

CLOUDEXIFY

PROJECT #: 230601

Instructor: Dr. Sarel Cohen



OUR TEAM



Ori
Braverman



Shachar
Levy



Maayan
Mashhadi

THE RESEARCHERS



Dr. Sarel Cohen



Moshik Hershkovitch

Senior Research Scientist @ IBM
Research
Ph.D. Candidate @ TAU



Daniel Cohen

M.Sc. Student
Supervised by Dr. Sarel
Cohen & Moshik
Hershkovitch



Yarin Pinyan

Director of Security
Research @ Startup



THE PROBLEM

- Lack of efficient secondary indexing for unsorted data.
- Existing indexing methods are inadequate for handling unsorted data effectively.
- Difficulty in retrieving specific information from unsorted datasets.
- Absence of a straightforward solution for organizing and accessing unsorted data efficiently.

TARGET & INTENTION

- 01** Address the challenge of efficient secondary indexing for unsorted data.
- 02** To improve query performance, increased flexibility, and reduced data storage overhead in DynamoDB tables, by creating secondary index using LSI.
- 03** Extend a novel approach that enhances the retrieval of specific information from unsorted datasets.
- 04** Colaborate with IBM research and startups for industrial aplicativity.
- 05** Improve the overall performance and effectiveness of secondary indexing methods.



SOLUTION

Permutation Vector

For mapping predictions into the underlying unsorted base data.

Learned Index

Which maps a lookup key to a bounded search range.

Fingerprint Vector


To prune unnecessary accesses during equality lookups.



ARCHITECTURE AND THE PROCESS OF LSI

The local search uses the permutation vector to locate corresponding entries in the base data.

The fingerprint vector prunes unnecessary accesses.



OTHER SOLUTIONS AND BENCHMARKS

BTree is a balanced tree-based index structure that organizes data in a sorted order. It is commonly used in traditional database systems. BTree supports efficient insertion, deletion, and search operations. It is suitable for both disk-based and memory-based systems.

RobinHash is a hash table-based index structure that uses a robin-hood hashing algorithm. It provides fast lookup and insertion operations, making it suitable for in-memory databases. RobinHash consumes more space compared to other index structures.

LSI

LSI is a learned index structure that can index unsorted data. It uses machine learning techniques to build a model that approximates the distribution of the data. LSI achieves a good trade-off between space efficiency and lookup performance.

BTree

The Adaptive Radix Tree (ART) is a data structure designed for efficient indexing and retrieval of key-value pairs in databases. It adapts its structure dynamically based on the data distribution, resulting in efficient search operations. ART is known for its low memory overhead and high performance.

