Karlsruhe Institute of Technology

Hochschule Karlsruhe
Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

The Eight International Conference on Advances in Databases, Knowledge, and Data Applications

June 26 - 30, 2016 - Lisbon, Portugal

GraphSM Workshop

Graph Storage Considerations

Andreas Schmidt

**Department of Informatics and
Business Information Systems
University of Applied Sciences Karlsruhe
Germany**

**Institute for Applied Computer Sciences
Karlsruhe Institute of Technologie
Germany**

Observation

The time to answer a query is mainly determined by the amount of data that has to be loaded from disk / from main memory

Example Szenario

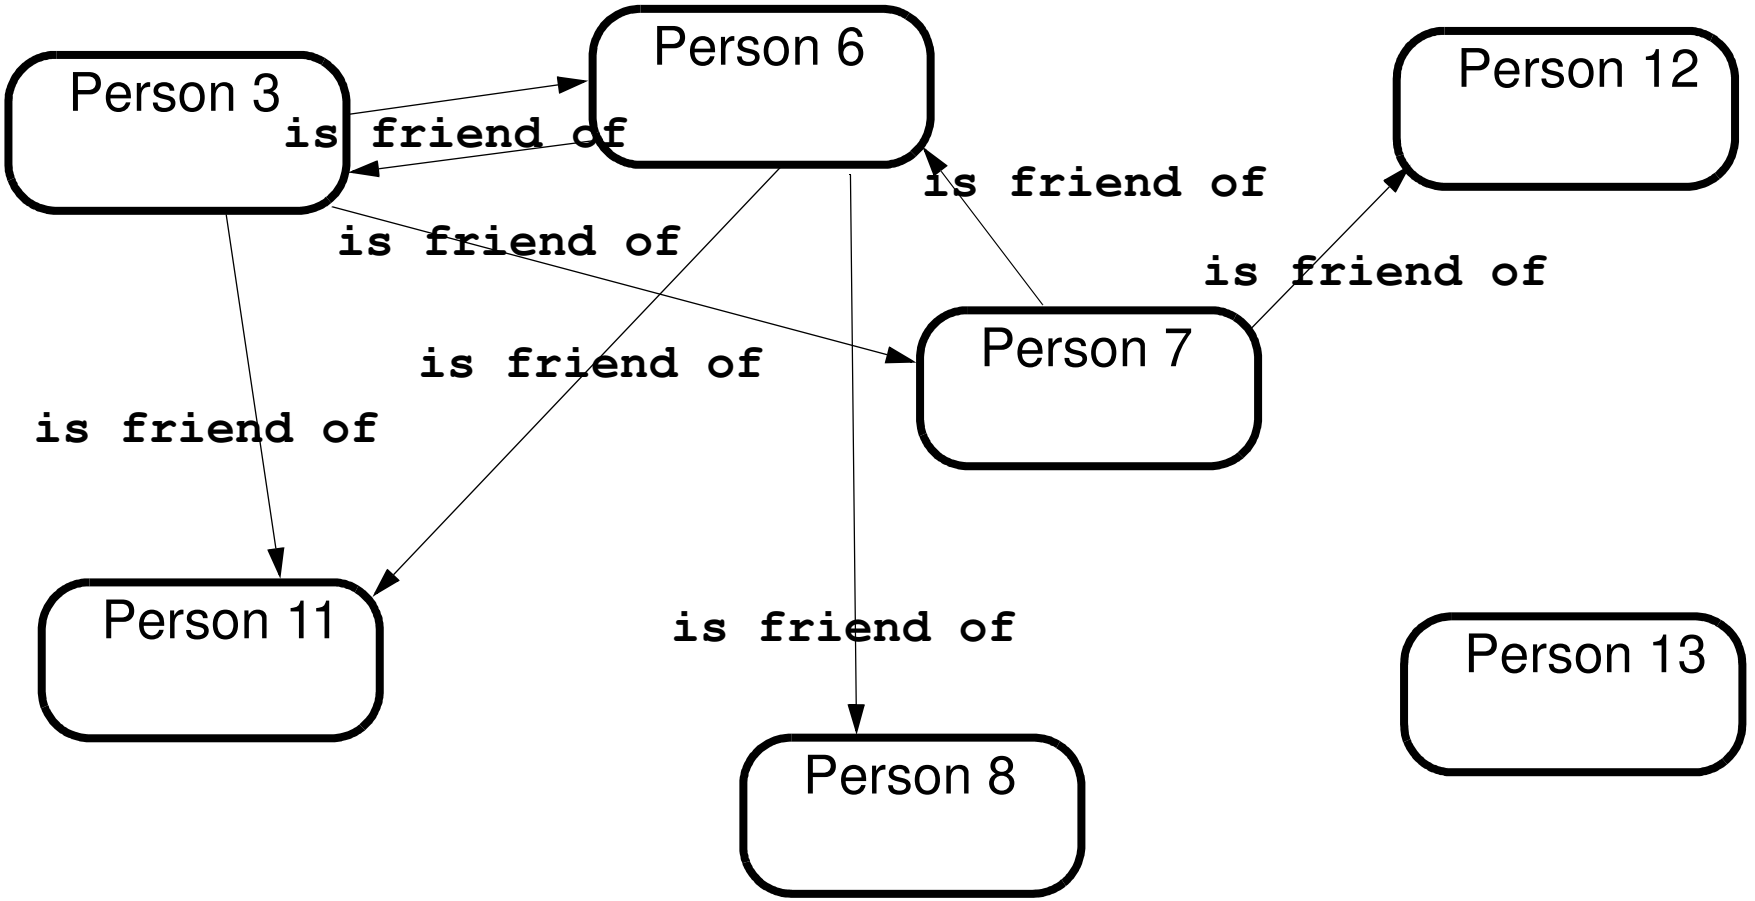

Example Query

- Example Query: Show all my friends of grade n
  - Dataset 1: 1000 users, on average 50 friends/person

