Taming Latency In Data Center Applications

Ph.D. Defense of Dissertation

Mohan Kumar Advisor: Taesoo Kim

Motivation: Importance of Latency

Speed Matters for Google Web Search

Jake Brutlag Google, Inc.

The Cost of Latency

DOI:10.1145/3015146

Microsecond-scale I/O means tension

Latency Critical In Data Center Applications

retail giant found that every 100ms of latency cost them 1' same time, a study by Tabb Group revealed that a broker (\$4 million in revenues per millisecond if its electronic tradi 5ms behind the competition.

Although low-latency connectivity has become more com Amazon, Tabb and others began this conversation about I ago, low latency is still a critical element of business succe

Attack of the Killer Microseconds

Data Center Applications



Contemporary Data Center Characteristics

- Optimized network microsecond round-trip time
- Moving from 10/25 Gbps to 100/200 Gbps network
- Software running in servers induce high latency:
 - \succ 66% of the inter-rack latency [1]
 - > 81% of the intra-rack latency [1]

Data Center Applications - Server Latency



Thesis Statement

System abstractions and optimizations are needed at different levels of the software stack, from the software services running in the user space and the kernel to the software running on SmartNICs, to reduce the latency and improve the throughput of current data-center applications.

Data Center Applications - Server Latency



Taming Application Latency- Thesis

- Xps Extensible Protocol Stack:
 - > Abstraction in kernel and user-space protocol stacks, and SmartNICs
 - \succ Reduces Redis latency by up to 73.3%
- LATR Lazy Translation Coherence:
 - > Kernel mechanism for free operations, page migration and swapping
 - \blacktriangleright Reduces Apache latency by up to 26.1%

Taming Application Latency - Thesis

- Dyad Untangling Logically-Coupled Consensus:
 - > Abstraction in SmartNIC for consensus
 - \succ Reduces timestamp server latency by up to 79%

Dyad: Untangling Logically-Coupled Consensus

Motivation - Consensus Algorithms

OPINION

The Amazon Outage in Perspective Inevitable, ^-

FACT #1: THE EFFECTS OF DOWNTIME ARE EXPENSIVE.

The most recent Ama; Instead, CIO.com colu



ecially with The average total cost of unplanned application downtime per lications nents. - 61 OF HILLOW HO 60 F HILLOW

Failures are inevitable and expensive

An endless stream of twe last week's Amazon Web as an indictment of publiwork at other cloud prov shortcomings. Still others have to be sure to hamm negotiations, just to ensu

per hour.



s were used

363 01

Consensus Algorithms

- Consensus Algorithms:
 - > Provides high availability by state machine replication
 - ➤ Keeps data consistent linearizable
 - > Consensus algorithms:
 - Multi-Paxos/Viewstamp Replication (VR)
 - Raft and Zookeeper Atomic Broadcast (ZAB)

Consensus Algorithms - Applications



Distributed Services

- ➤ Timestamp Servers
- ► Key-value stores
- > Database
- Lock managers





Dyad: Untangling Logically-Coupled Consensus

- Background
- Overview
- Design and Evaluation
- Conclusion

Consensus – VR Data Operation



Consensus – ZAB or Raft Data Operation



Replicas in a Data Center





[1] Understanding PCIe performance for end host, SIGCOMM'18

Consensus – VR Data Operation







Consensus Latency - Increasing Replicas



Consensus Operations

- Data Operation:
 - \succ Critical path for handling a client request
- Control Operations:
 - Recovery application recovery after failure
 - ➤ View Change new replicas joining/leaving the group, new leader
 - ➤ Heartbeats health status messages exchanged across replicas

Cost of Consensus - Summary

- Every client request has high consensus overhead
- Consensus algorithms share resources with application
- Consensus overhead increases with increasing replicas

Consensus - Existing Research

• Network approaches:

>

>

Rely on Network Guarantees

• Hardware approach:

Logically Coupled Consensus

sts

es

Dyad: Untangling Logically-Coupled Consensus

- Background
- Overview
- Design and Evaluation
- Conclusion

Logically-Coupled Consensus



Dyad: Untangling Logically-Coupled Consensus



Dyad: Untangling Logically-Coupled Consensus



Dyad: Classifying Consensus Operations

- Data Operation SmartNIC:
 - > Critical path for handling a client request
- Control Operations Host:
 - Recovery application recovery after failure
 - ➤ View Change new replicas joining/leaving the group, new leader
 - ➤ Heartbeats health status messages exchanged across replicas

Dyad: Untangling Logically-Coupled Consensus

- Background
- Overview
- Design and Evaluation
- Conclusion

Dyad: Data Operations





Dyad – Direct Cost



Dyad – Indirect Cost



Direct and indirect cost reduced by Dyad



Dyad: SmartNIC Primitives

- Hardware Filtering:
 - Specify packet format in domain-specific language (P4)
 - \succ Filter messages based on the header and payload
 - \succ Filters are applied to messages coming from the network and the host
Dyad: SmartNIC Primitives

- Packet Processing:
 - > Filtered messages invoke request/consensus/response handlers
 - ➤ Handlers drop/forward/modify a packet
 - ➢ Generate new packets















