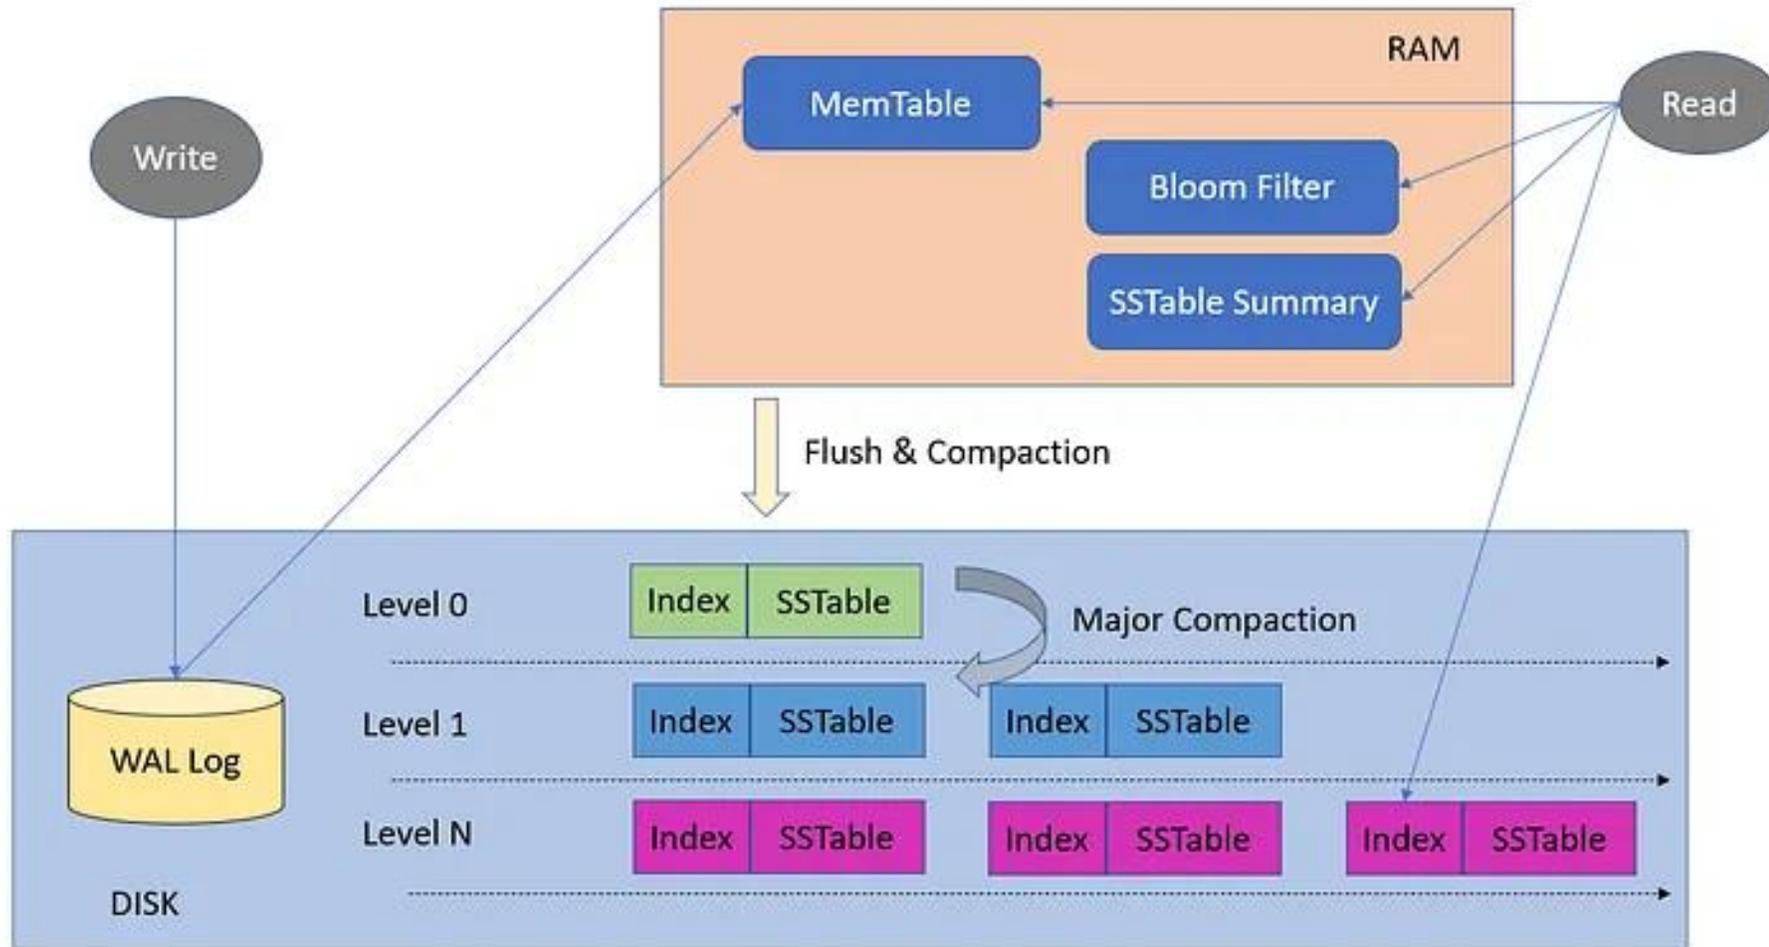
Figure 3-5. An SSTable with an in-memory index.

- An LSM (log-structured merge) tree is a data structure used in modern databases for efficient handling of high-volume writes by buffering data in memory before periodically flushing it to disk



LSM-tree

Log structured Merge Tree



Diverse range of LSM use cases

