

Performance of Transactions in RAMCloud

SEDCL Retreat

May, 2015

Collin Lee and **Seo Jin Park**
Stanford University



Overview

- **Review of transactions on RAMCloud**
- **Performance micro-benchmark**
 - Latency
 - Scalability
 - OCC performance in Contention
- **TPC-C benchmark**
 - Latency of New-Order transaction
 - Throughput scalability of New-Order transaction

Transactions Goals

What are we trying to build?

- **Multi-object atomic updates**
- **Tolerate client failures**
- **Performance**
 - Low-latency
 - Large scale: 1M+ clients
- **Simple programmer interface**
- **Non-goals and assumption:**
 - No long running transactions
 - Small commit sets: 100 objects or less

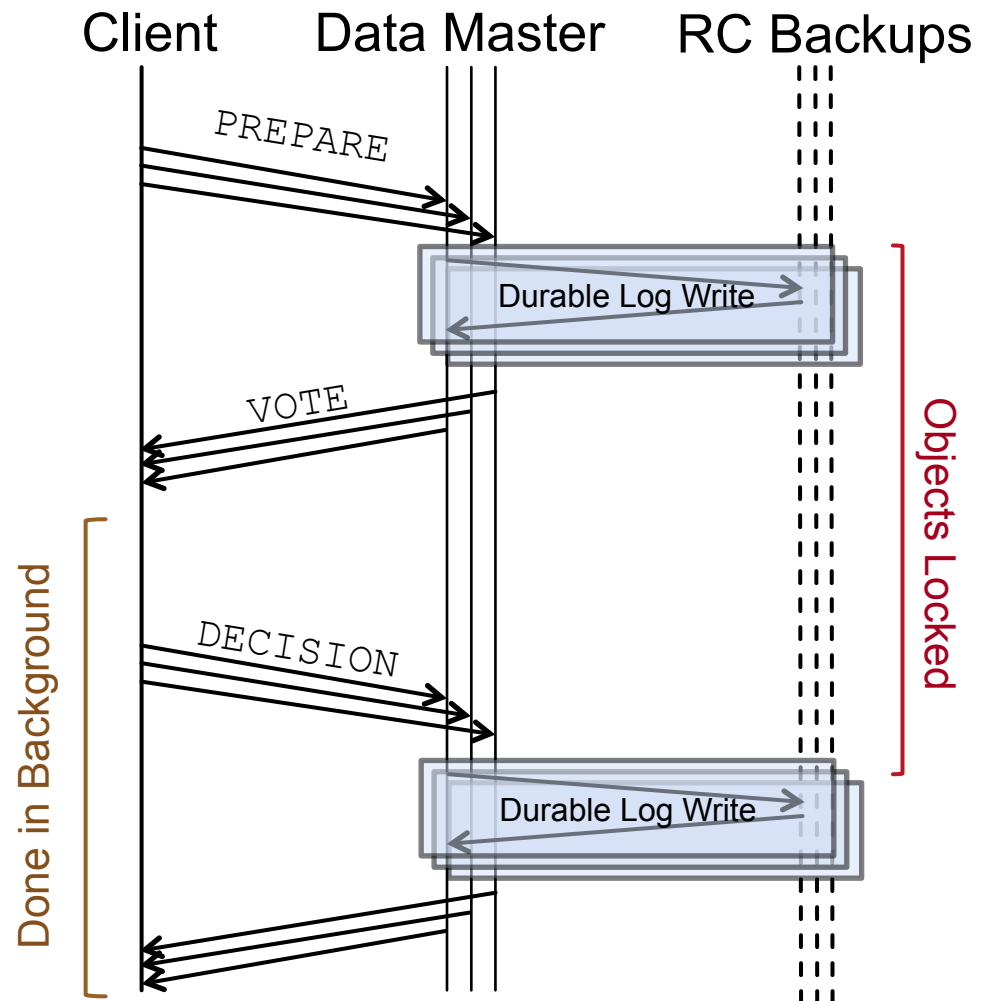
Transaction Client API

```
class Transaction {  
    read(tableId, key) => blob  
    write(tableId, key, blob)  
    delete(tableId, key)  
    commit() => COMMIT or ABORT  
}
```

