Performance, Concurrency, Stability tests with a MySQL based catalog

Zhen Xie

CERN/EP-CMC

Outline

- MySQL based catalog
- MySQL overview
- Performance tests
  - Data integrity, transaction, locking
  - Read, write speed
  - Storage efficiency
- Concurrency tests
- Stability tests

MySQL based catalog

- File catalog: lnf->pfn lookup
- Why MySQL? Free and reported positive experience by other experiments
- Likely to be used by the September prototype
- A simplified data model for the performance test:

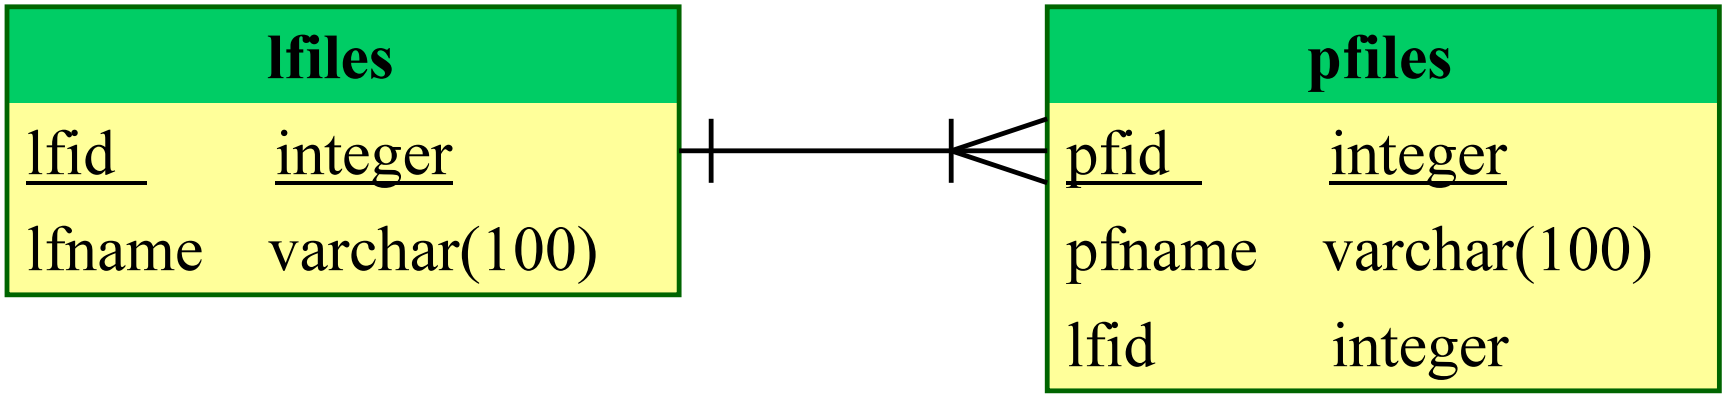

MySQL overview

- Multi-threaded client/server
- Server on a single machine

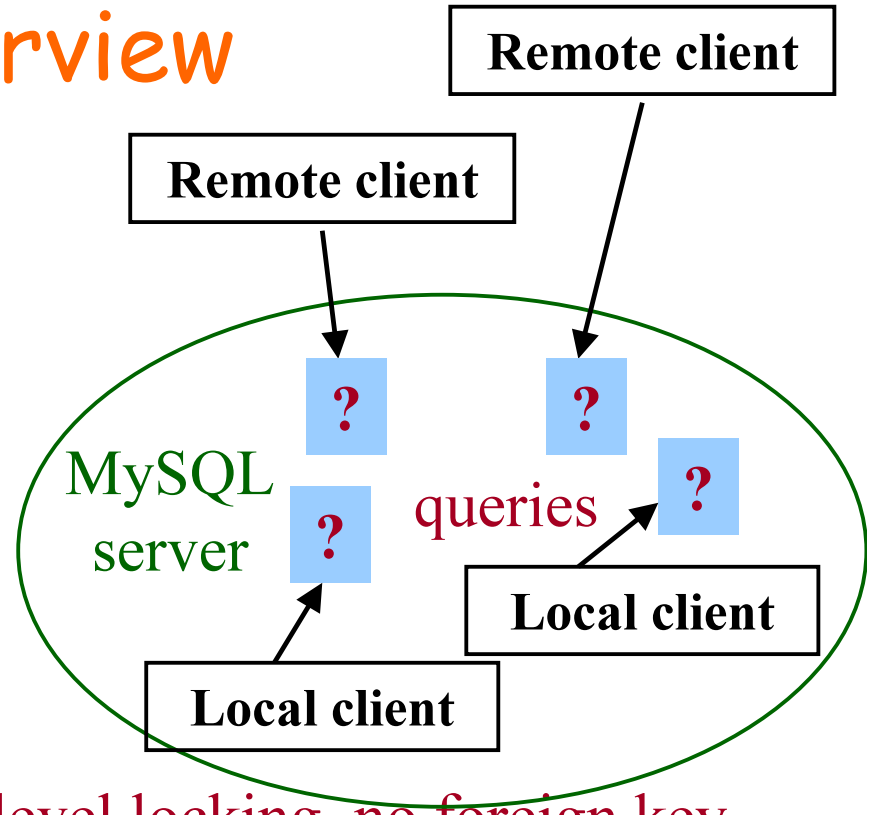

- Tree type of tables:
  - MyISAM: no transaction, table-level locking, no foreign key constraint
  - BDB: transaction, page-level locking, no foreign key constraint
  - InnoDB: transaction, row-level locking, foreign key constraint, hot backup

Tests on InnoDB table (1)

- Foreign key constraint: yes
- Transactional:
  SET AUTOCOMMIT, COMMIT, ROLLBACK
- Transaction model:
  - No-locking multiversion consistent read:
    SELECT…FROM
  - Row-level read lock:
    SELECT…FROM…LOCK IN SHARE MODE
  - Row-level write lock: INSERT INTO…
  - Table lock: LOCK TABLES…
  - Automatic deadlock detection

Tests on InnoDB table(2)

- Insert time: 3.5ms/insert
- Delete time: 3.5ms/delete
- Lookup time
  - Table lfiles size: 10^3, 10^4, 10^5, 10^6 rows
  - lfn:pfn=1:1
    - 0.6-0.8 ms , select 1 out of 1
    - 1-2 ms, select 1 out of 100
    - 8 ms, select 1 out of 1000
  - lfn:pfn=1:10
    - 3-4ms, select 1 out of 1
    - 5ms, select 1 out of 100

Tests on InnoDB tables (3)

- Oracle-like physical and logical structure
- Default page size 16KB
- Preliminary test result:
  - 1.28 < table size(page)/data size(page) < 5