grade	execution time MySQL (sec.)	execution time neo4j db (sec.)	number results
2	0.028	0.04	~900
3	0.213	0.06	999
4	10.273	0.07	999
5	92.613	0.07	999

Source: Neo4j in Action, Manning Verlag, 2015

Example Query

- Example Query: Show all my friends of grade n
  - Dataset 2: 1.000.000 users, on average 50 friends/person

grade	query time MySQL (sec.)	query time neo4j (sec.)	number results
2	0.016	0.010	~2500
3	30.267	0.168	~125.000
4	1,543.505	1.359	~600.000
5	not finished	2.132	999.999

Source: Neo4j in Action, Manning Verlag, 2015

Relational Database - Data Represenation ?

- Table is_friend_of

person	friend
3	6
3	7
3	11
6	3
6	8
6	11
6	17
7	6
7	12
...	...

- Index is_friend_of(person, friend)
  Complexity: O(log(n))

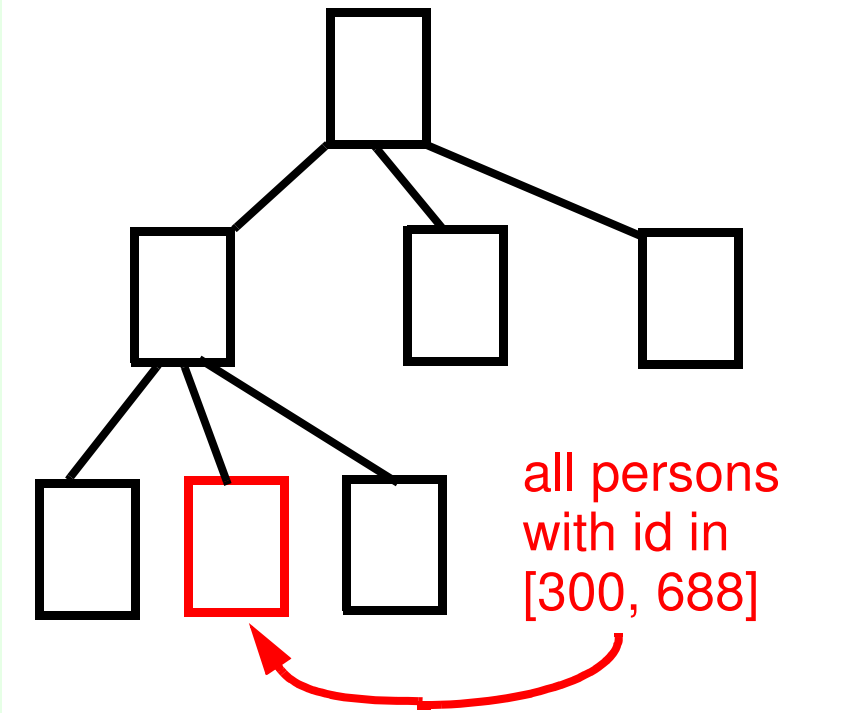

Graph Database - Data Representation (extract)