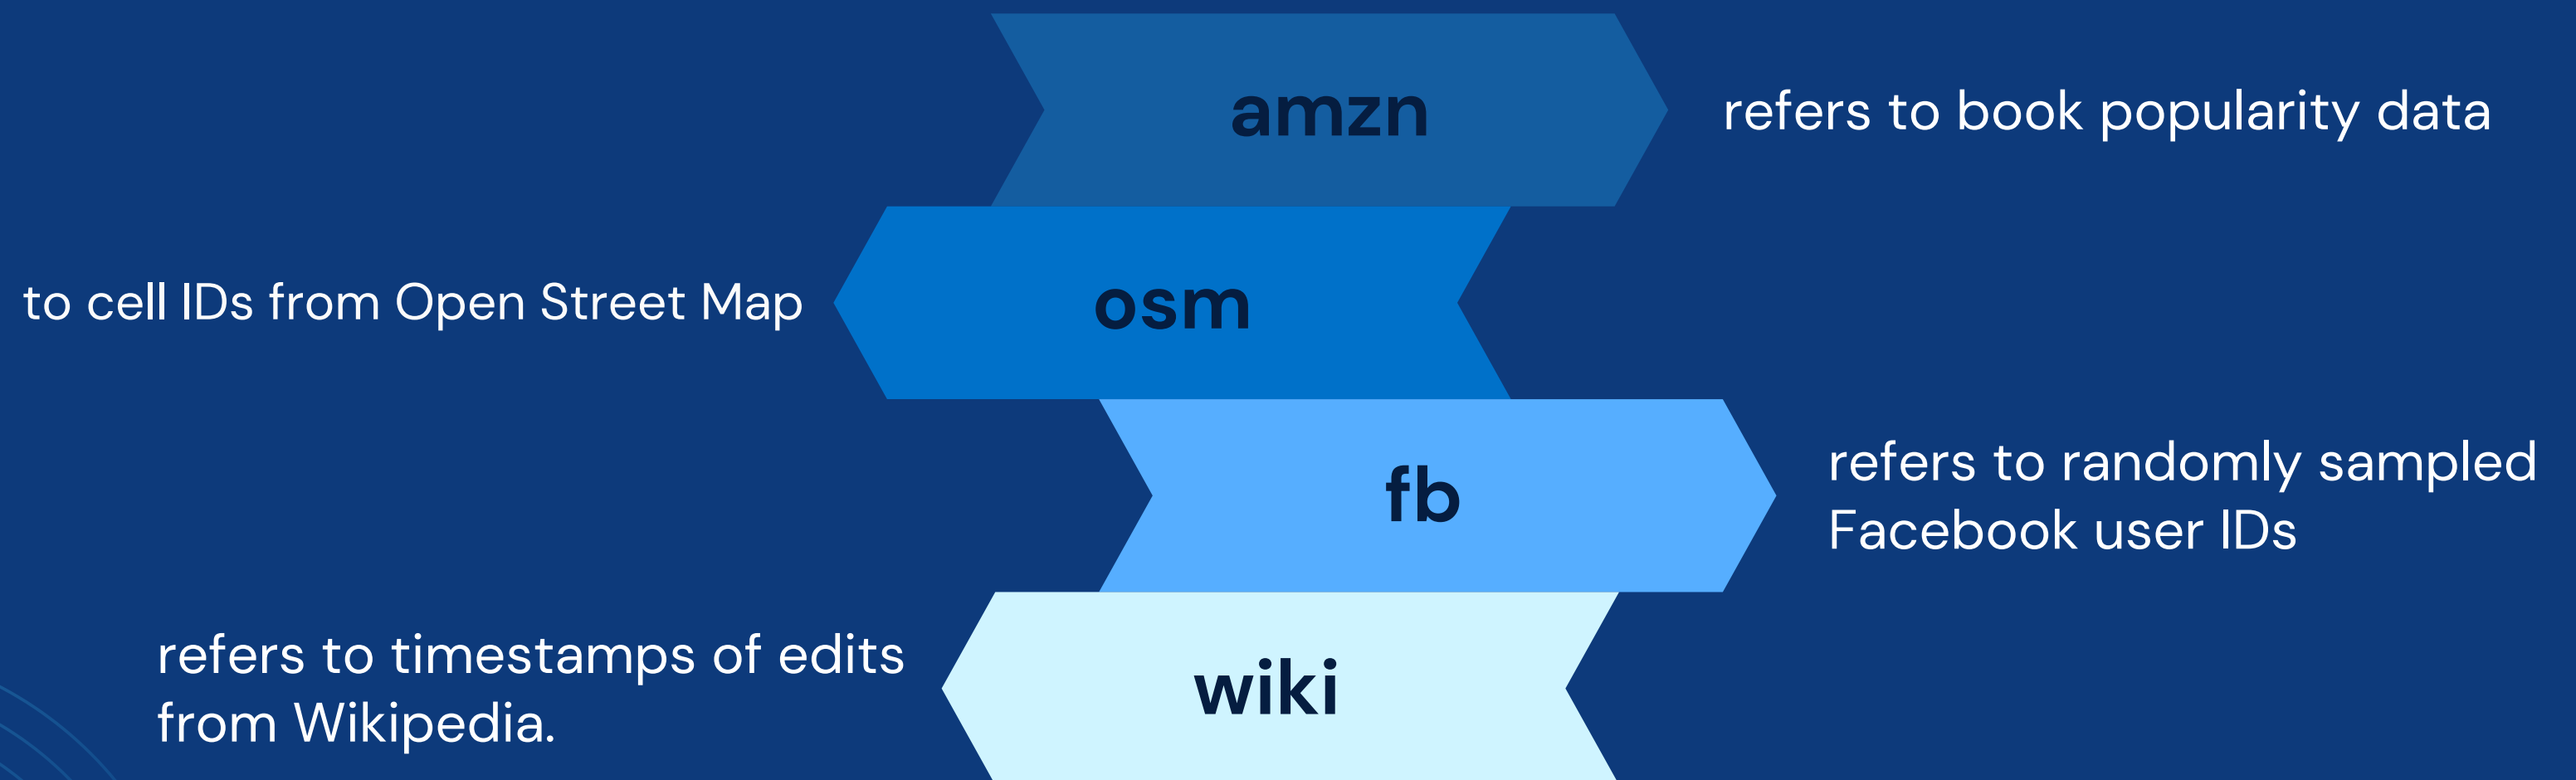
ART

RobinHash

compare their performance in terms of build times, lookup latency, and space usage.

Comparison Between Other Solutions

Datasets used to evaluate LSI:

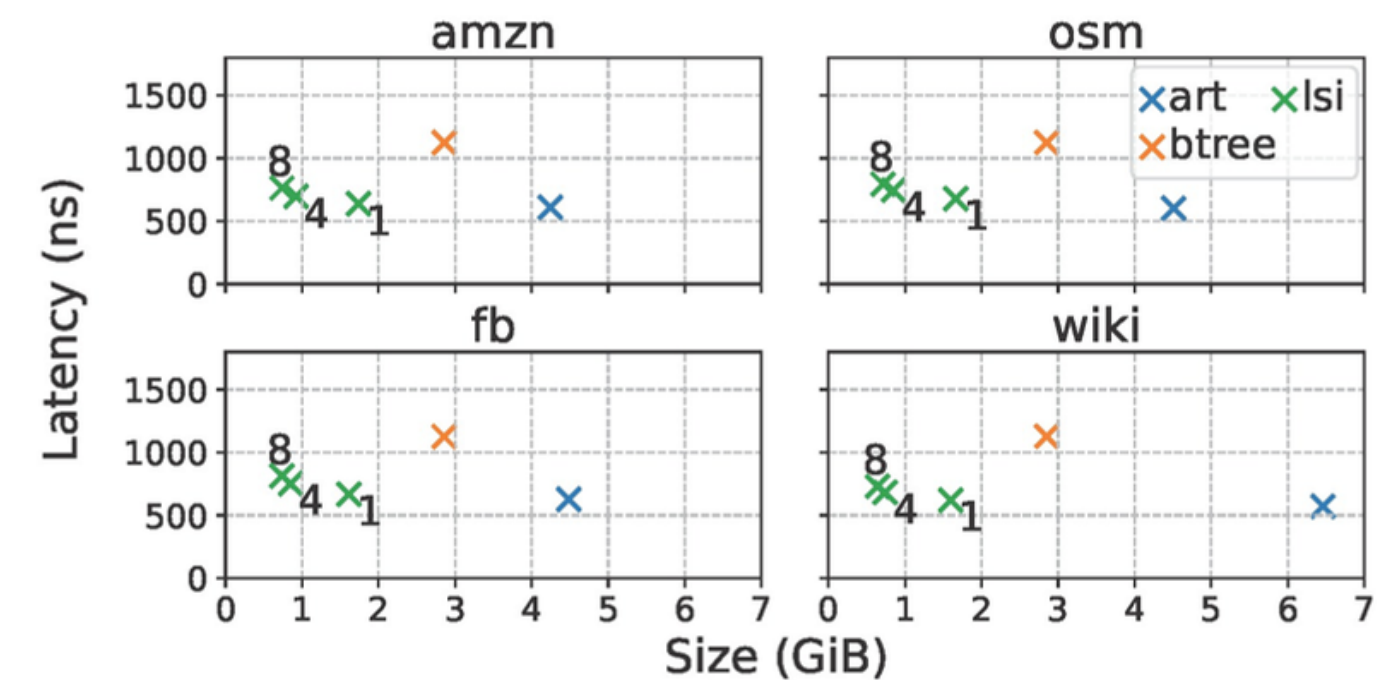



COMPARISON BETWEEN OTHER SOLUTIONS

Build times in seconds for different index structures.

The BTree index structure achieves the lowest build times, followed by RobinHash and LSI with an error bound of 8. ART, on the other hand, has the highest build times.


lower-bound lookups using non-existing keys. The graph represents the latency (in nanoseconds) of the lookups, and the text annotations on the graph indicate the error bounds.



