

CSE 444: Database Internals


Lectures 27 NewSQL

Slides from Andrew Pavlo
Brown University

References

- **Scalable SQL and NoSQL Data Stores**, Rick Cattell, SIGMOD Record, December 2010 (Vol. 39, No. 4)
- **The end of an Architectural Era: (It's Time for a Complete Rewrite)**, M. Stonebraker et. al. VLDB '07
- Online documentation: H-Store

VOTER BENCHMARK *Japanese "American Idol"*




TRANSACTION

1. Check whether user has already voted.
2. Insert new vote entry.
3. Update vote count for contestant.

3

VOTER BENCHMARK *Japanese "American Idol"*




4

Argument from VLDB'07 paper

- Popular DBMSs based on designs from 70's
- But computer architectures are changing
- And applications have new requirements
- Past 40 years have seen extensions to DBMS design but no major re-design


5

TRADITIONAL DBMS *Measured CPU Cycles*

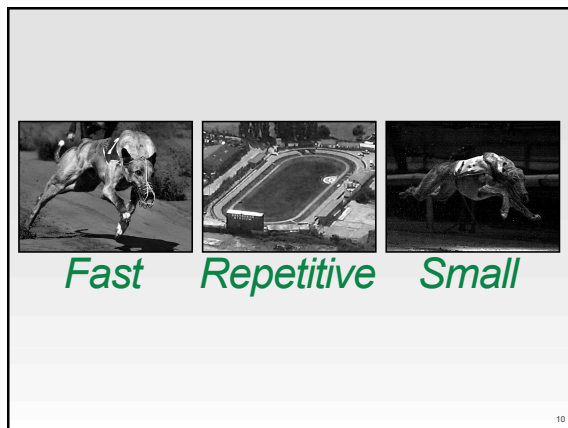


OLTP THROUGH THE LOOKING GLASS,
AND WHAT WE FOUND THERE
SIGMOD, pp. 891-896, 2006

6



CAN YOU SCALE UP WITHOUT GIVING UP TRANSACTIONS?



Optimization

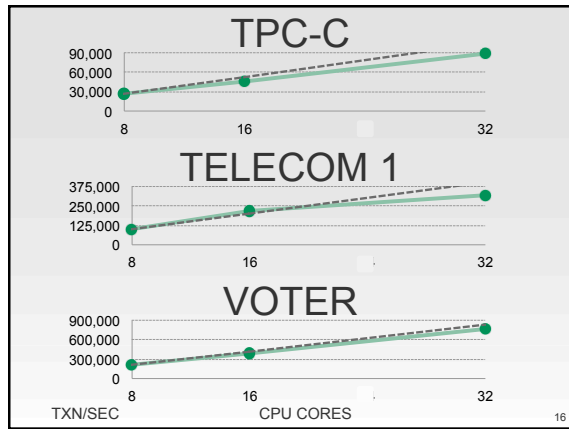

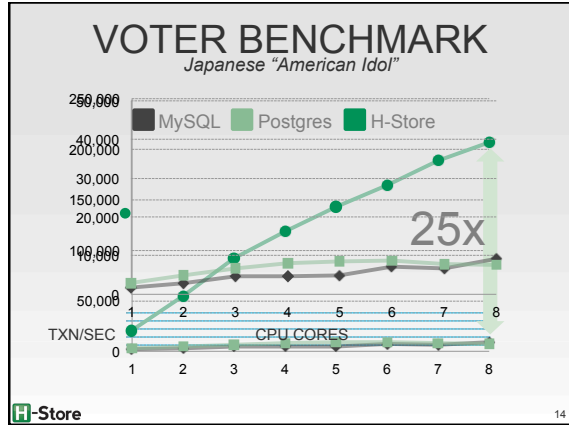
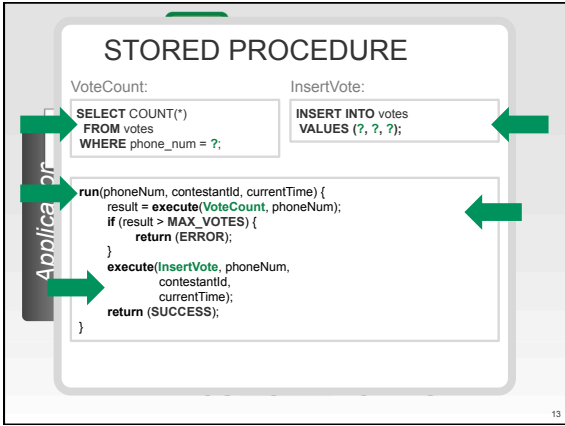
USE A LIGHTWEIGHT SYSTEM *DESIGNED* FOR OLTP TRANSACTIONS.

H-STORE: A HIGH-PERFORMANCE, DISTRIBUTED MAIN MEMORY TRANSACTION PROCESSING SYSTEM
Proc. VLDB Endow., vol. 1, iss. 2, pp. 1492-1499, 2008.

X DISK ORIENTED
MAIN MEMORY STORAGE

i CONCURRENT EXECUTION
SERIAL EXECUTION

/ HEAVYWEIGHT RECOVERY
COMPACT LOGGING



Distributed Transactions

- Discussion based on VLDB'07 paper

