SQLite a File-Based Database Management System

Outline

- Introduction
- Architecture
- Benchmarks
- Conclusion

Introduction

- Popular file-based DBMS
- Row oriented storage format
- Several characteristics for popularity
  - Cross-Platform
  - Compact and self-contained
  - Reliable
  - Fast
- Designed for OLTP transaction
- Online transaction processing (OLTP)
  - E.g. view details of an order
- Online analytical processing (OLAP)
  - E.g. massive data requests for analysing

Architecture

- 4 Modules
  - Core module
  - SQL Compiler module
  - Backend module
  - Accessoires module

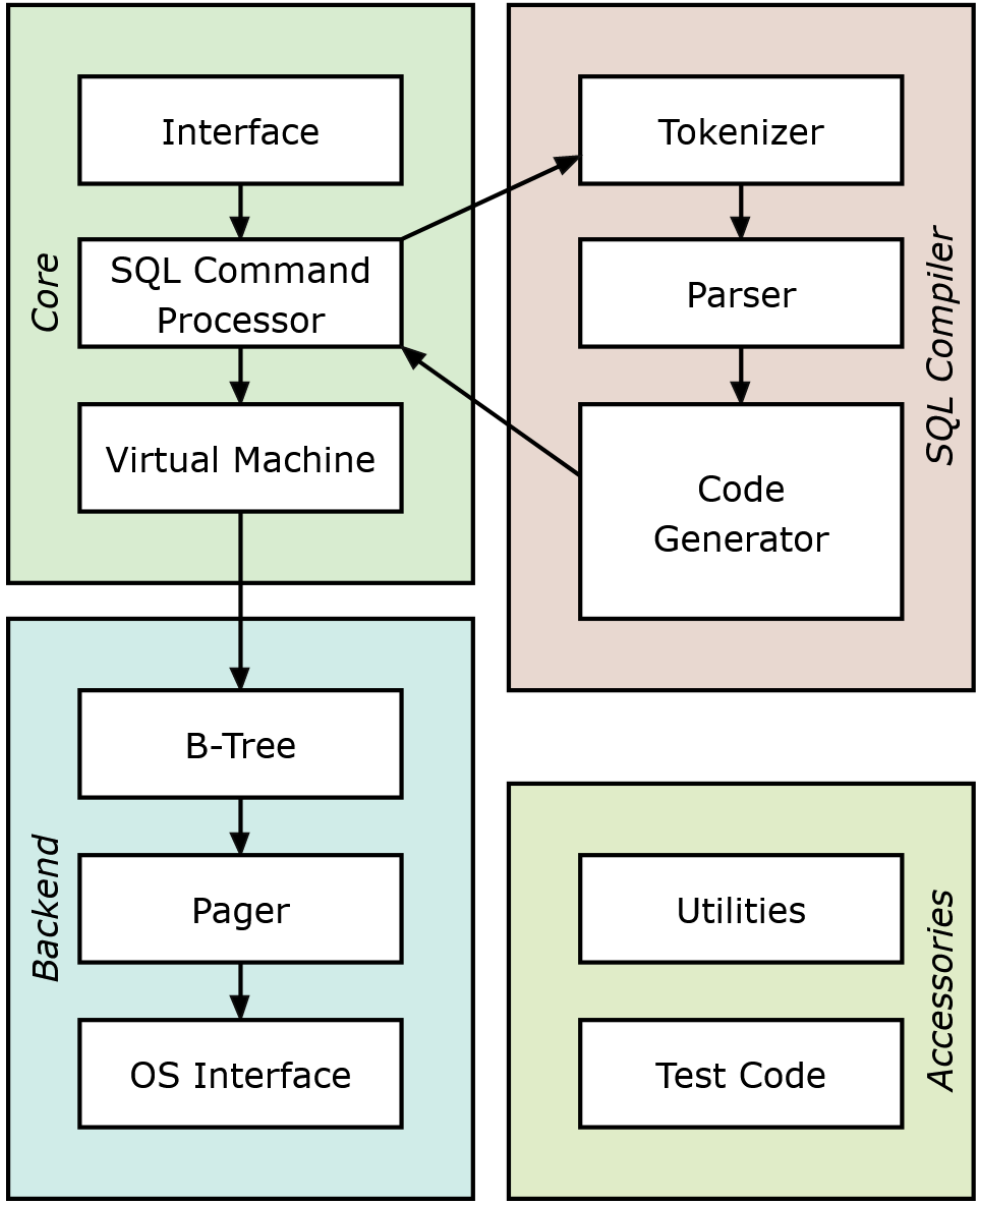

Source: https://www.sqlite.org/arch.html; Last accessed 20.05.2023

SQL Compiler module

- Tokenizer
  - Splits SQL statements into token
- Parser
  - Gives meaning to tokens
- Code generator
  - Generates runnable bytecode program

