

# Log-Based CRDT for Edge Applications


**Nazmus Saquib**   Chandra Krintz   Rich Wolski

Department of Computer Science  
University of California, Santa Barbara


September 27, 2022



# Multi-Tiered Cloud Architecture




# Multi-Tiered Cloud Architecture



heterogeneity  
unstable network connectivity  
unreliable power source

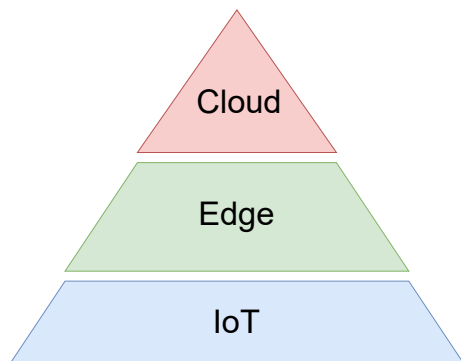
# Multi-Tiered Cloud Architecture



heterogeneity  
unstable network connectivity  
unreliable power source

seamless integration of data throughout all layers  
require *minimization/avoidance of coordination*:

# Multi-Tiered Cloud Architecture




heterogeneity  
unstable network connectivity  
unreliable power source

seamless integration of data throughout all layers  
require *minimization/avoidance of coordination*:


Conflict-Free Replicated Data Types **CRDTs**

# Conflict-Free Replicated Data Type (CRDT)




shared data type

# Conflict-Free Replicated Data Type (CRDT)



shared data type  
allows concurrent update

# Conflict-Free Replicated Data Type (CRDT)




shared data type

allows concurrent update

strong eventual consistency



# Conflict-Free Replicated Data Type (CRDT)



shared data type

allows concurrent update

strong eventual consistency

no complex coordination overhead

# Why Another type of CRDT?

CRDT

does not support non-commutative operation

Op-based CRDTs require exactly-once causal delivery

State-based CRDTs require data type-specific merge

# Why Another type of CRDT?

our work: Log-Structured CRDT (**LSCRDT**)

CRDT

does not support non-commutative operation

Op-based CRDTs require exactly-once causal delivery

State-based CRDTs require data type-specific merge

Log-Structured CRDT (LSCRDT)

supports non-commutative operation

no special delivery requirement

generic merge step

# Why Another type of CRDT?

our work: Log-Structured CRDT (**LSCRDT**)

CRDT

does not support non-commutative operation

Op-based CRDTs require exactly-once causal delivery

State-based CRDTs require data type-specific merge

Log-Structured CRDT (LSCRDT)

supports non-commutative operation

no special delivery requirement

generic merge step

additional LSCRDT features: **operation reversal** and **versioning**

# Data Types Using Logs

## log of operations (OpLog)

operation 1

operation 2

.


.

.

operation n

# Data Types Using Logs


## log of operations (OpLog)




fields in a single entry:

# Data Types Using Logs


## log of operations (OpLog)



·  
·  
·




fields in a single entry:



# Data Types Using Logs


## log of operations (OpLog)




.

.

.



fields in a single entry:





# Data Types Using Logs

## log of operations (OpLog)

operation 1
operation 2

.

.

.

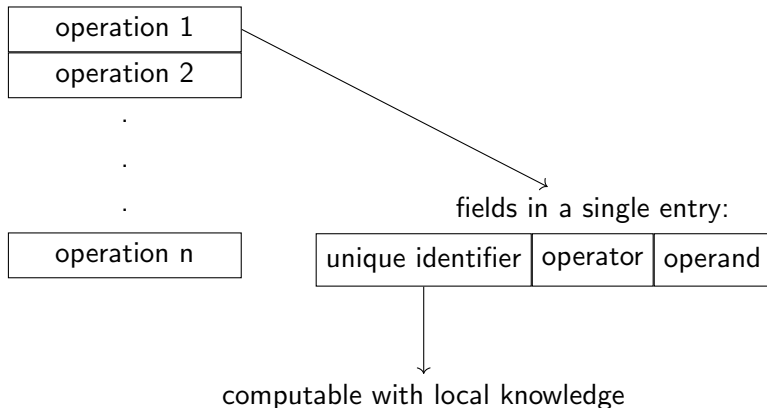
operation n
-------------

fields in a single entry:

unique identifier	operator	operand
-------------------	----------	---------


# Data Types Using Logs

## log of operations (OpLog)



# Data Types Using Logs

## log of operations (OpLog)



computable with local knowledge

version stamps (Lamport timestamps):

counter (one plus highest value) + replica ID

## Replica A OpLog

## Replica A OpLog

1	1A	assign	10
---	----	--------	----

## Replica A OpLog

1	1A	assign	10
2	2A	assign	5

## Replica A OpLog

1	1A	assign	10
2	2A	assign	5
		.	
		.	
		.	
9	9A	assign	7

# LSCRDT Registers

## Replica A OpLog

1	1A	assign	10
2	2A	assign	5

·  
·  
·

read the tail for current value

→ 9

9A	assign	7
----	--------	---



# LSCRDT Registers

## Replica A OpLog

read a specific version

1	1A	assign	10
→ 2	2A	assign	5

·  
·  
·

read the tail for current value

→ 9	9A	assign	7
-----	----	--------	---

# LSCRDT Registers

## Replica A OpLog

read a specific version

→

1	1A	assign	10
2	2A	assign	5
		⋮	
		⋮	
		⋮	
9	9A	assign	7

read the tail for current value

→

9	9A	assign	7
---	----	--------	---

Optimization: maintain a map from version stamp to sequence number

## Replica A OpLog

## Replica A OpLog

1	1A	inc	10
2	2A	dec	5
3	3A	dec	3

## Replica A OpLog

1	1A	inc	10	10
2	2A	dec	5	5
3	3A	dec	3	2

Maintain extra column per entry for cumulative sum

## Replica A OpLog

## Replica A OpLog

1	1A	add	10
2	2A	add	5
3	3A	add	3
4	4A	rem	5

## Replica A OpLog

1	1A	add	10
2	2A	add	5
3	3A	add	3
4	4A	rem	5

## Replica A Checkpoint

1	3A	{10, 5, 3}
---	----	------------

Maintain checkpoint entry at a specific interval



# Replication Overview

OpLog(A)