COMPARISON BETWEEN OTHER SOLUTIONS

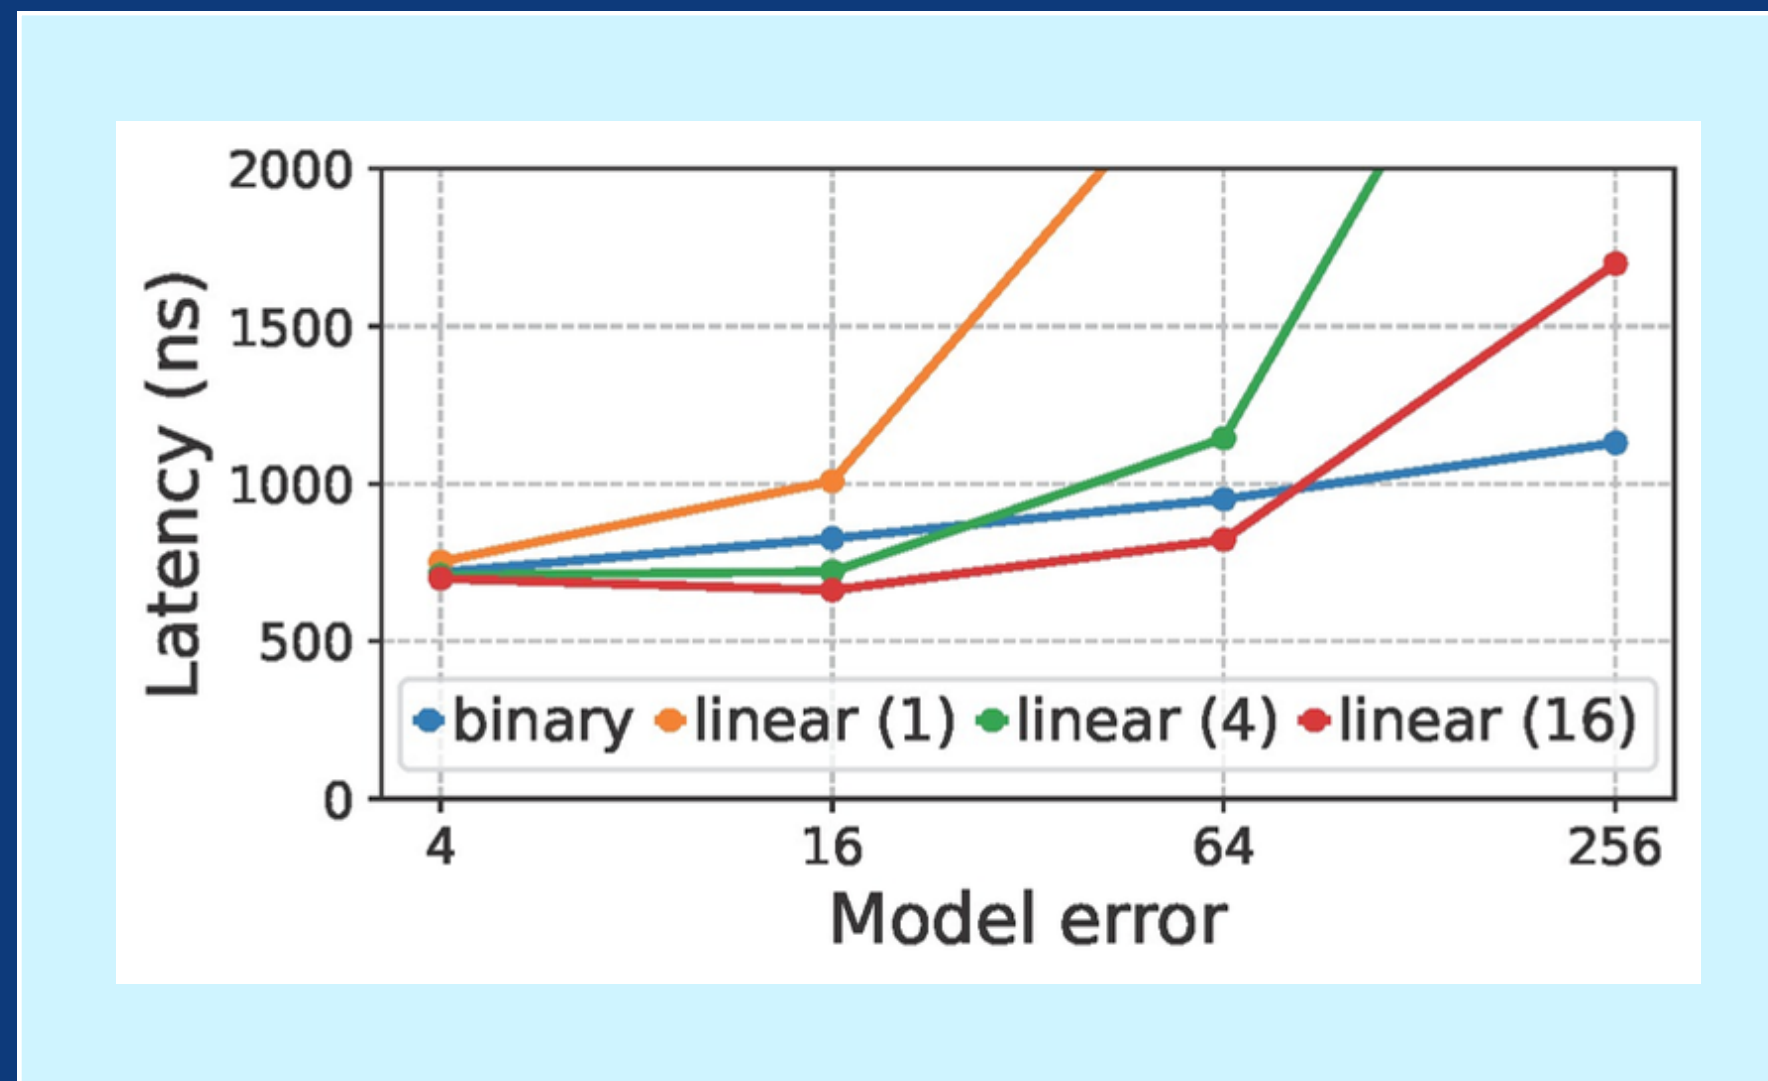

Equality lookups on the amzn dataset comparing LSI to RobinHash.

RobinHash has a bit less latency than LSI's. However, RobinHash also consumes 4 times the amount of space compared to LSI.