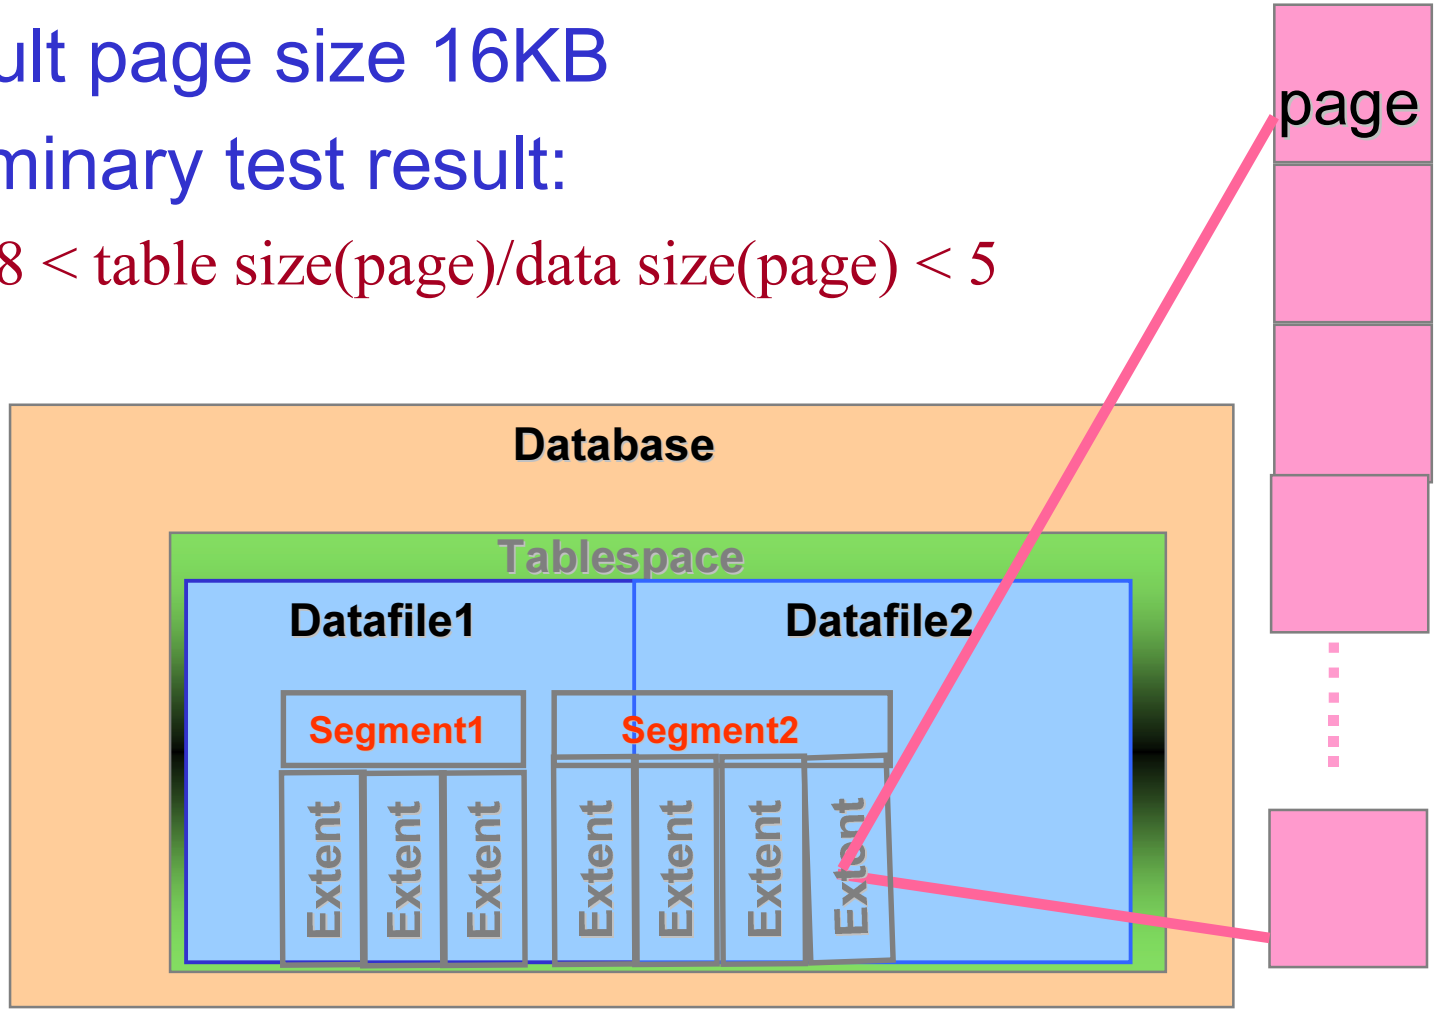

Concurrency tests

- Maximum number of clients:
  - OS limit on file descriptors
  - server memory: each client claims a thread-stack
- Read/write performance with large number of connections
  - Standalone tests planned (by Maria Girone): code ready, 1 server, 10 client nodes
  - Keep track of the performance problem with the CMS production book-keeping system

Stability tests

- Client crash test
  - Automatic rollback of uncommitted transactions
- Server crash test
  - When the server restarts, InnoDB automatically checks logs and performs a rollforward of the database to the present
- More complex tests forseen

Conclusions

- This simple MySQL based catalog prototype
  - Scaling up to a few million catalog entries
  - Performing hundreds of file lookups per second
- Tests with hundres clients sucessful, but very large number of connections may require changes to server OS configuration.
- Concurrency under investigation also in CMS production setup. Good concurrency with InnoDB can be achived but requires control of transaction duration.
- Reliability tests are being performed, will soon be concluded.