

# The FuzzyLog: A Partially Ordered Shared Log

*Joshua Lockerman, Jose M. Faleiro, Juno Kim, Soham Sankaran*

*Daniel J. Abadi, James Aspnes, Siddhartha Sen, Mahesh Balakrishnan*


Yale University    UC Berkeley    UC San Diego    Cornell University

University of Maryland College Park    Microsoft Research    Facebook

(work done at Yale University)

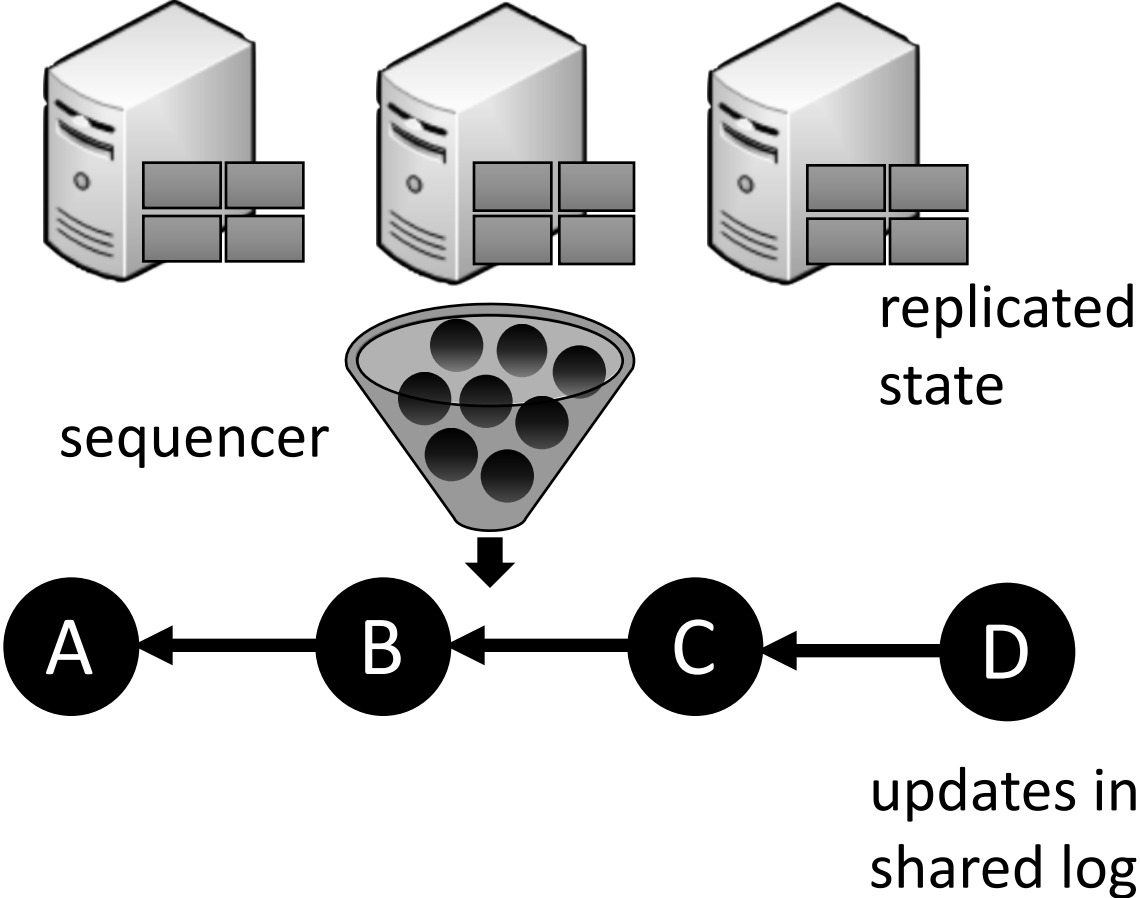
# distributing state

- consistency
- durability
- scalability
- availability



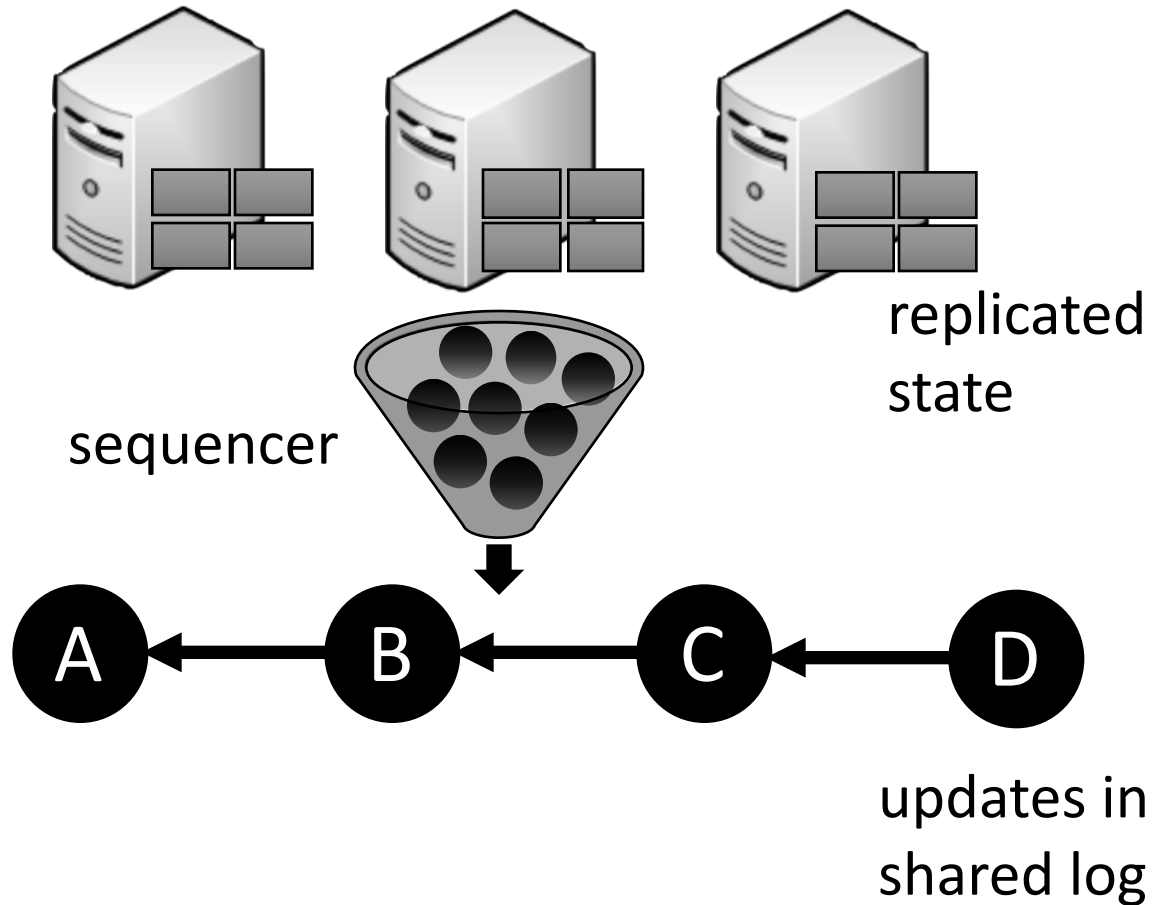
... complex systems!

shared log → simple systems

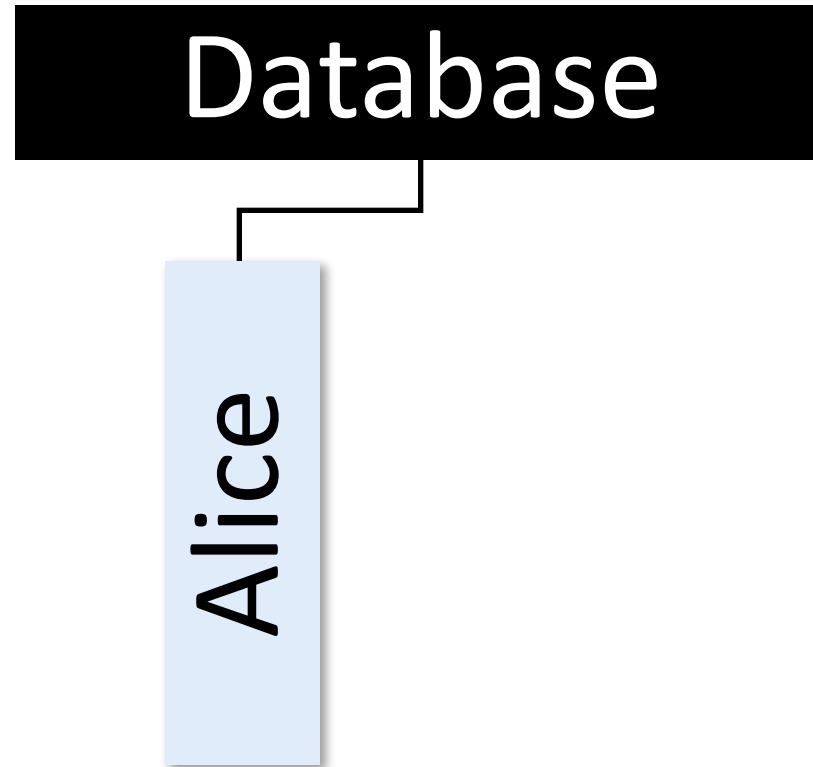


# shared log $\rightarrow$ simple systems


- ✓ consistency
- ✓ durability
- ✗ scalability
- ✗ availability (network partitions)