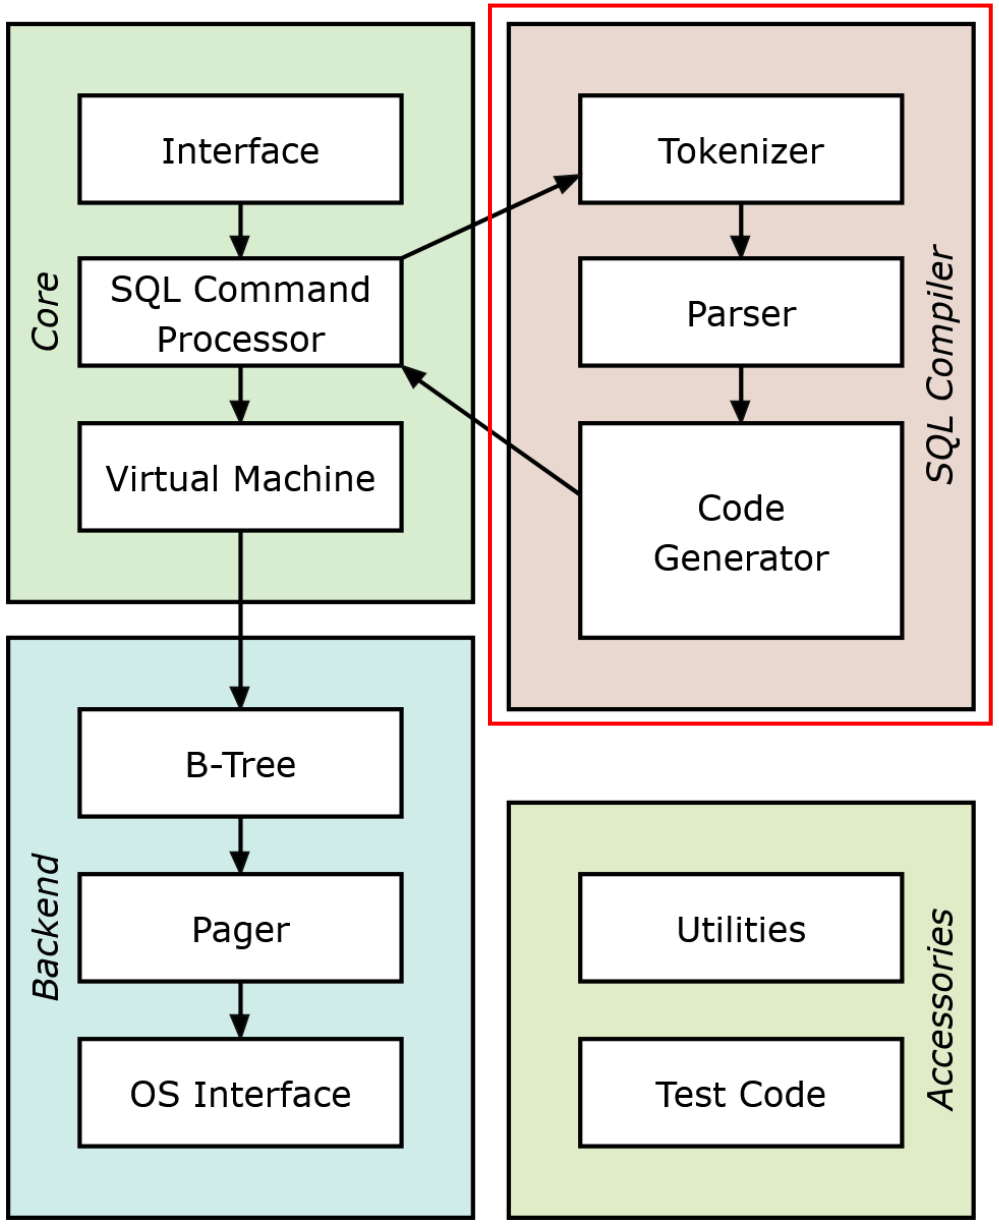

Source: https://www.sqlite.org/arch.html; Last accessed 20.05.2023

Bytecode Example

```
$ sqlite3 ex1.db
sqlite> explain delete from tbl1 where two<20;
addr  opcode         p1    p2    p3    p4             p5  comment
----  -------------  ----  ----  ----  -------------  --  -------------
0     Init           0     12    0                    00  Start at 12
1     Null           0     1     0                    00  r[1]=NULL
2     OpenWrite      0     2     0     3              00  root=2 iDb=0; tbl1
3     Rewind         0     10    0                    00
4       Column         0     1     2                    00  r[2]=tbl1.two
5       Ge             3     9     2     (BINARY)       51  if r[2]>=r[3] goto 9
6       Rowid          0     4     0                    00  r[4]=rowid
7       Once           0     8     0                    00
8       Delete         0     1     0     tbl1           02
9     Next           0     4     0                    01
10    Noop           0     0     0                    00
11    Halt           0     0     0                    00
12    Transaction    0     1     1     0              01  usesStmtJournal=0
13    TableLock      0     2     1     tbl1           00  iDb=0 root=2 write=1
14    Integer        20    3     0                    00  r[3]=20
15    Goto           0     1     0                    00
```

Source: https://www.sqlite.org/opcode.html; Last accessed 20.05.2023

Core module

- Execution Engine
  - Designed as VM (VDBE)
  - Executing SQL Statements as bytecode program
- Interacts with Backend and SQL Compiler

