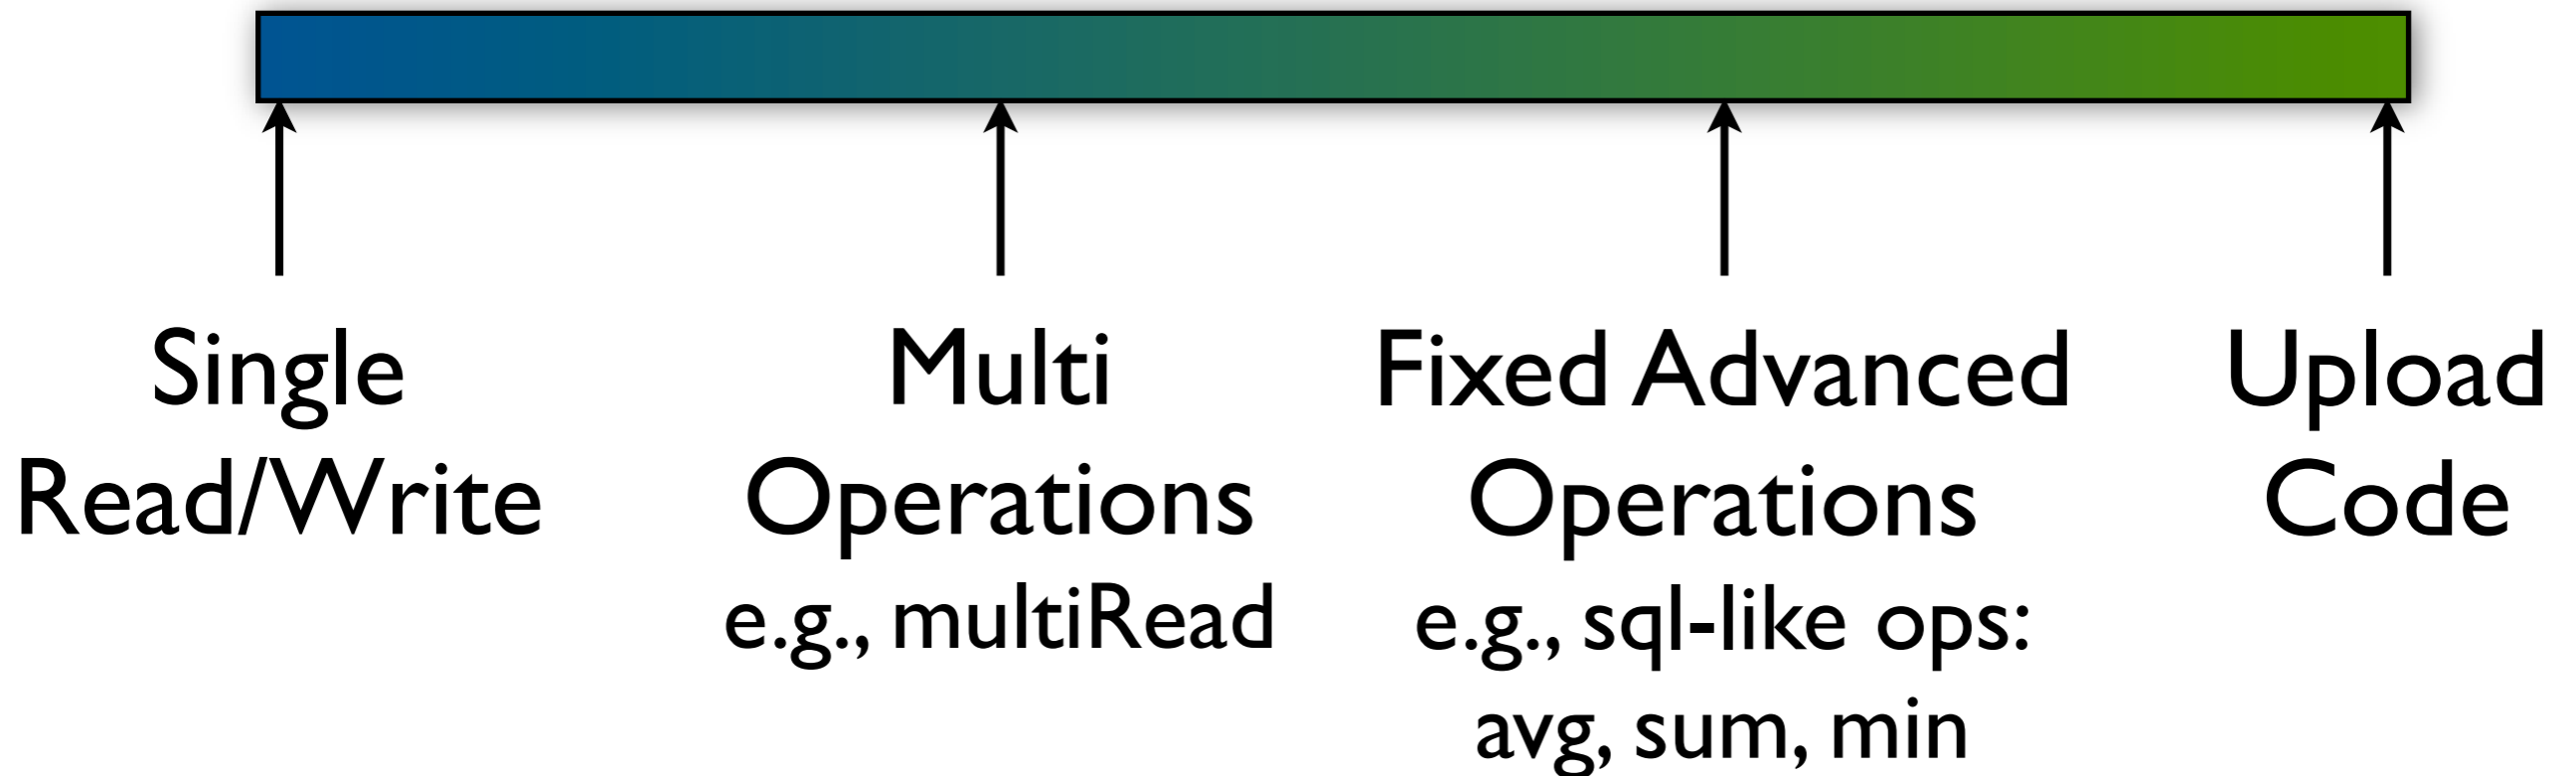


RAMCloud: multiRead