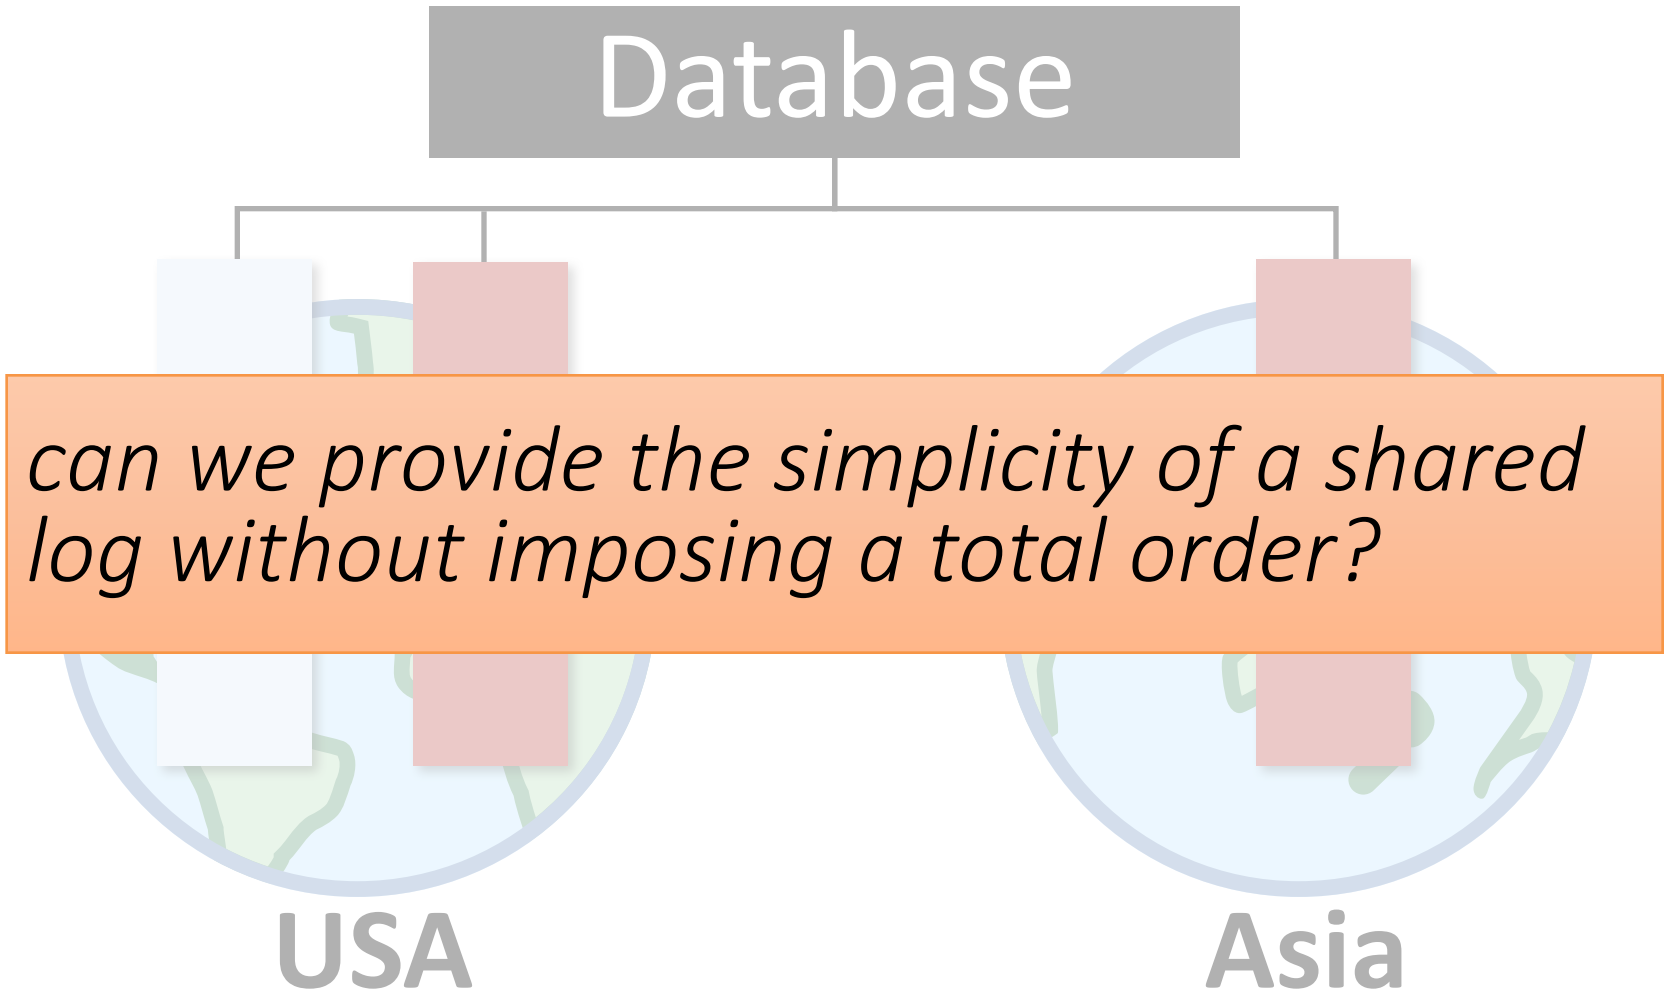
applications are partially ordered



applications are partially ordered

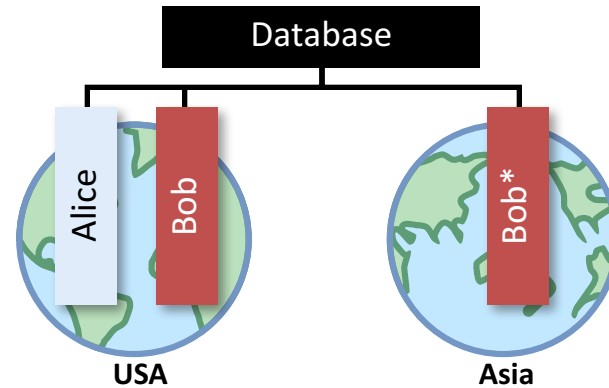


applications are partially ordered



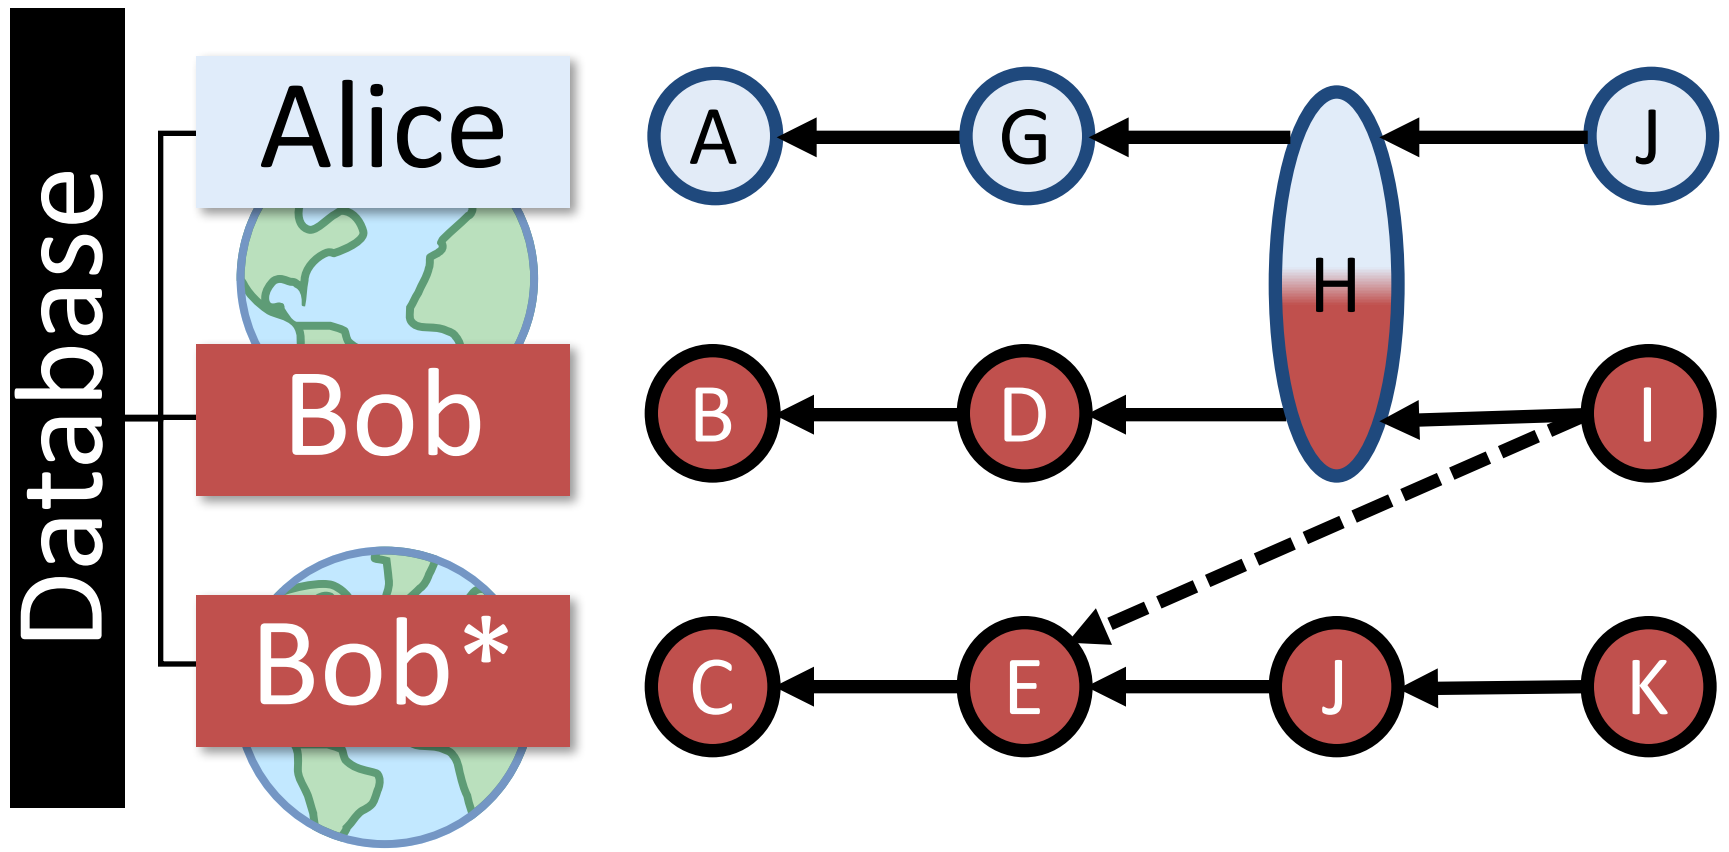