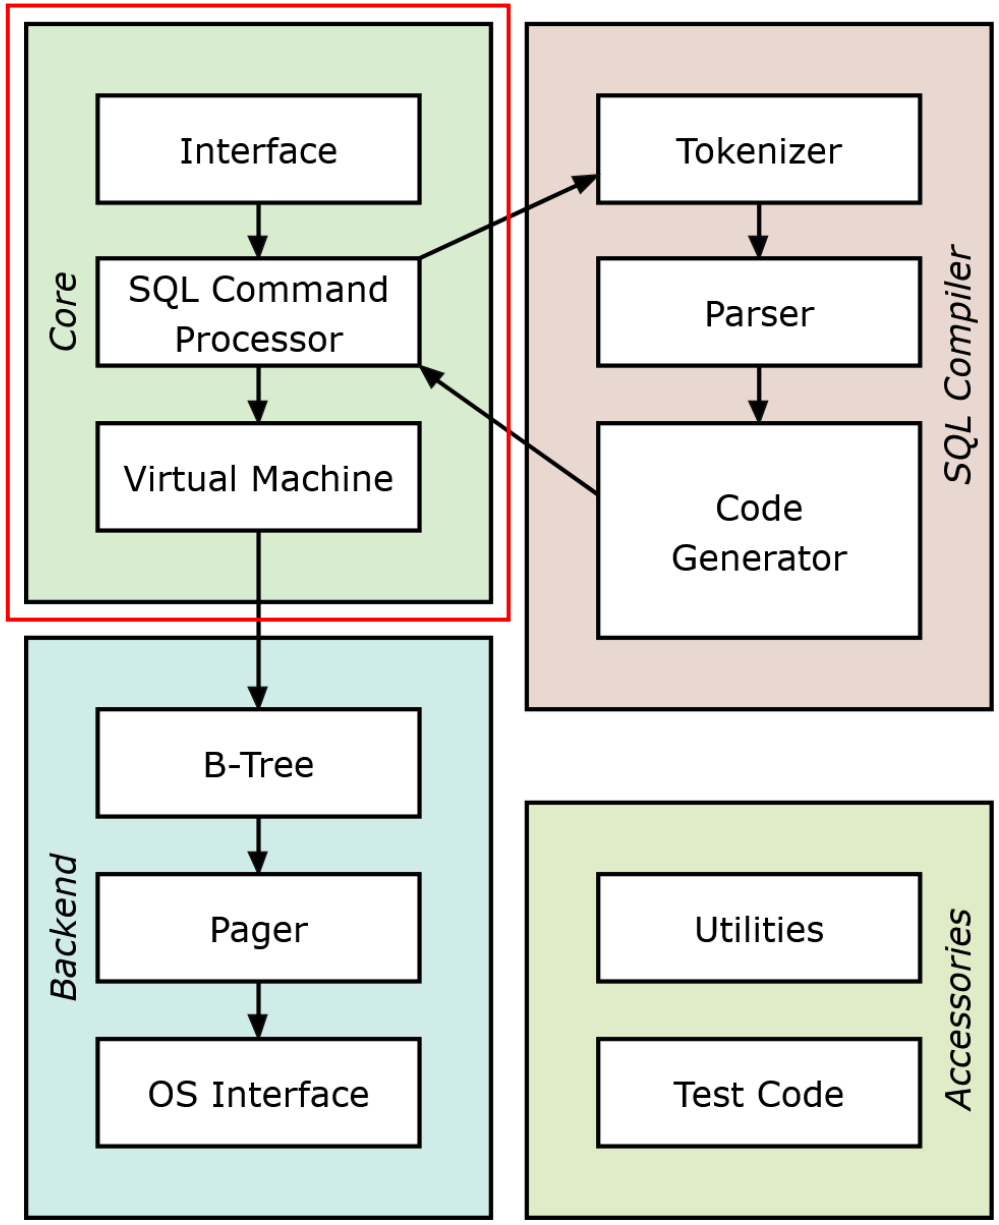

Source: https://www.sqlite.org/arch.html; Last accessed 20.05.2023

Backend module

- B-Tree
  - Table B-Tree
  - Index B-Tree
  - Consists of leaf pages and interior pages
- Pager
  - Responsible for reading, writing and caching pages
  - Makes sure requests are handled quickly, safely and efficiently
- OS Interface
  - Usage of VFS