Ankita Kejriwal

The Big Question

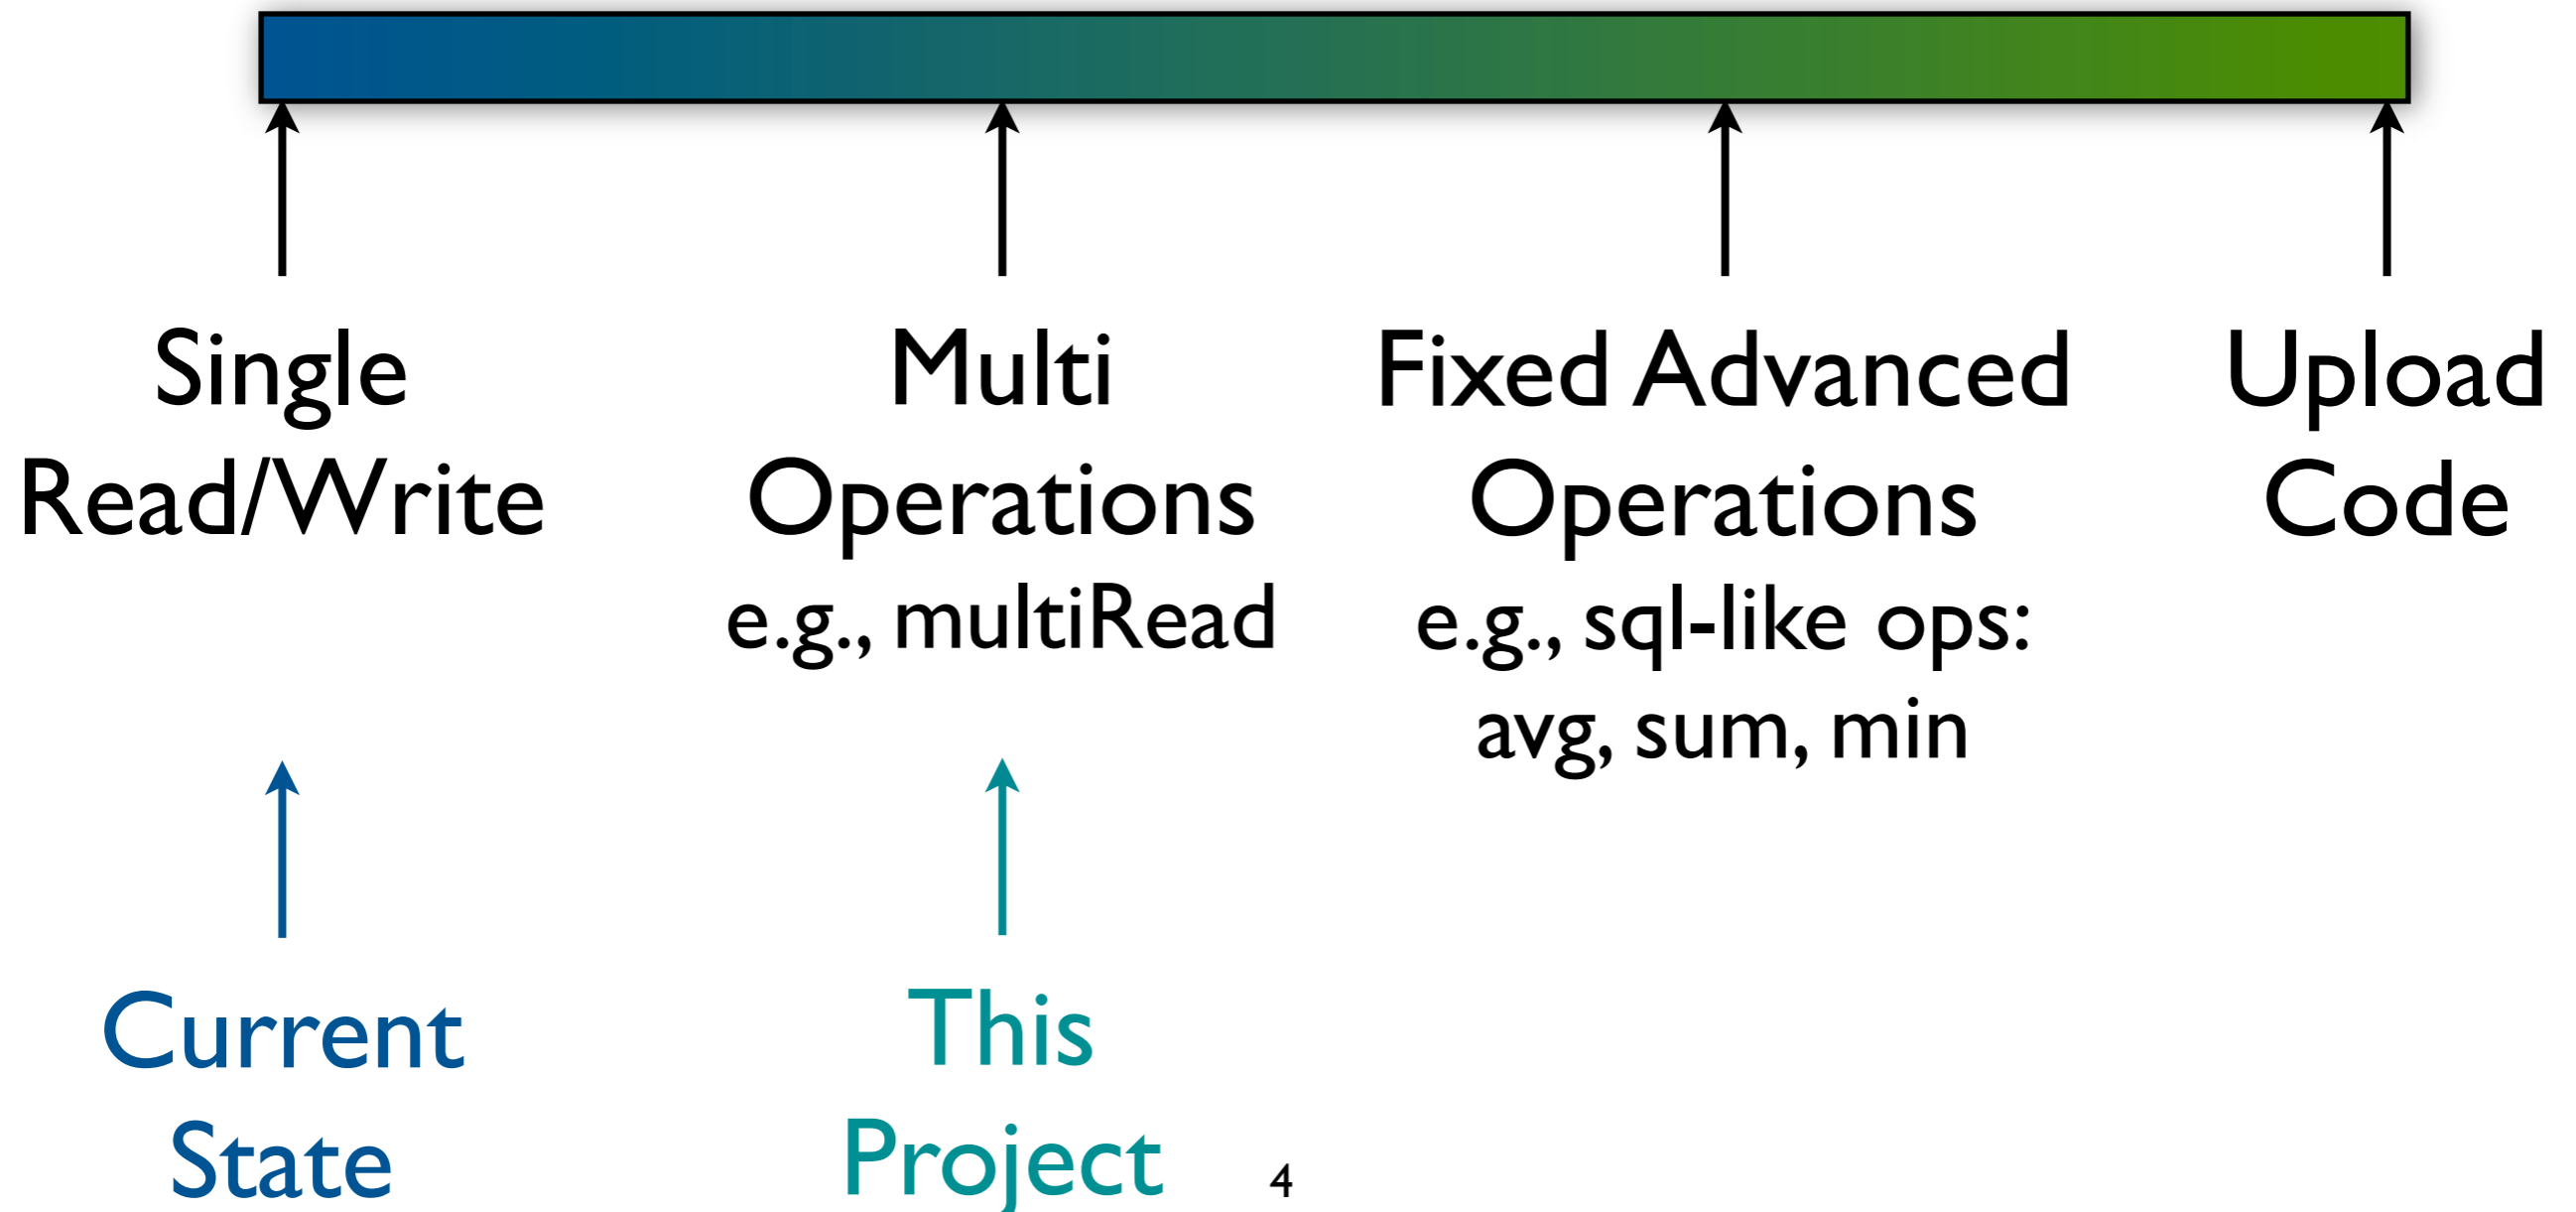
What is the ideal distribution of work between the client and the server?