OpLog(B)



# Replication Overview

OpLog(A)



merge step

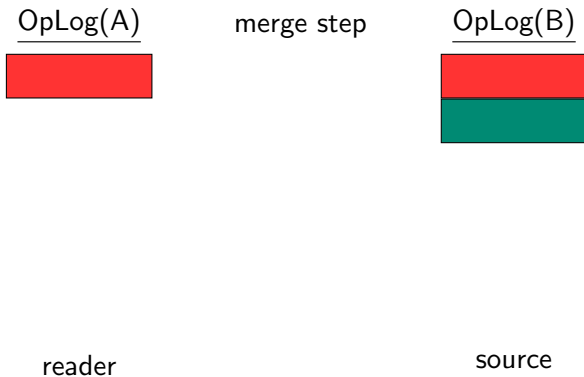
OpLog(B)




reader

source


# Replication Overview




# Replication Overview




# Replication Overview




# Replication Overview



# Replication Overview



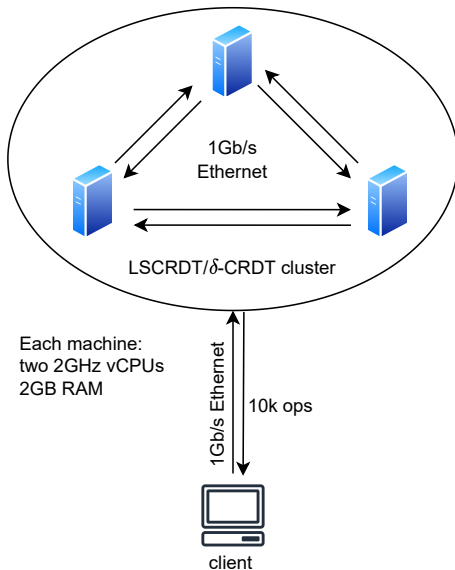
# Replication Overview



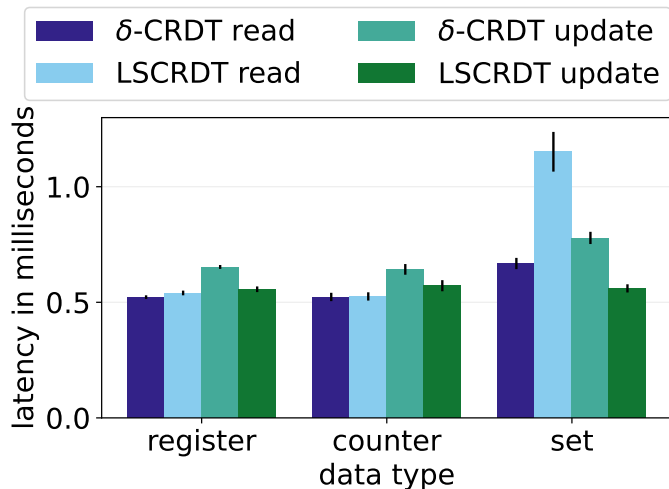
both replicas have **same operations** in the **same order** – convergence!



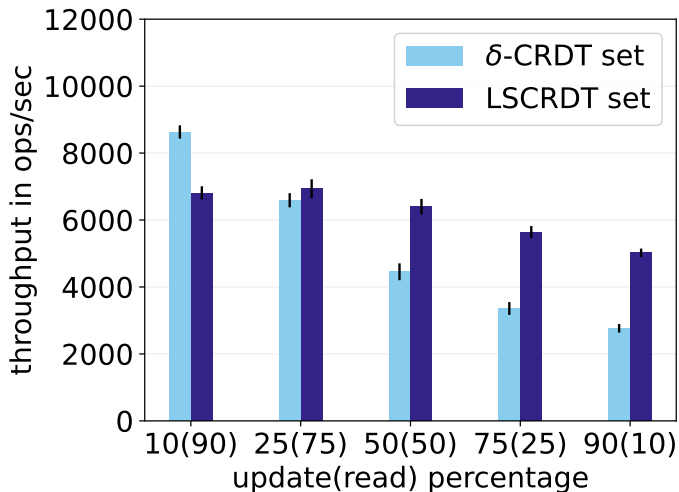
# Experimental Setup



# Single Node Latency



# Throughput for Set



# Summary

- LSCRDT provides a uniform way to reverse operations.

# Summary

- LSCRDT provides a uniform way to reverse operations.
- It removes many restrictions commonly associated with CRDTs.

# Summary

- LSCRDT provides a uniform way to reverse operations.
- It removes many restrictions commonly associated with CRDTs.
- It tracks version histories of data types.

# Summary

- LSCRDT provides a uniform way to reverse operations.
- It removes many restrictions commonly associated with CRDTs.
- It tracks version histories of data types.
- Results suggest LSCRDT can provide 1.8x throughput of  $\delta$ -CRDT.

# Summary

- LSCRDT provides a uniform way to reverse operations.
- It removes many restrictions commonly associated with CRDTs.
- It tracks version histories of data types.
- Results suggest LSCRDT can provide 1.8x throughput of  $\delta$ -CRDT.

nazmus@cs.ucsb.edu