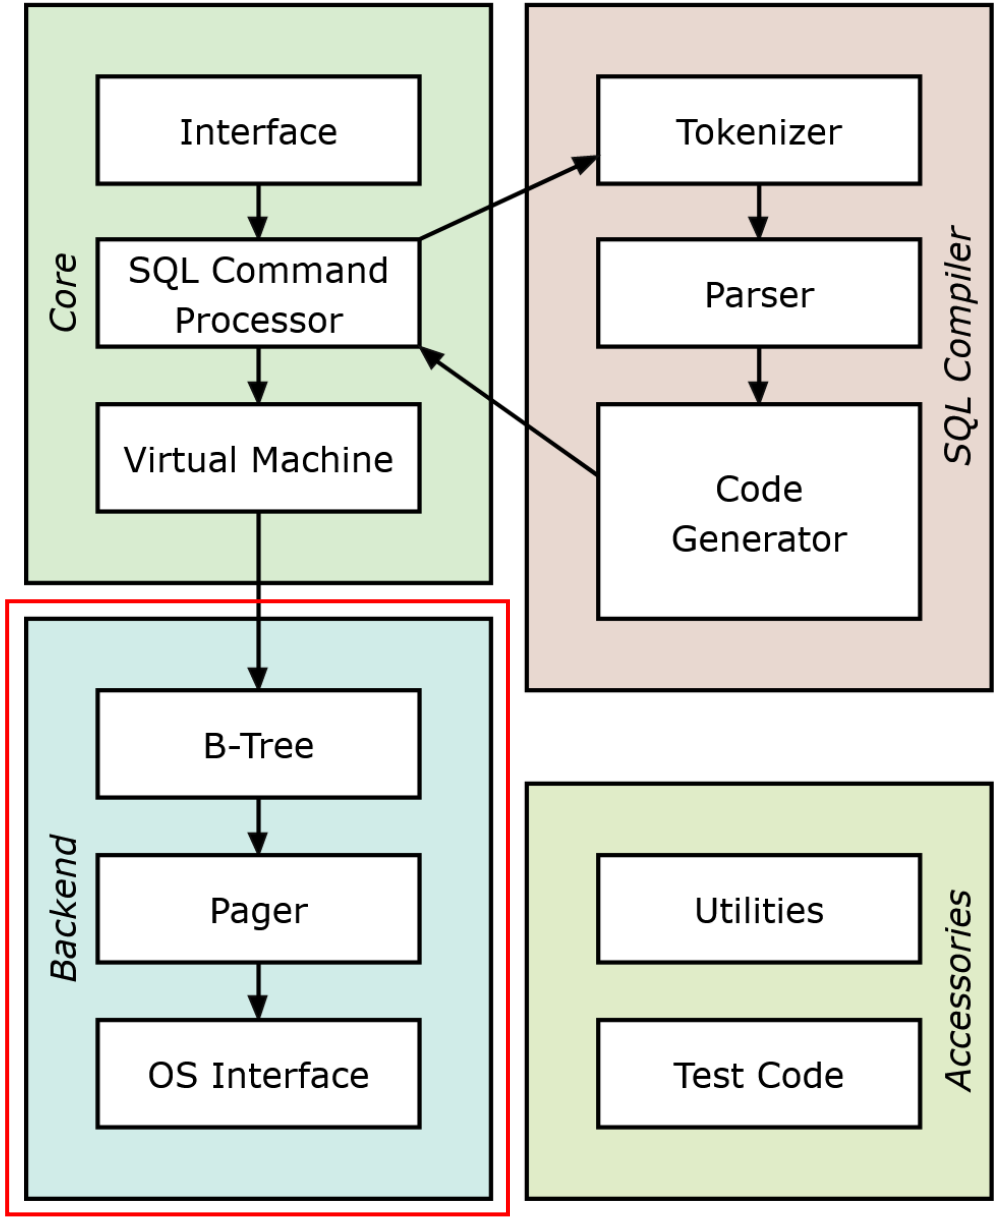

Source: https://www.sqlite.org/arch.html; Last accessed 20.05.2023

OLTP Benchmark

- TATP Benchmark
  - Random generated transactions with fixed probability
  - Vary Subscriber table size and scale other tables proportionally
  - 10 sec. warmup
  - 60 sec. evaluation
  - Transactions per sec.

