



One Size
Almost
Fits AllTM

Application






Disk-Oriented Overhead

Measured CPU Cycles



OLTP THROUGH THE LOOKING GLASS,
AND WHAT WE FOUND THERE
SIGMOD, pp. 981-992, 2008.





H-Store



**DISK ORIENTED
MAIN MEMORY STORAGE**



**CONCURRENT EXECUTION
SERIAL EXECUTION**




**HEAVYWEIGHT RECOVERY
COMPACT LOGGING**



H-STORE: A HIGH-PERFORMANCE, DISTRIBUTED
MAIN MEMORY TRANSACTION PROCESSING SYSTEM
Proc. VLDB Endow., vol. 1, iss. 2, pp. 1496-1499, 2008.

STORED PROCEDURE

VoteCount:






```
SELECT COUNT(*)  
FROM votes  
WHERE phone_num = ?;
```

InsertVote:



```
INSERT INTO votes  
VALUES (?, ?, ?);
```




```
run(phoneNum, contestantId, currentTime) {  
    result = execute(VoteCount, phoneNum);  
    if (result > MAX_VOTES) {  
        return (ERROR);  
    }  
    execute(InsertVote, phoneNum,  
           contestantId,  
           currentTime);  
    return (SUCCESS);  
}
```


EXECUTION ENGINES

TS

YCSB

50% Reads / 50% Writes






Anti- Caching



ANTI-CACHING: A NEW APPROACH TO SWAPPING IN
MAIN MEMORY OLTP DATABASE SYSTEMS

VLDB 2014.


Application



Primary Storage




Anti-Cache



YCSB (2x Memory)

50% Reads / 50% Writes



Anti-Caching

- Overhead Reduction.
- In-Memory Compression.
- Vertical Partitioning.



Next-Generation CPUs

- Concurrency control
evaluation on +1000 cores
simulated architecture.
- Multiple cores per partition.





H-Store



S-Store

S-Store



- Integrate distributed stream processing operators directly inside of an OLTP system.



Application



Primary Storage



Stream Operator

Stream Operator

Stream Operator

Stream Operator

Output

Continuous Queries

Non-Blocking Analytics

- Asynchronous Snapshots.
- In-Memory Column Store.
- Query Decomposition.



Transactions on Streams

- ACID properties for streams.
- Open Questions:
 - Transaction boundaries?
 - Interaction of streams & tables?





H-Store



S-Store



N-Store

N-Store

- Transaction processing on non-volatile memory.
 - Memristors / ReRam
 - Phase Change Memory
 - Spintronics



ISTC
BIG DATA



Non-Volatile Memory


- NVM Emulator from Intel.
- Tunable Latencies.
- Accessed as a filesystem.



ISTC
BIG DATA



4 Socket Platform with 3 DIMMs/channel



■ NVM ■ DRAM



Non-Volatile Memory

- Stage 1: H-Store + Anti-Cache
- Stage 2: H-Store + mmap()
- Stage 3: N-Store



ISTC
BIG DATA



Escape From Planet **Stonebraker**

(i.e., Andy Needs to Get Tenure)


Beyond the 'Stores

- Non-Partitionable Workloads.
- The Poor Man's Spanner.
- Scientific Databases.



- **Self-Organization**
- **MVCC**

Conclusion





**Andy
Pavlo**



**Justin
DeBrabant**



**John
Meehan**



**"The Thrill"
Stonebraker**



**Stan
Zdonik**



**Xiangyao
Yu**



**George
Bezerra**



**Sam
Madden**



**Nesime
Tabul**



**Ugur
Cetintemel**



**Tim
Kraska**



**Saman
Amarasinghe**