- **Optimistic concurrency control**

Transaction Commit

- **Client-driven 2PC**
- **RPCs:**
 - `PREPARE () => VOTE`
 - `DECISION ()`
- **Client blocked time: 1RTT + 1D**
- **Decisions sent in the background**



Performance of Transactions

- **Measured performance on Infiniband with (100B object) and 3-way replication.**

1. Transaction Commit Latency

- Single server < 16 μ s; 5 servers: ~23 μ s


2. Transaction Throughput

- Single server: 67k txns/sec; 5 servers: 40k txns/sec

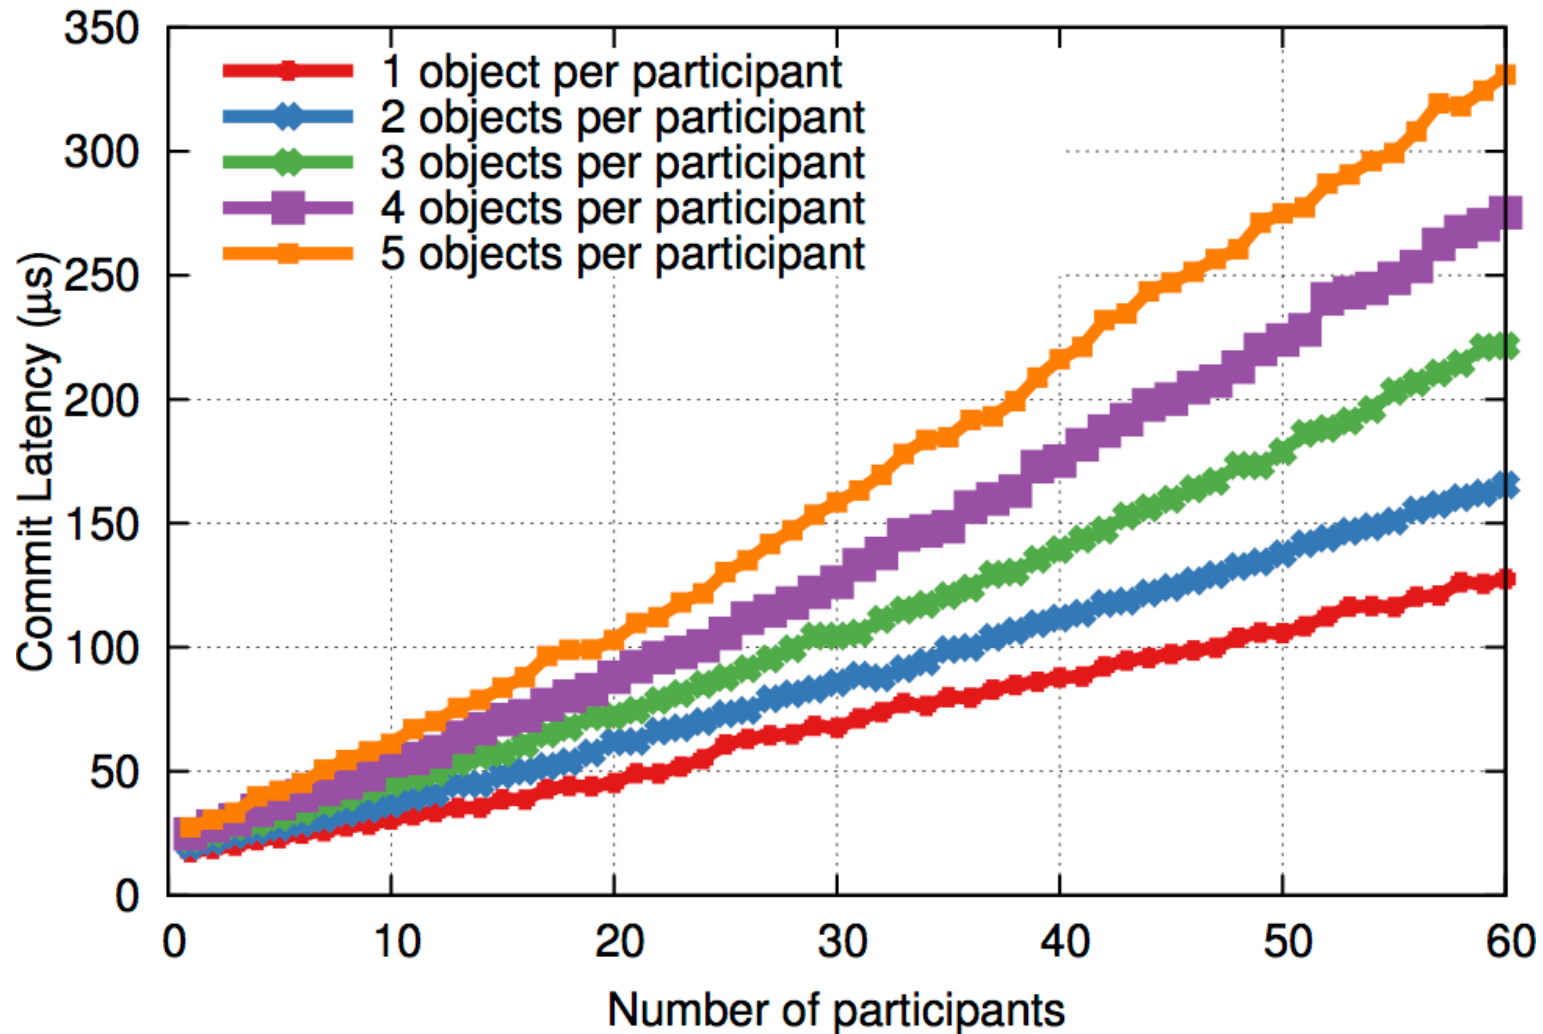
3. TPC-C benchmark

Transaction Commit Latency

- Single server < $16\mu\text{s}$; 5 servers: $\sim 23\mu\text{s}$




Transaction Commit Latency(2)