# two sources of partial order

- data shards
- geo-replication





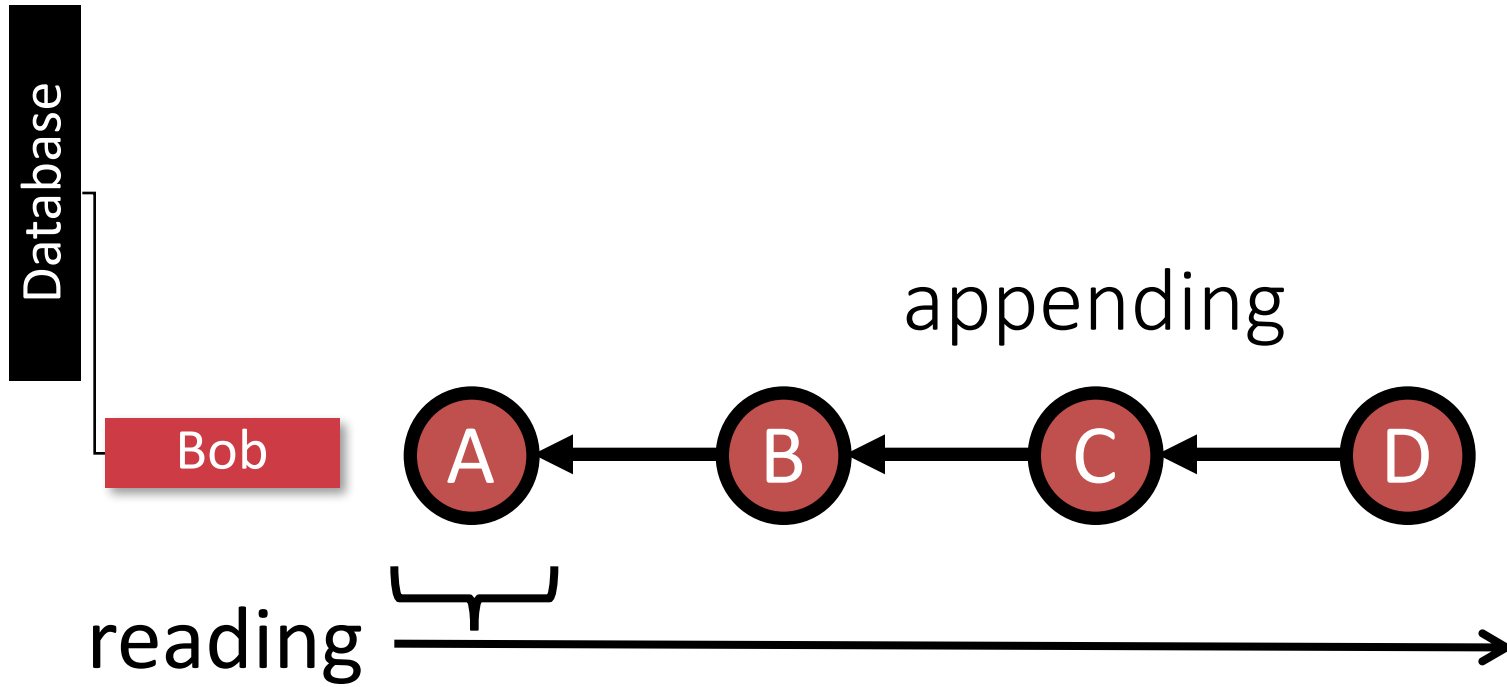
# FuzzyLog – a partially ordered log




# outline

- abstraction
- applications
- implementation
- evaluation

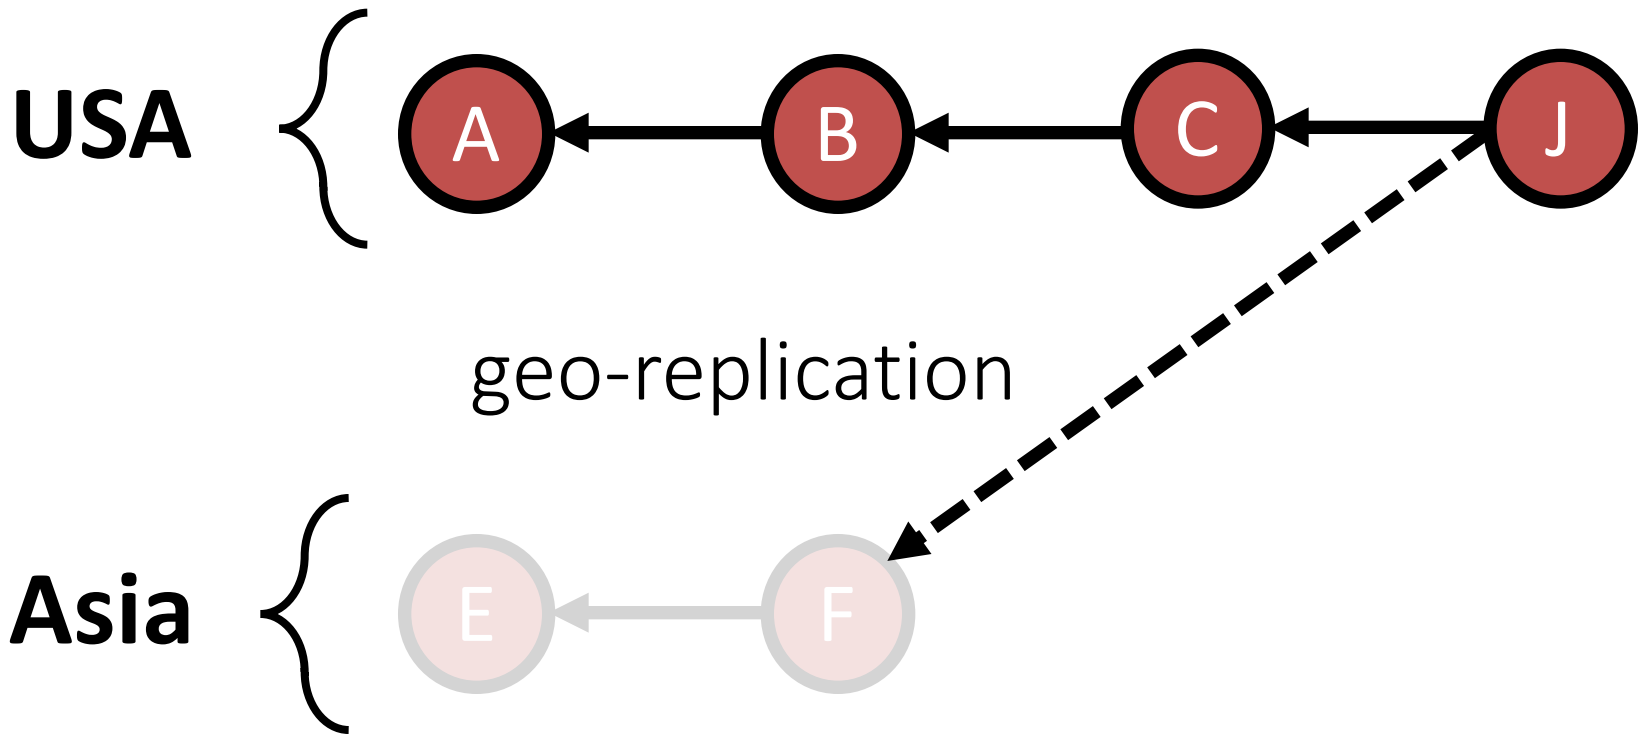