Dyad: Timestamp Server with 5 replicas

- VR - VR-batching - Dyad-leader



Packets Per Second (PPS * 1K)

> Reduce latency by up to 76%, Improves throughput by 5.8x







Context switch





Dyad: Ordering on Replica SmartNIC

- Ordering and Logging:
 - \succ Logs ordered by the sequence number in prepare message
 - Prepare message are processed and dropped on the SmartNIC
- Ordered Execution:
 - \succ Commit messages forwarded to the host processor
 - > The request is appended to the commit message by SmartNIC





Dyad: Timestamp Server with 5 replicas

- VR - VR-batching - Dyad-leader - Dyad-all



Packets Per Second (PPS * 1K)

 \succ Reduce latency by 30 µs

Dyad: Consensus Latency

System	Consensus latency (µs)	% reduction
VR	350	N/A
VR-batching	409	N/A
Dyad-Leader	48	86%
Dyad-All	17	95%
		1 •

Timestamp server - 5 replicas

Dyad: CPU Usage Timestamp Server



Packets Per Second (PPS * 1K)

 \succ Reduce CPU usage by up to 70% on the leader

Dyad: Control Operations



Dyad: Application Failures



[1] Simple Testing Can Prevent Most Critical Failures, OSDI'14

Dyad: Detecting Application Failures



Dyad: Detecting Application Failures

- Measure host RTT for each request
- Computed weighted average of host RTTs
- Detect failure response not within host RTT threshold

Application Recovery - VR



Dyad: Application Recovery

- Recovery using logs on SmartNIC
- Two stage recovery:
 - \succ Recover logs from the SmartNIC
 - \succ Recover remaining logs from other replicas



Dyad: SmartNIC Failure



Dyad: System Failure



[1] Simple Testing Can Prevent Most Critical Failures, OSDI'14

Dyad: System Recovery

- SmartNIC Failure:
 - > Detected on the host using heartbeat/client messages
 - Existing VR recovery: fetch remaining logs from other replicas
- System Failure:
 - \succ Existing VR recovery: fetch logs from other replicas
 - > Dyad supports logging to disk from host (Raft)

Dyad: Reliable Connection

- Dyad Supports Raft:
 - \succ Using TCP connection to replicas
 - > TCP stack specifically decode Raft headers and payload
 - ➤ Host application logs client commands to disk for persistence

Dyad: Raft Latency



Dyad: Ease of Use

- Memcached:
 - \succ Enable consensus for Memcached
 - \sim 100 lines of code for data operations on replica
 - > Evaluate impact on latency and throughput



Dyad: Memcached Latency



Dyad: Untangling Logically-Coupled Consensus

- Motivation
- Background
- Overview
- Design and Evaluation
- Conclusion

Dyad: Conclusion

- SmartNIC abstraction for consensus
- Data operations performed on the SmartNIC
- Control operations performed on the Host
- Enables consensus as a service on SmartNICs

Thesis: Conclusion

- Xps Extensible Protocol Stack:
 - > Abstraction in kernel, user space, and SmartNIC
- Latr lazy TLB shootdown:
 - ➤ Kernel mechanism for TLB shootdown

System abstractions and optimizations are needed at different levels of the software stack to reduce the latency and improve the throughput of current data-center applications.

Thank you!

Backup Slides
Arrakis

Redis comparison with Arrakis



Latr - Apache





User-Space Stacks

User Space: Protocol processing

Systems	Latency (µs)	Mitigation
mTCP	~ 23	Batching
IX	~12	Batching
Arrakis	~2.6 - 6.3	None



Packets Per Second (PPS * 1K)

Context Switch

VR - Leader Context Switch



Packets Per Second (PPS)

Dyad - Parallelism

Dyad: Application Parallelism

- Without SmartNIC:
 - Sequence numbers are available in prepareok message
 - \succ Multi-thread execution by using the sequence number
- Dyad:
 - Request are ordered without containing the sequence number
 - SmartNIC appends the sequence number to the client request

Dyad: Parallelism Timestamp Server

— Dyad-leader **—** Dyad-all **—** Dyad-parallel



Packets Per Second (PPS * 1K)

 \succ Improves throughput by up to 2.1x

Reading Logs



Log Size (MB)

> Log read throughput ~256 MB with 16 threads

Direct Cost Formula

Cost of Consensus - Direct and Indirect



VR Recovery Data Transfer

Application Recovery - VR data transfer

Replicas	Log Size (MB)	Data transferred (MB)
3	100	200
5	100	400
7	100	600

False Positives RTT

Dyad: False Positives with Timestamp Server



RTT Threshold (* RTT)

 \succ RTT = ~96 µs

SmartNIC - Netronome

SmartNIC: Memory Hierarchy and Latency



94

Recovery Example

Dyad - Recovery Phase1



Dyad - Recovery Phase2



Raft - Logging to Disk

Dyad: Raft Latency with disk logging



Dyad - Future Work

Dyad: Future Work

- Logging to disk from SmartNIC:
 - \succ Possible with NVMe over fabric
 - > Possible over PCIe? ARM, FPGA, or NPU
- Optimize request handling:
 - Sending parsed requests to host