- Nodes (fixed record length):

```
01: node 1:
02: node 2:
03: node 3:1-3, ...
04: node 4:
05: node 5:
06: node 6:4-7, ...
07: node 7:8-9, ...
08:
09:
10:
11:
```

- Relations (fixed record length):

```
01: 3-> 6, is_friend_of, -,2
02: 3-> 7, is_friend_of, 1,3
03: 3->11, is_friend_of, 2,-
04: 6-> 3, is_friend_of, -,5
05: 6-> 8, is_friend_of, 4,6
06: 6->11, is_friend_of, 5,7
07: 6->17, is_friend_of, 6,-
08: 7-> 6, is_friend_of, -,9
09: 7->12, is_friend_of, 8,-
10:
11:
```

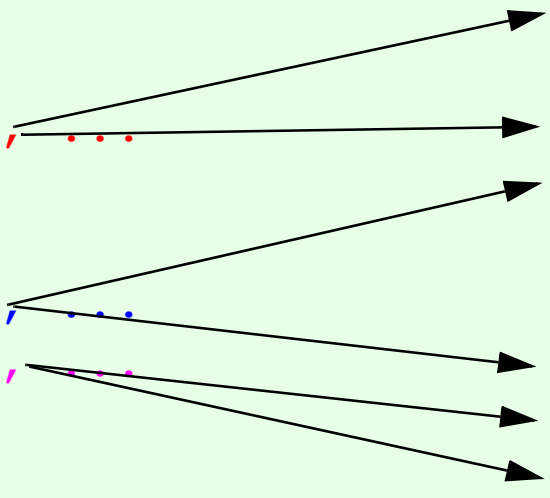

Complexity: O(1)

Other Example:

- Query: Return all pairs of persons who have at least 2 common friends, but did not know each other
- Problem with this query?

Numbers everyone should know (from [Lad09])

• L1 cache reference	0.5 ns
• Branch mispredict	5 ns
• L2 cache reference	7 ns
• Mutex lock/unlock	100 ns
• Main memory reference	100 ns
• Compress 1K bytes with Zippy	10,000 ns
• Send 2K bytes over 1 Gbps network	20,000 ns
• Read 1 MB sequentially from memory	250,000 ns
• Round trip within same datacenter	500,000 ns
• Disk seek	10,000,000 ns
• Read 1 MB sequentially from network	10,000,000 ns
• Read 1 MB sequentially from disk	30,000,000 ns
• Send packet CA->Netherlands->CA 1	50,000,000 ns

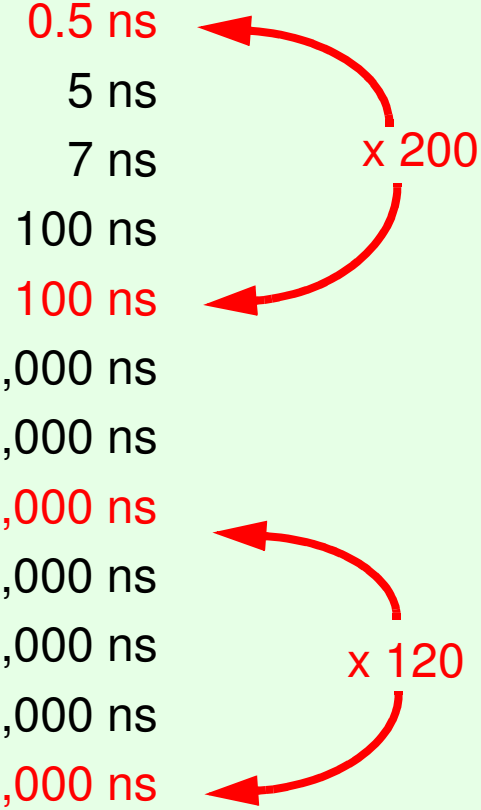

[Lad09] http://www.cs.cornell.edu/projects/ladis2009/talks/dean-keynote-ladis2009.pdf