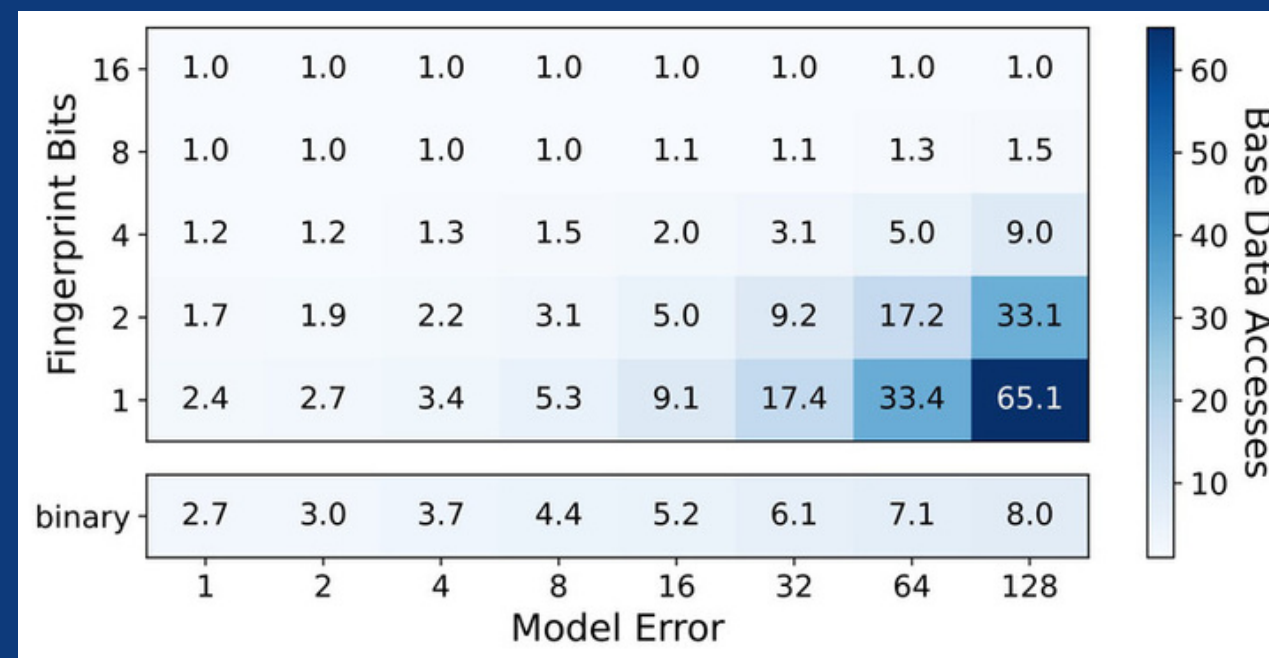
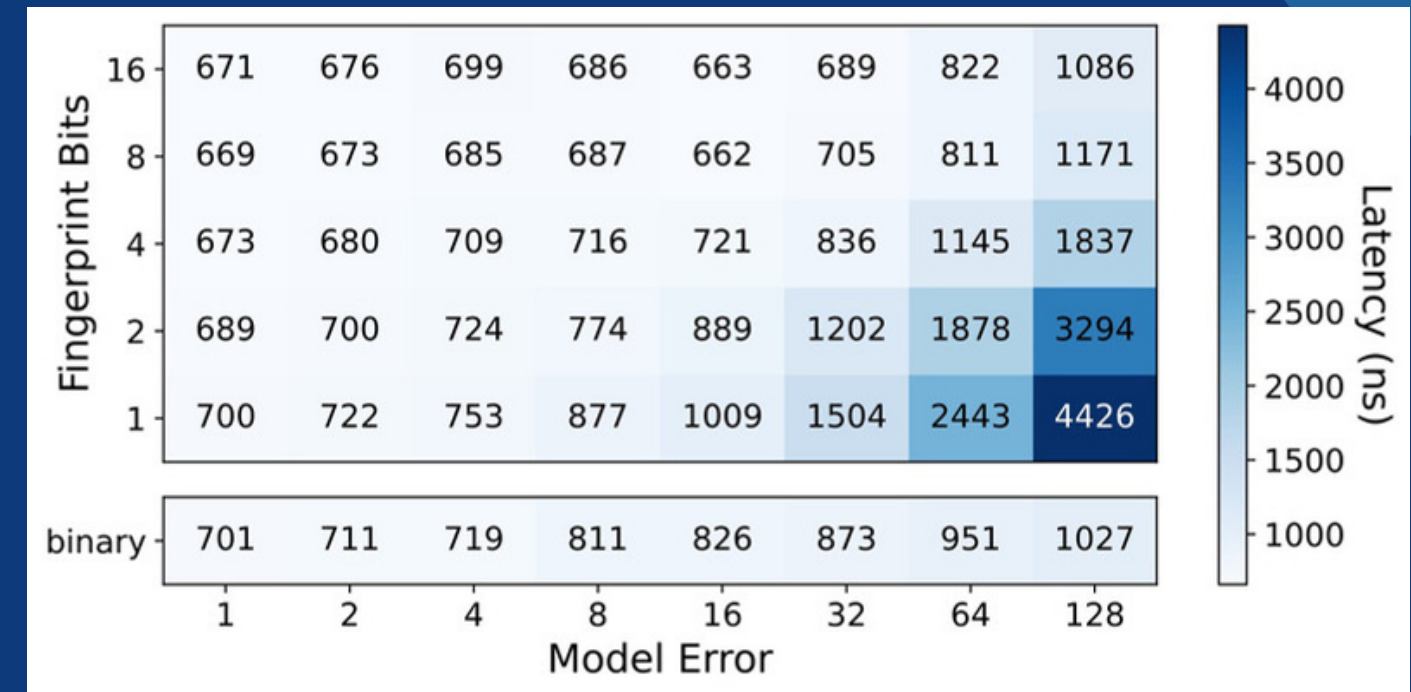