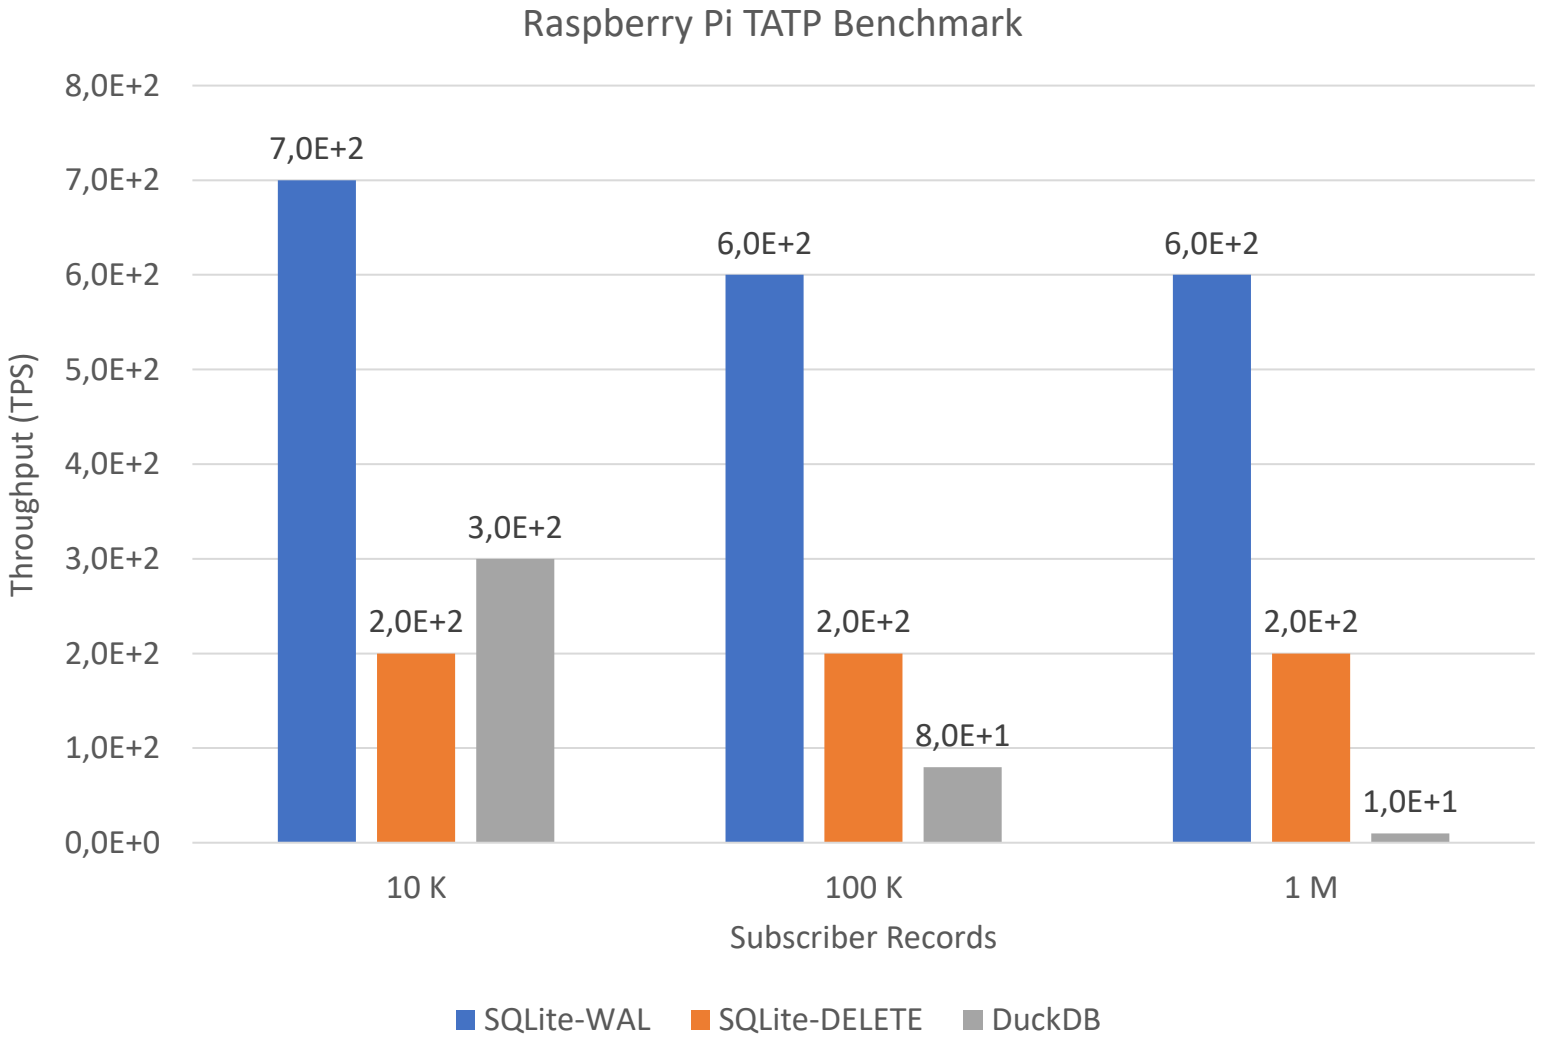

OLAP Benchmark

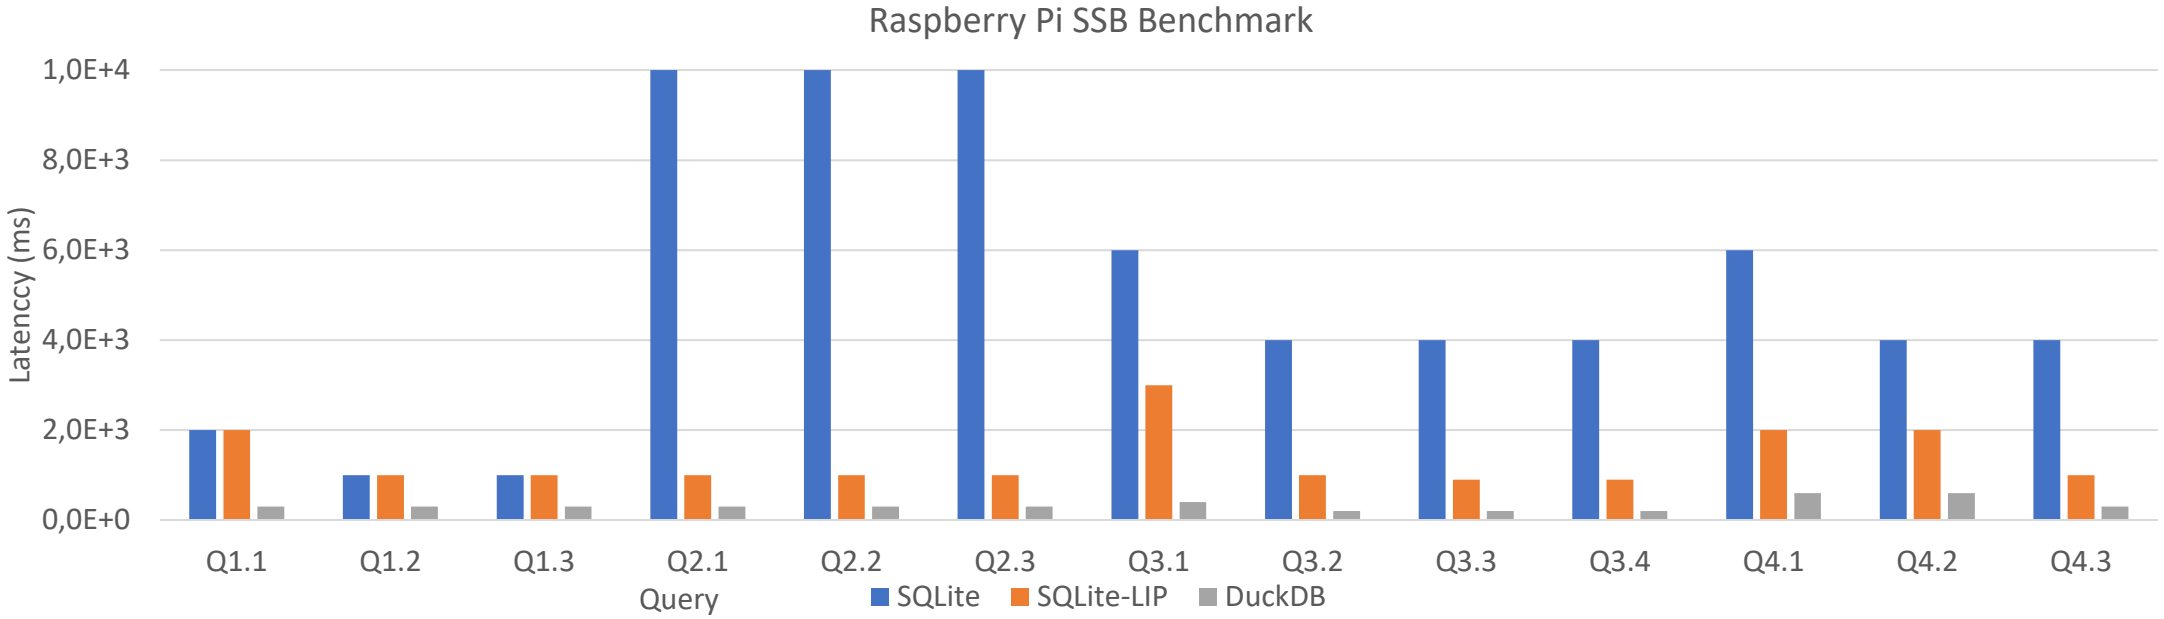

- Star Schema Benchmark (SSB)
  - Latency for a query

Conclusion

- SQLite well fitted for OLTP -> Not distributed application
- Potential for OLAP
- Continue to be an important DBMS for storing data in filesystem

Sources

- Paper:
  Kevin P. Gaffney, Martin Prammer, Larry Brasfield, D. Richard Hipp, Dan Kennedy, and Jignesh M. Patel. 2022. SQLite: Past, Present, and Future. Proc. VLDB Endow. 15, 12 (aug 2022), 3535–3547. https://doi.org/10.14778/3554821.3554842
- Website:
  https://www.sqlite.org/arch.html
  Last accessed 21.05.2023