Design Considerations

- Low level operations (Read-Write):
 - Arbitrary computations are possible
 - Simplifies server implementation
- Higher level operations:
 - Fewer RPCs exchanged => lower latency
 - Less network bandwidth
- Issues with higher level operations:
 - Need data locality
 - Potential security implications (with uploading code)

The Big Question



Motivation

- Batch multiple requests to a single master in one multiRead RPC
- Parallelize requests to multiple masters

multiRead API




Interesting Issues in Implementation

- Error handling
 - Individual read requests can fail
- RPC Packet Overflow
 - Possible Solution: Client sends another request for remaining objects


Benchmarking

- multiRead and read
- one table, one master
- multiple 100B objects
- random object Ids

Latency - Multiple Objects on 1 Master



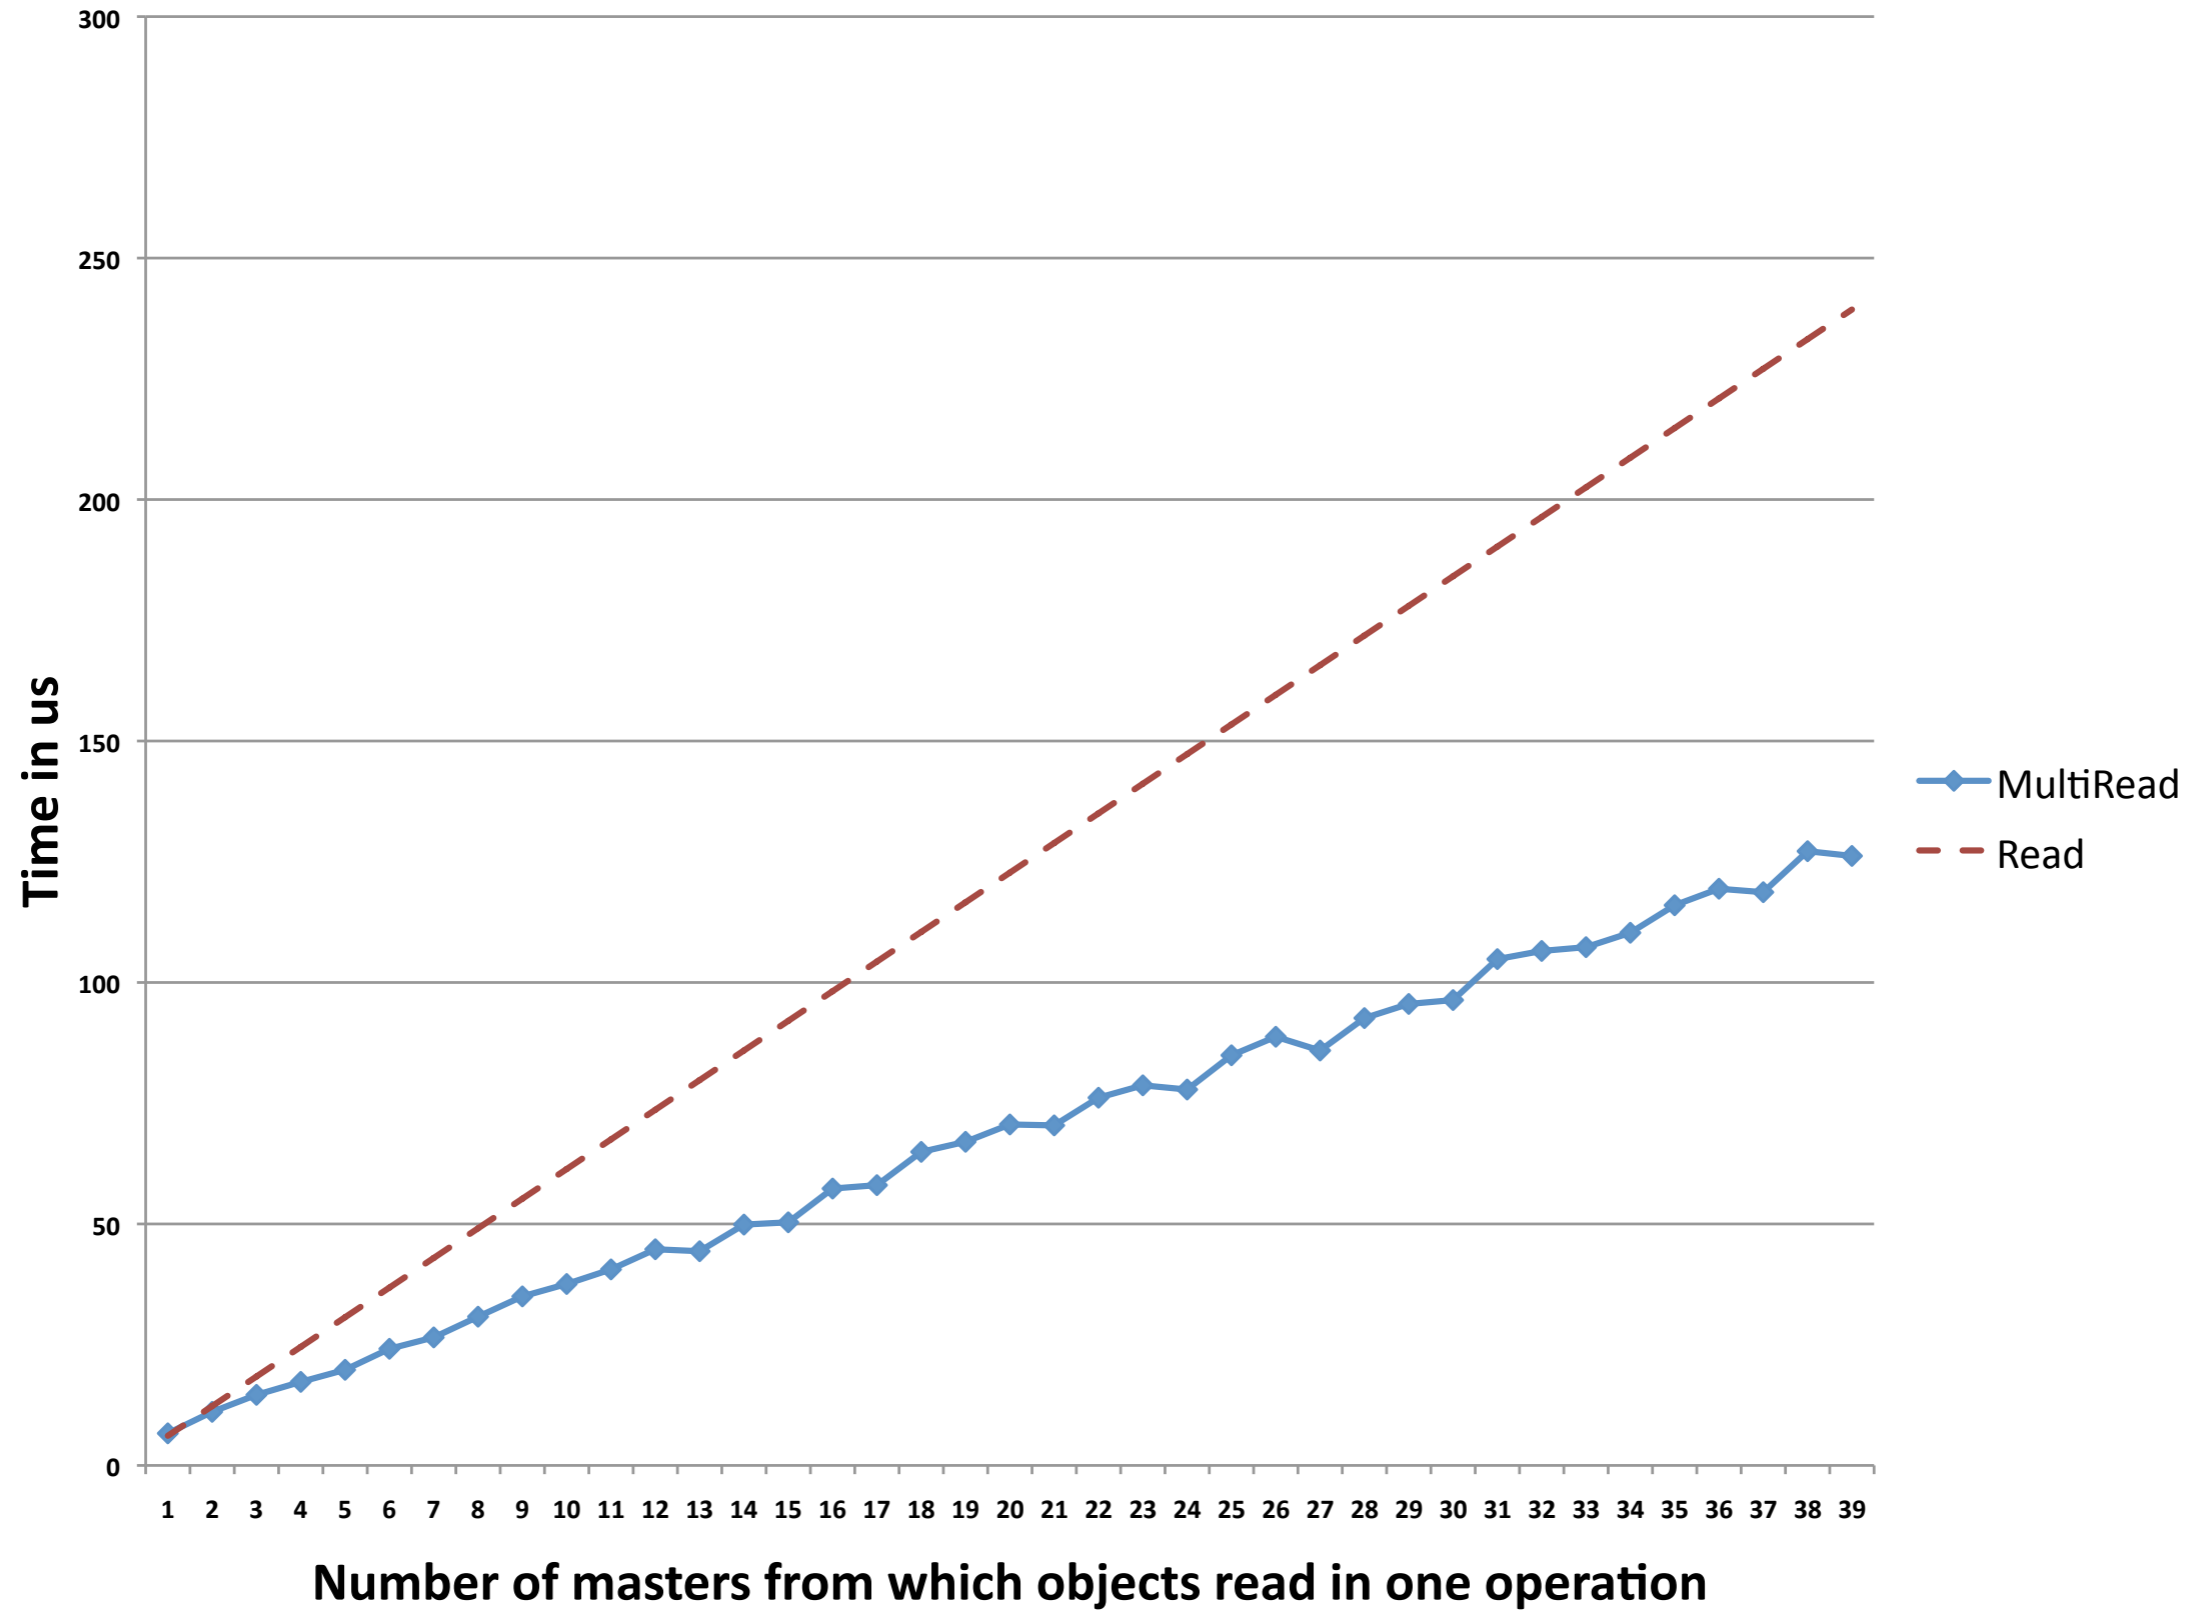
Throughput - Multiple Objects on 1 Master