COMPARISON BETWEEN SEARCH METHODS

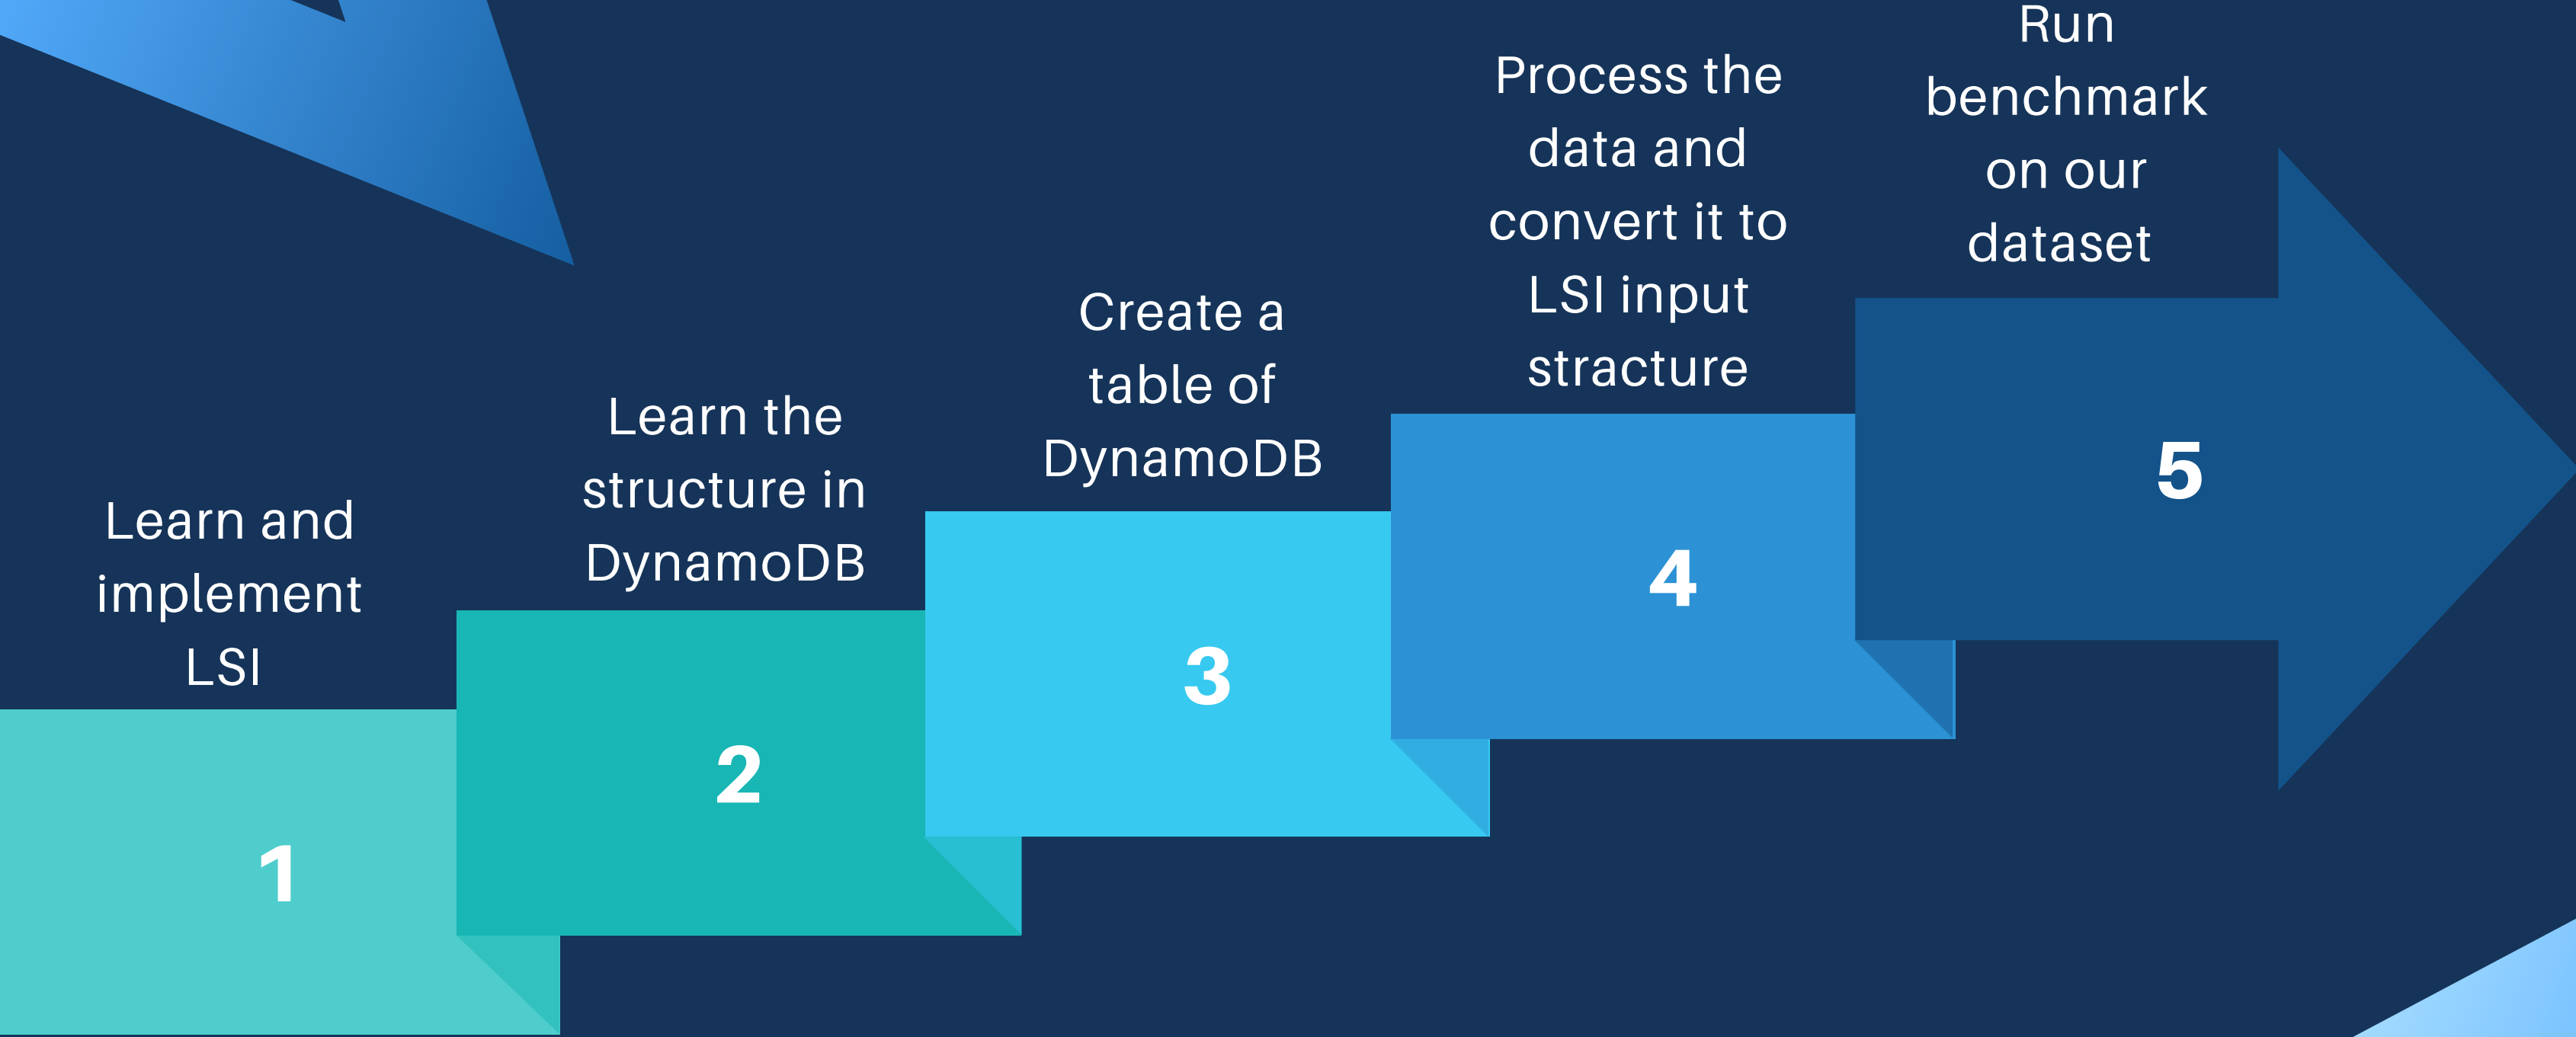
Compares binary search and linear search with different fingerprint sizes. The brackets in the figure indicate the number of fingerprint bits used for each search method.