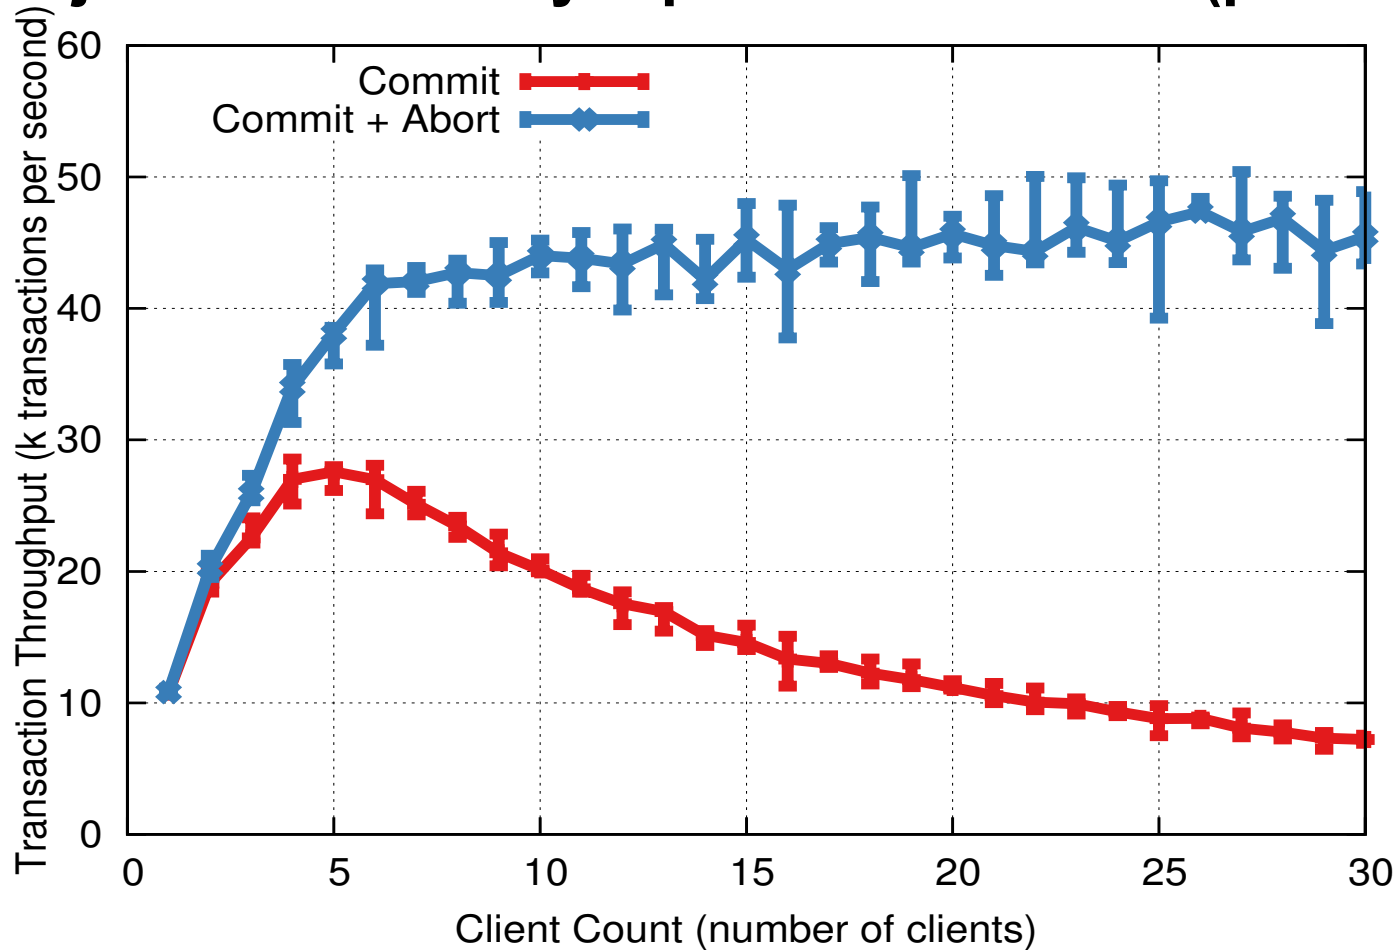
Transaction Throughput

- Single server: 67k txns/sec; 5 servers: 40k txns/sec



TX Throughput with contention


- 5 Objects selected by Zipfian distribution ($\rho=0.99$)