one color per data shard




one color per data shard



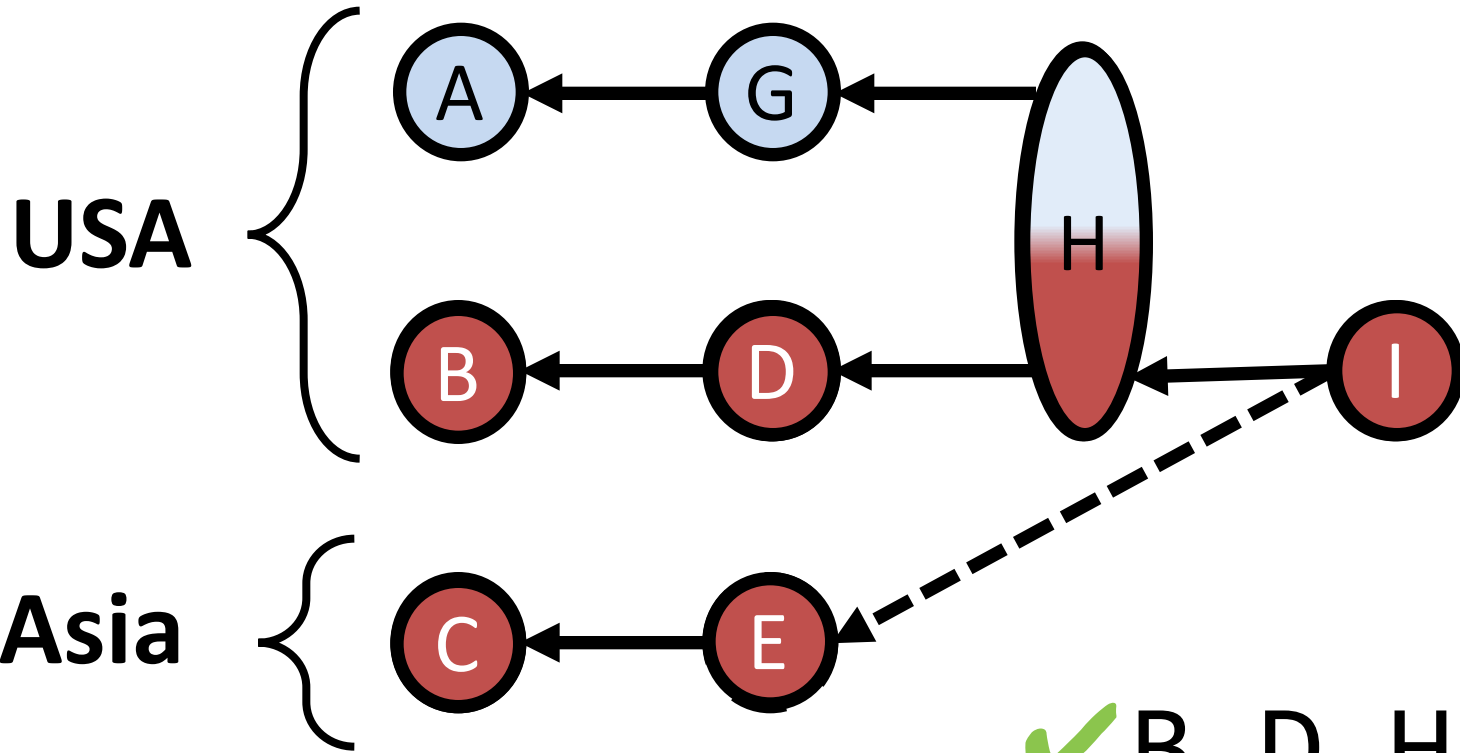
each color is in multiple regions



```
void append(colors, data);
```




```
void sync(callback);
```



✓ B D H C E I

✗ D B H I E C



# recap: the FuzzyLog API

- append a node to one or more colors
  - each color is totally ordered within a region
  - each color is causally ordered across regions.
- sync nodes in a single color
  - nodes are seen in log-order