Sequential vs. Random read (from [Jac09])

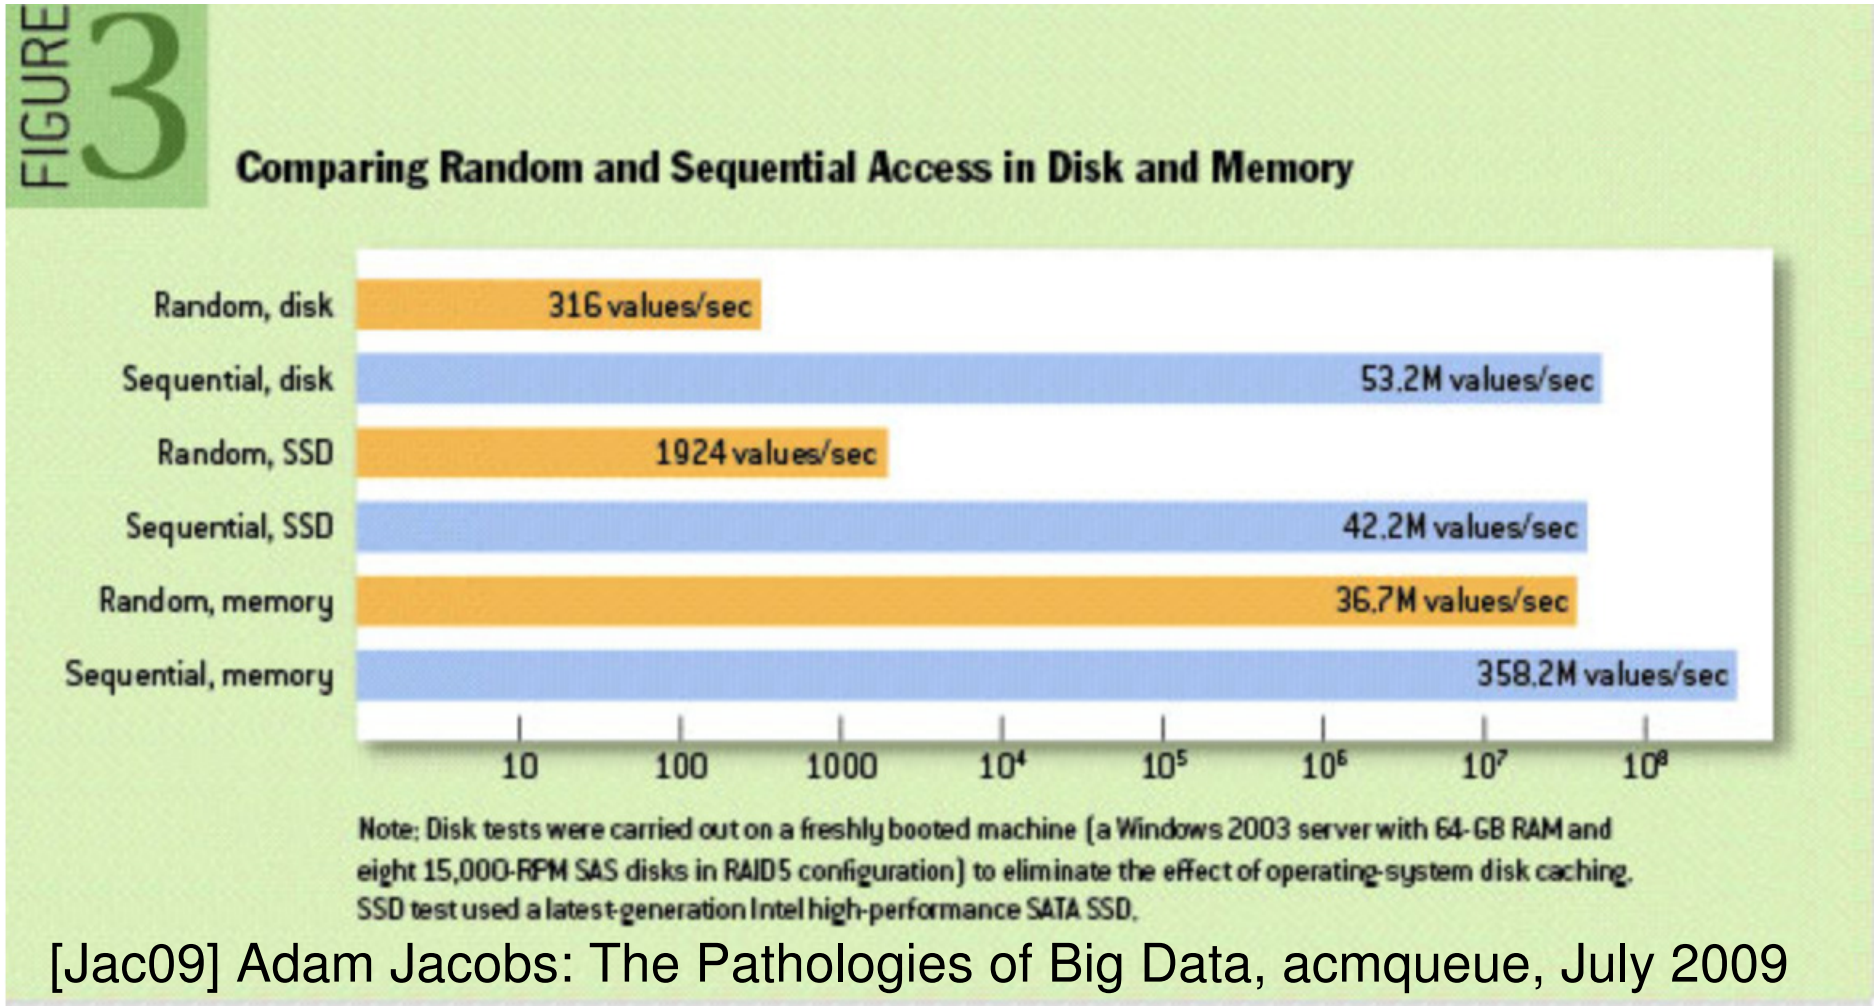

[Jac09] Adam Jacobs: The Pathologies of Big Data, acmqueue, July 2009

Cache & Cacheline

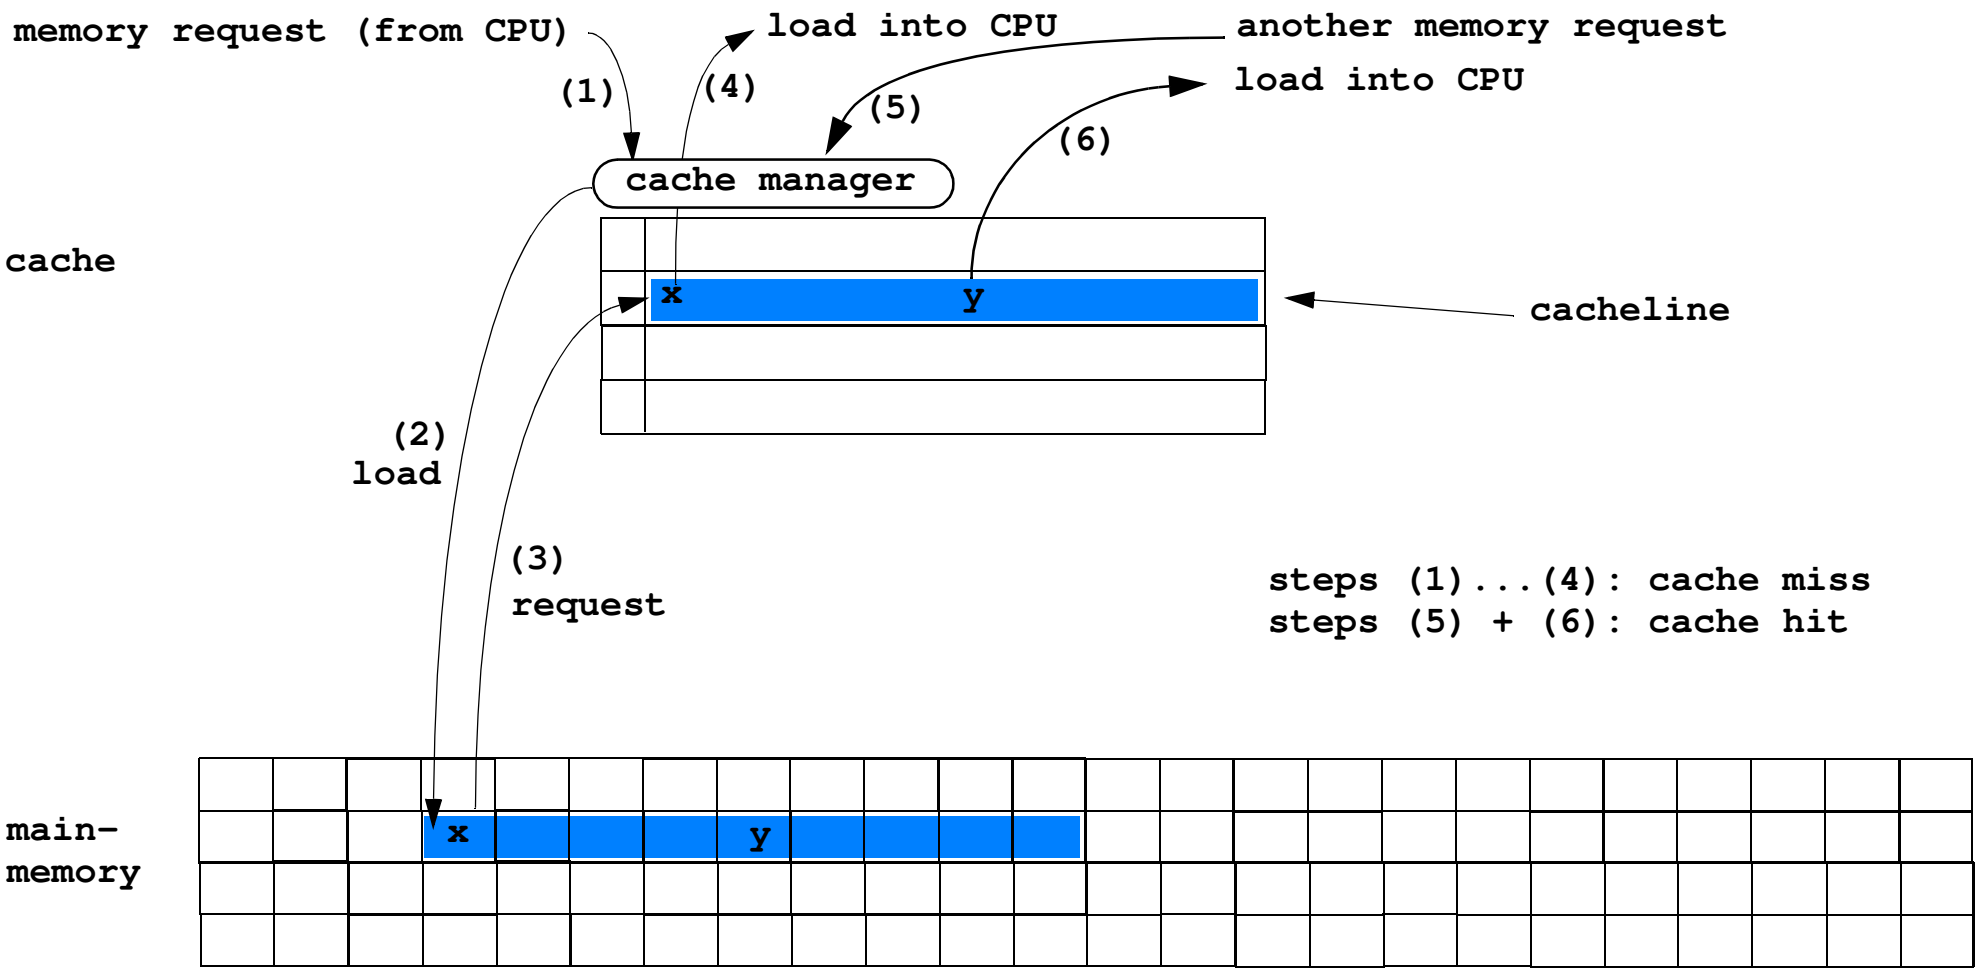

Resumee

- Native Graph databases support index-free adjacency
- Choose the right database type, according for your application demands
- Knowing the underlying storage model can help to chose the right data-base type
- Data from the first level cache can be accesed 200 times faster, compared to data that resides in main memory (cache consciousness)
- Sequential access is between one and five orders of magnitude faster than random access

that’s all ....

thanks