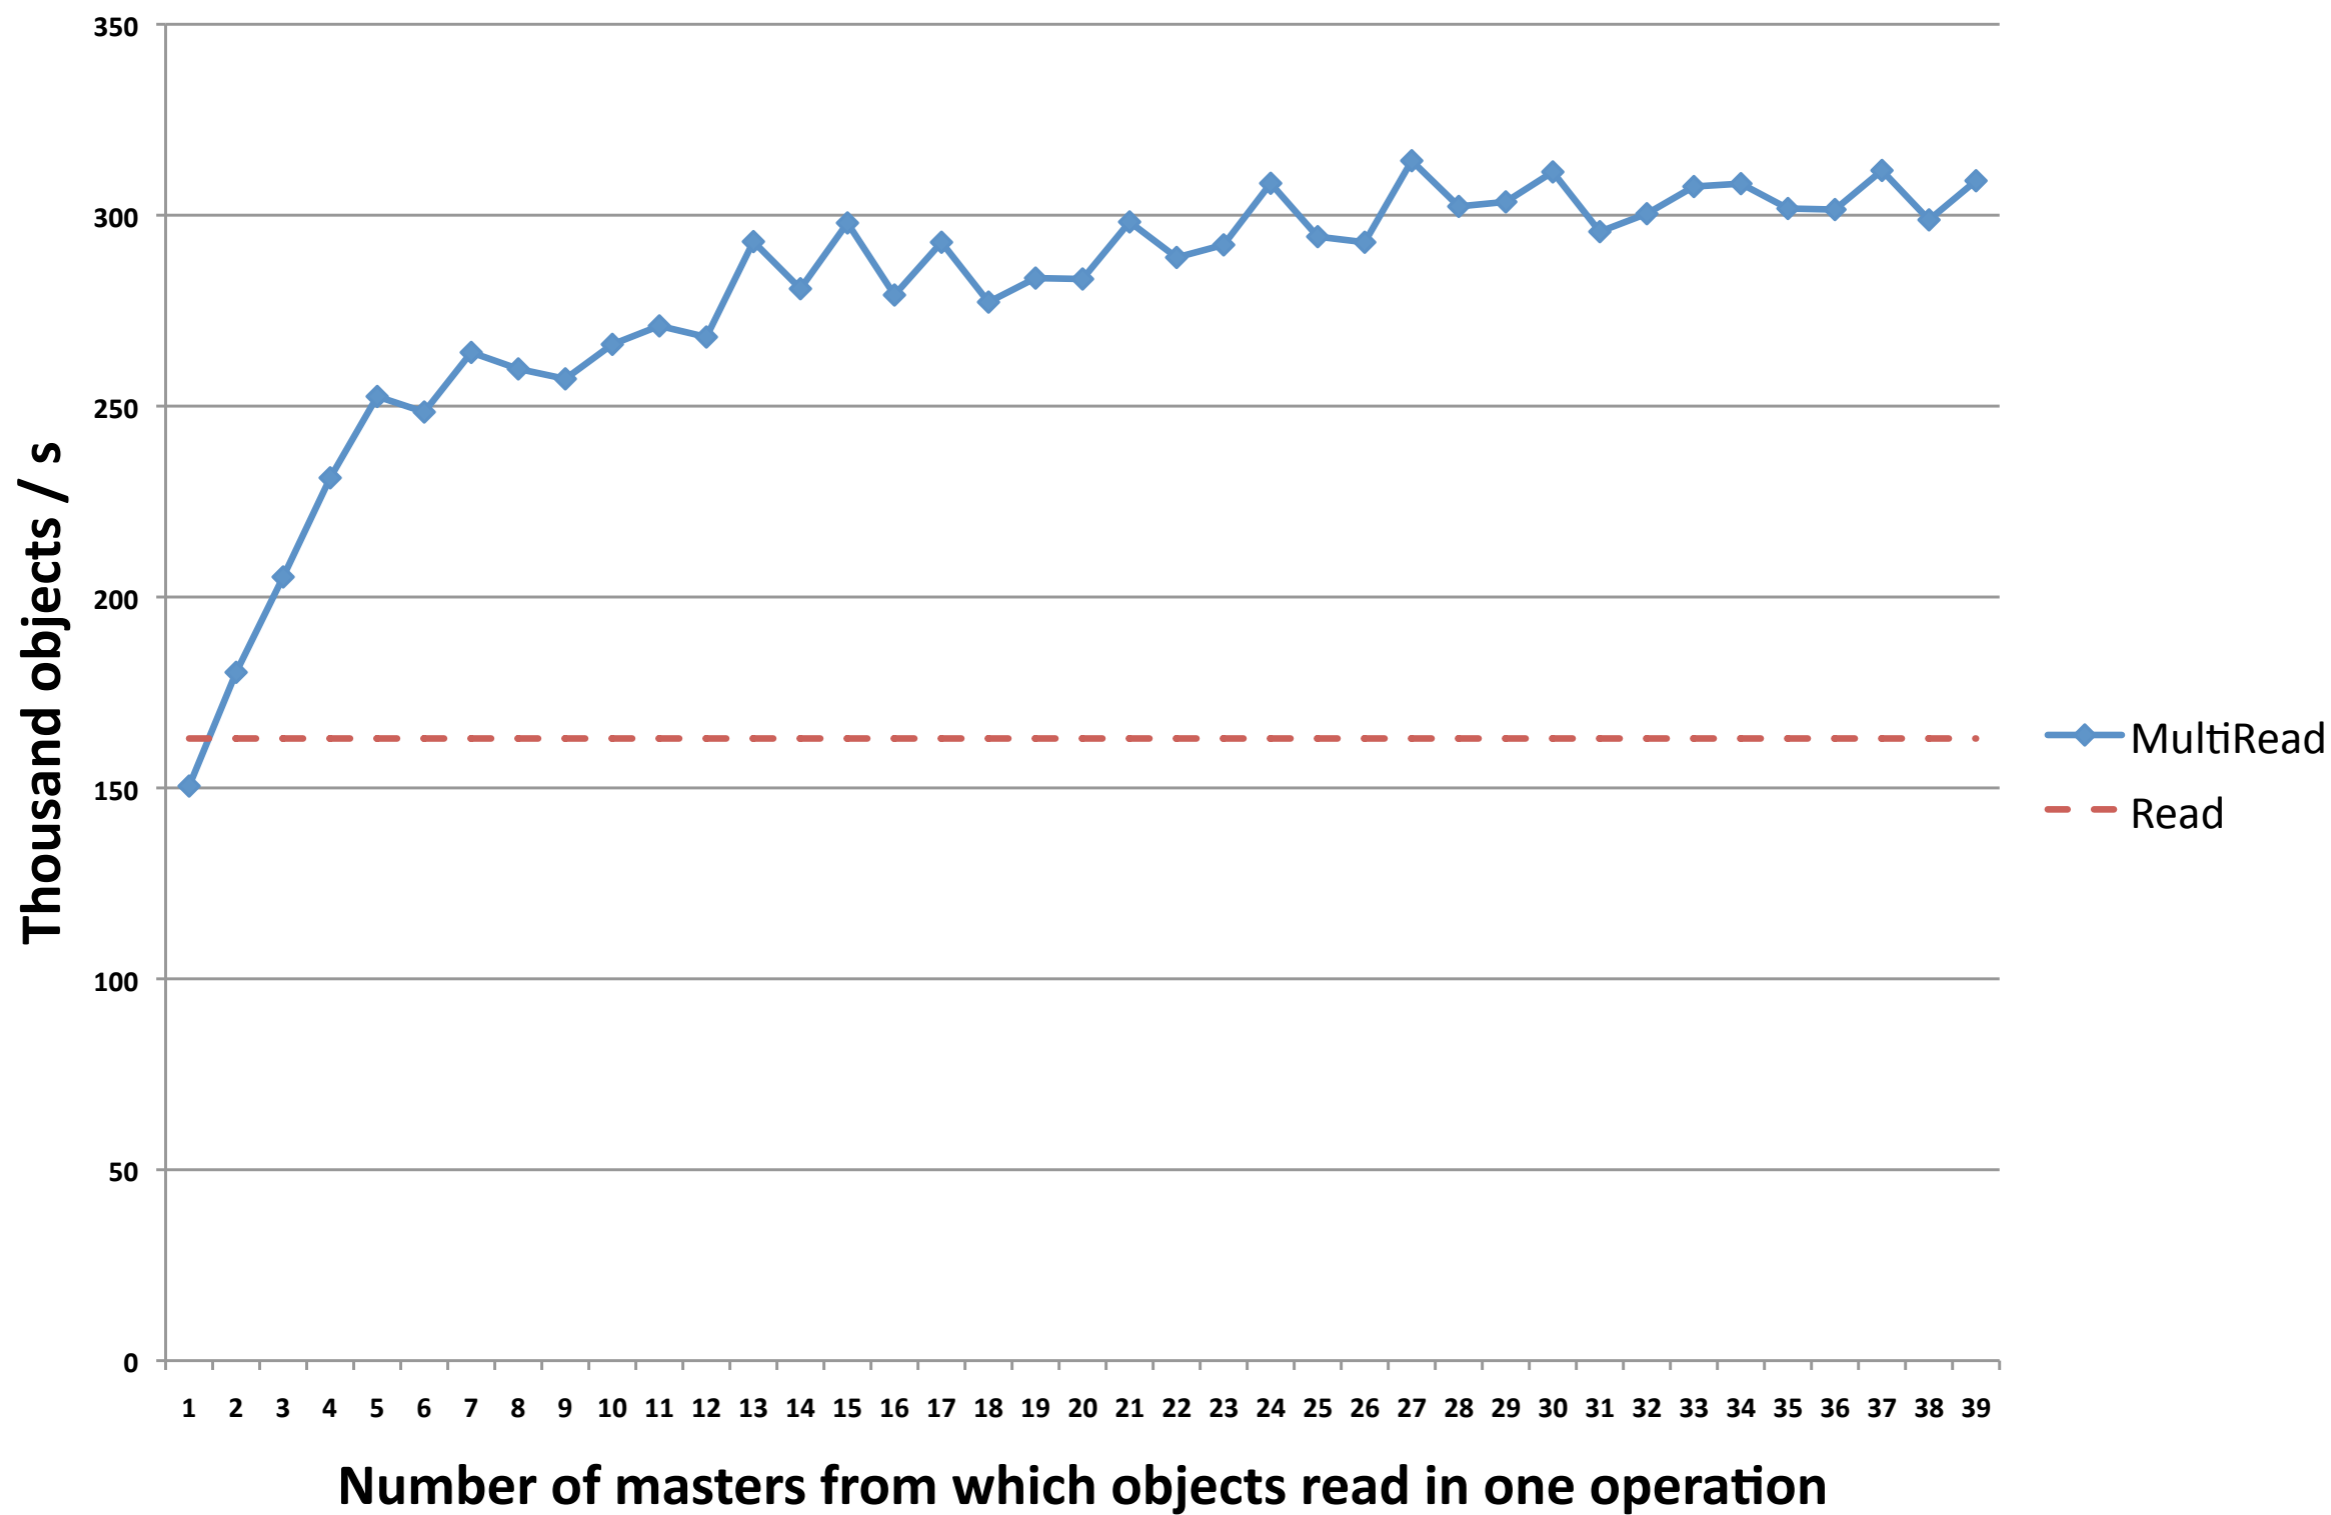
Benchmarking

- multiRead and read
- multiple tables, multiple masters
- one 100B object on each master

Latency - 1 Object on each Master



Throughput - 1 Object on each Master



Summary & Future Work

- Summary:
 - Batched requests to 1 server: up to 5.2x speedup over read
 - Parallel requests to multiple servers: up to 1.9x speedup over read
- Future Work:
 - MultiRead - Understanding what limits performance and whether it can be improved; RPC Overflow; More benchmarks; Optimizations
 - Design, Implementation and Benchmarking of MultiWrite and other advanced operations

Questions?