ANALYSIS OF FINGERPRINT VECTOR



OUR WORK PROCESS





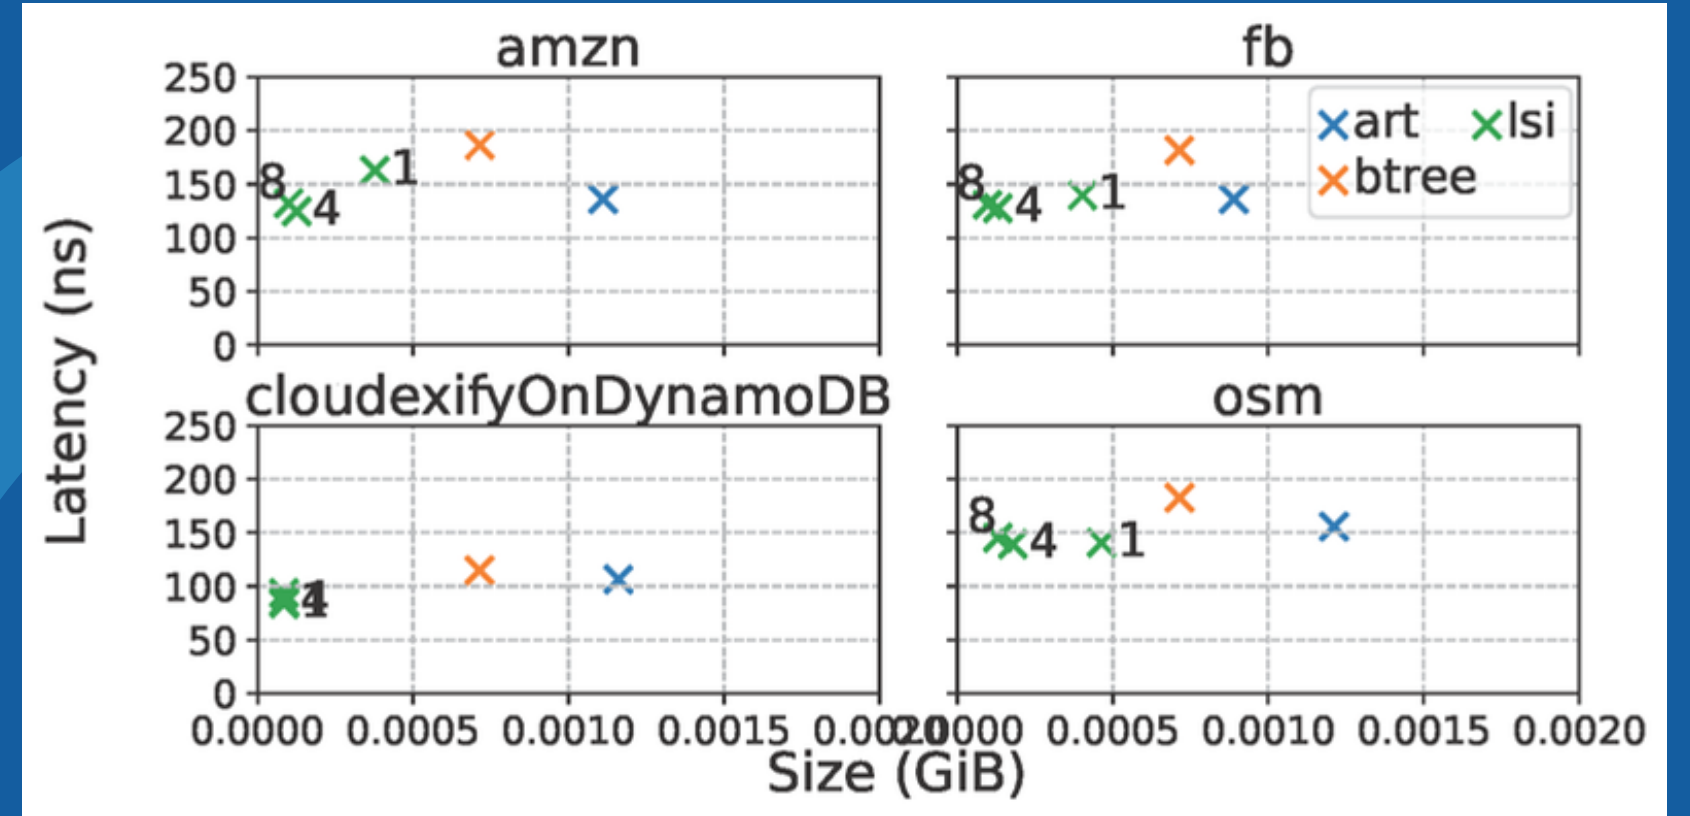

DYNAMO-DB



Amazon
DynamoDB

- Amazon DynamoDB is a fully managed NoSQL database service offered by Amazon Web Services (AWS), designed to provide high availability, scalability, and low-latency access to data.
- Its underlying structure is based on a key-value store, where each item in a table is uniquely identified by a primary key.
- DynamoDB supports secondary indexes, which enhance query flexibility.
- These secondary indexes use a B-tree data structure for efficient indexing, making it possible to retrieve data quickly based on attributes other than the primary key.

DynamoDB Benchmark results



Thank's For Watching

Connect with us.

 [https://github.com/shacharlevy/
Cloudexify](https://github.com/shacharlevy/Cloudexify)