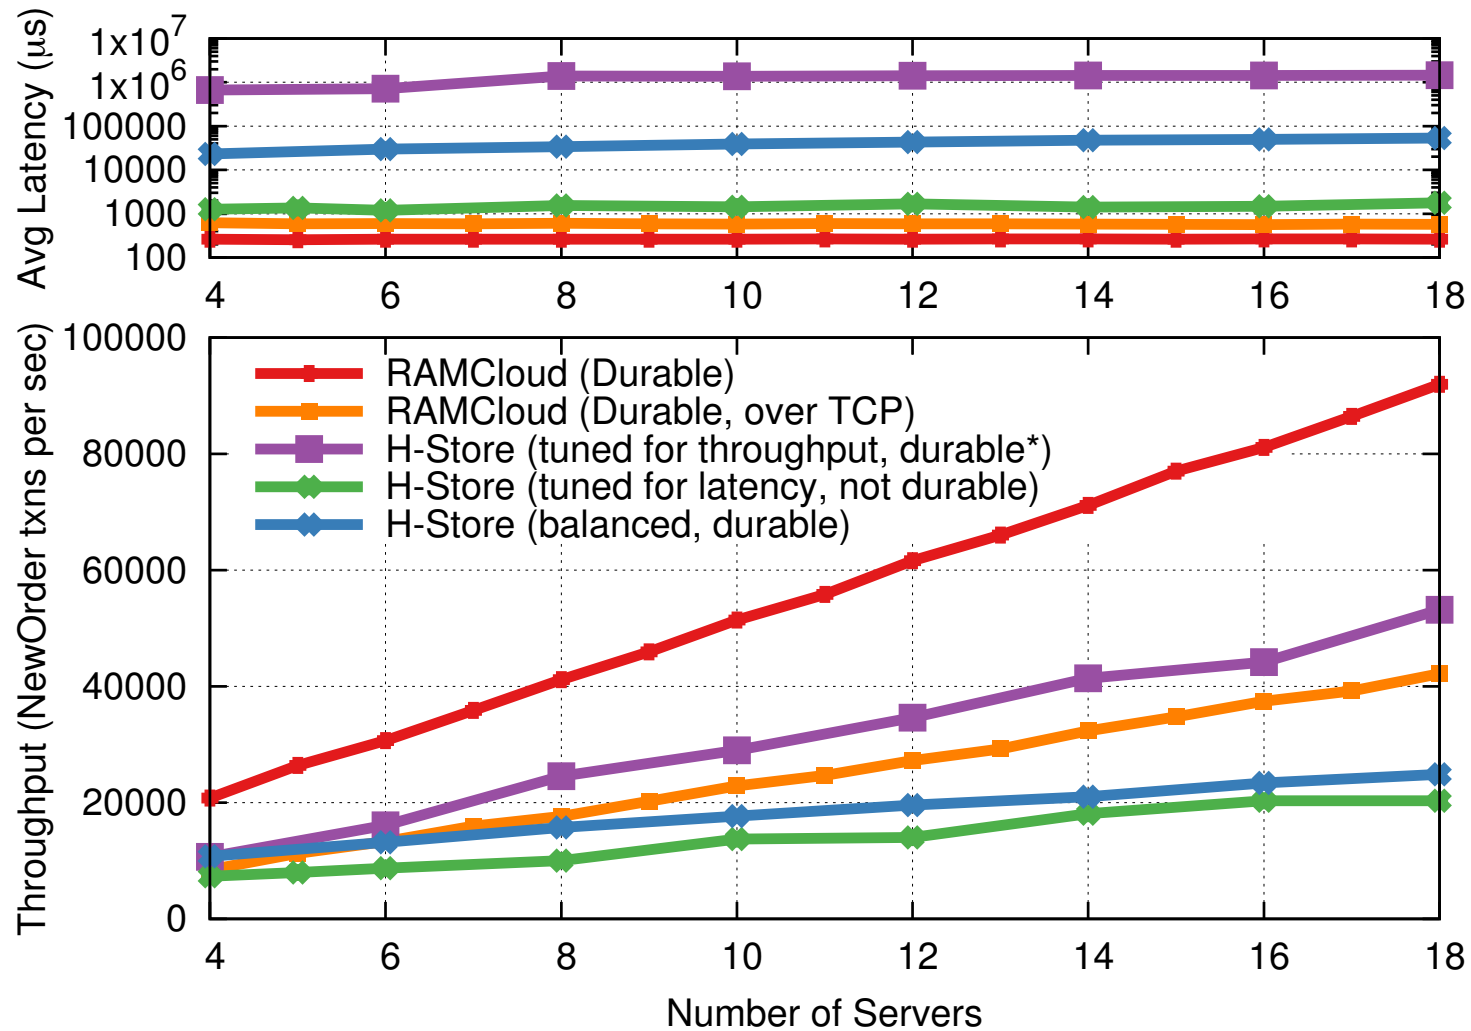
TPC-C Benchmark

- **TPC-C is an industry standard OLTP workload.**
 - Compares with other state-of-art in-memory DBMS.
- **Modified TPC-C for RAMCloud benchmark**
 - No client wait time.
 - No 30-days space requirement.
- **New-Order transaction performs on average 23 selections, 11 updates, and 12 insertions.**
- **Latency is measured from end to end.**

TPC-C NewOrder Latency



TPC-C NewOrder Throughput



Conclusion

- **The low-latency transaction on RAMCloud actually showed 1RTT + 1D latency for small commit set.**
- **Scales well, although as commit set increases, latency increases.**
- **OCC causes high abort rate in contention.**
- **Outperforms a state-of-art transaction system using a popular TPC-C benchmark.**

Questions