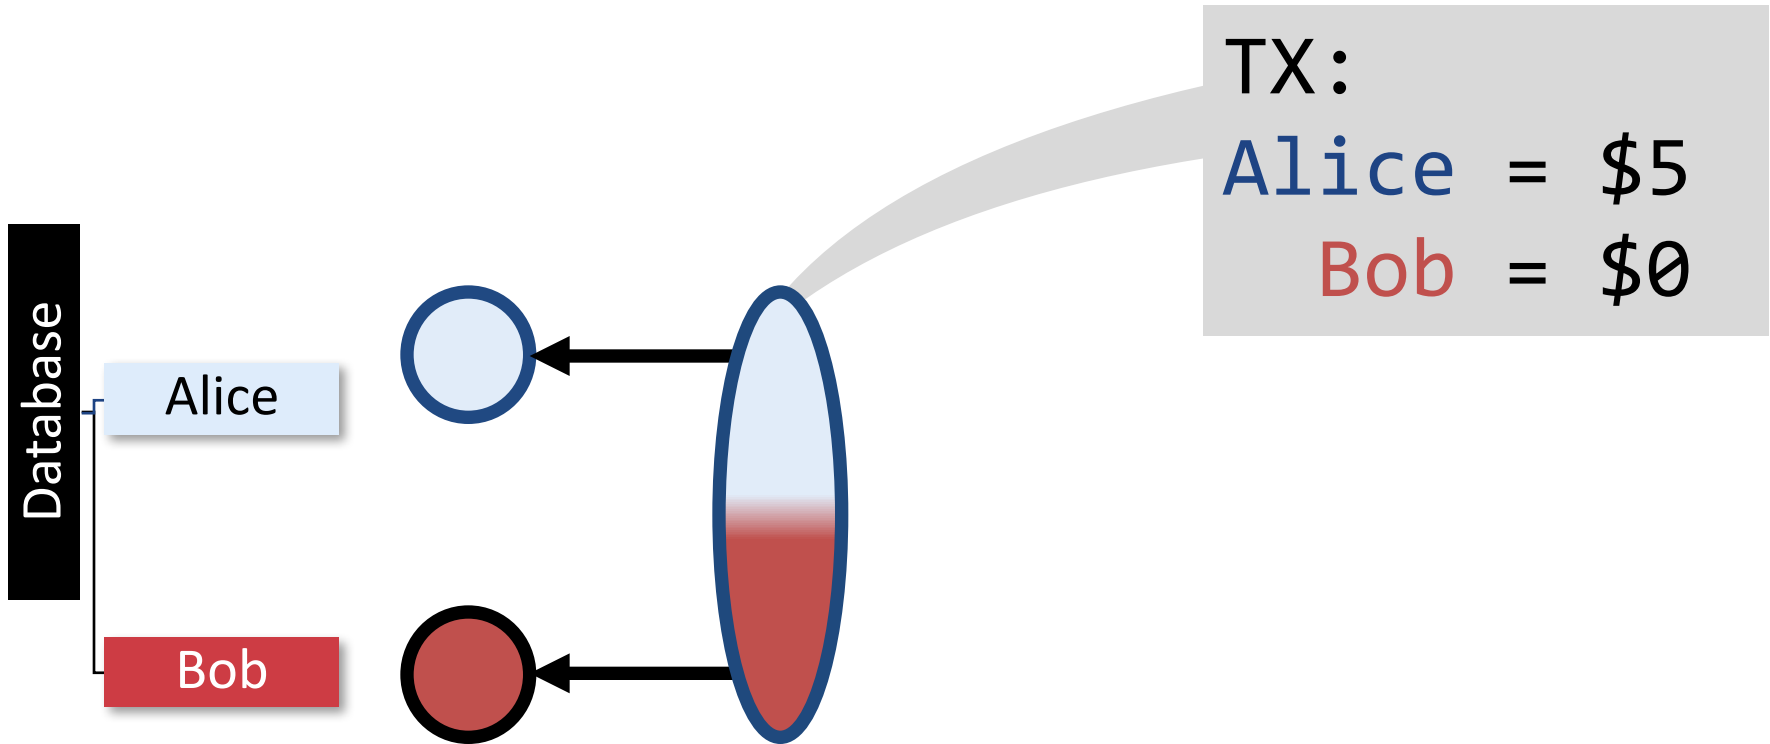
# outline

- abstraction
- **applications**
- implementation
- evaluation

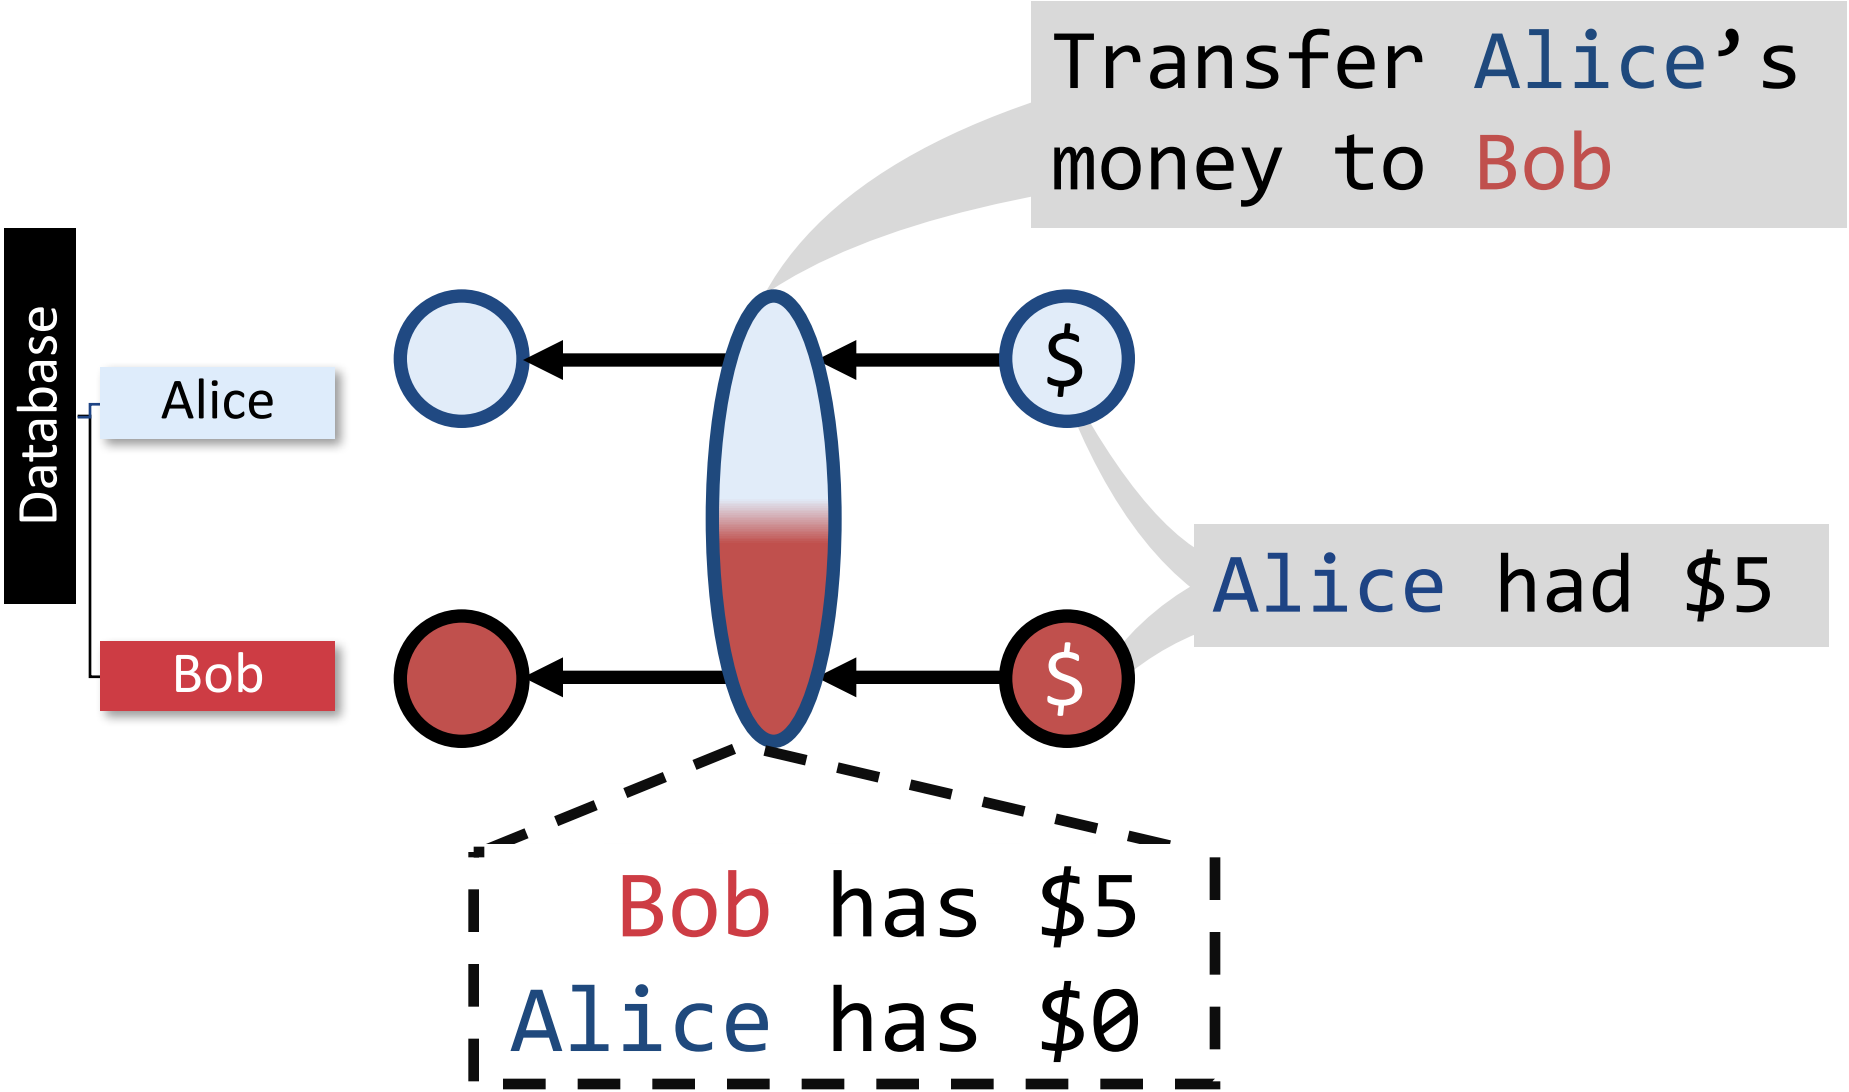
# applications

- transactions: AtomicMap
- best-effort consistency: CAPMap

# AtomicMap: write-only tx




# AtomicMap: read-write tx




# applications

- transactions: AtomicMap
- best-effort consistency: CAPMap


# CAPMap: best-effort consistency



# CAPMap: best-effort consistency



# CAPMap: best-effort consistency









# outline

- abstraction
- applications
- **implementation**
- evaluation



# implementation

- single region
- multiple regions

# FuzzyLog implementation







application servers




FuzzyLog storage servers

# FuzzyLog mechanism: multi-appends




application servers




FuzzyLog storage servers

# FuzzyLog mechanism: multi-appends




# FuzzyLog mechanism: multi-appends

## Phase 1




# FuzzyLog mechanism: multi-appends

## Phase 1



# FuzzyLog mechanism: multi-appends


## Phase 1





# FuzzyLog mechanism: multi-appends

## Phase 1



# FuzzyLog mechanism: multi-appends

Phase 1

Max = 12

Client A

**multi-appends are assigned  
unique timestamps.**

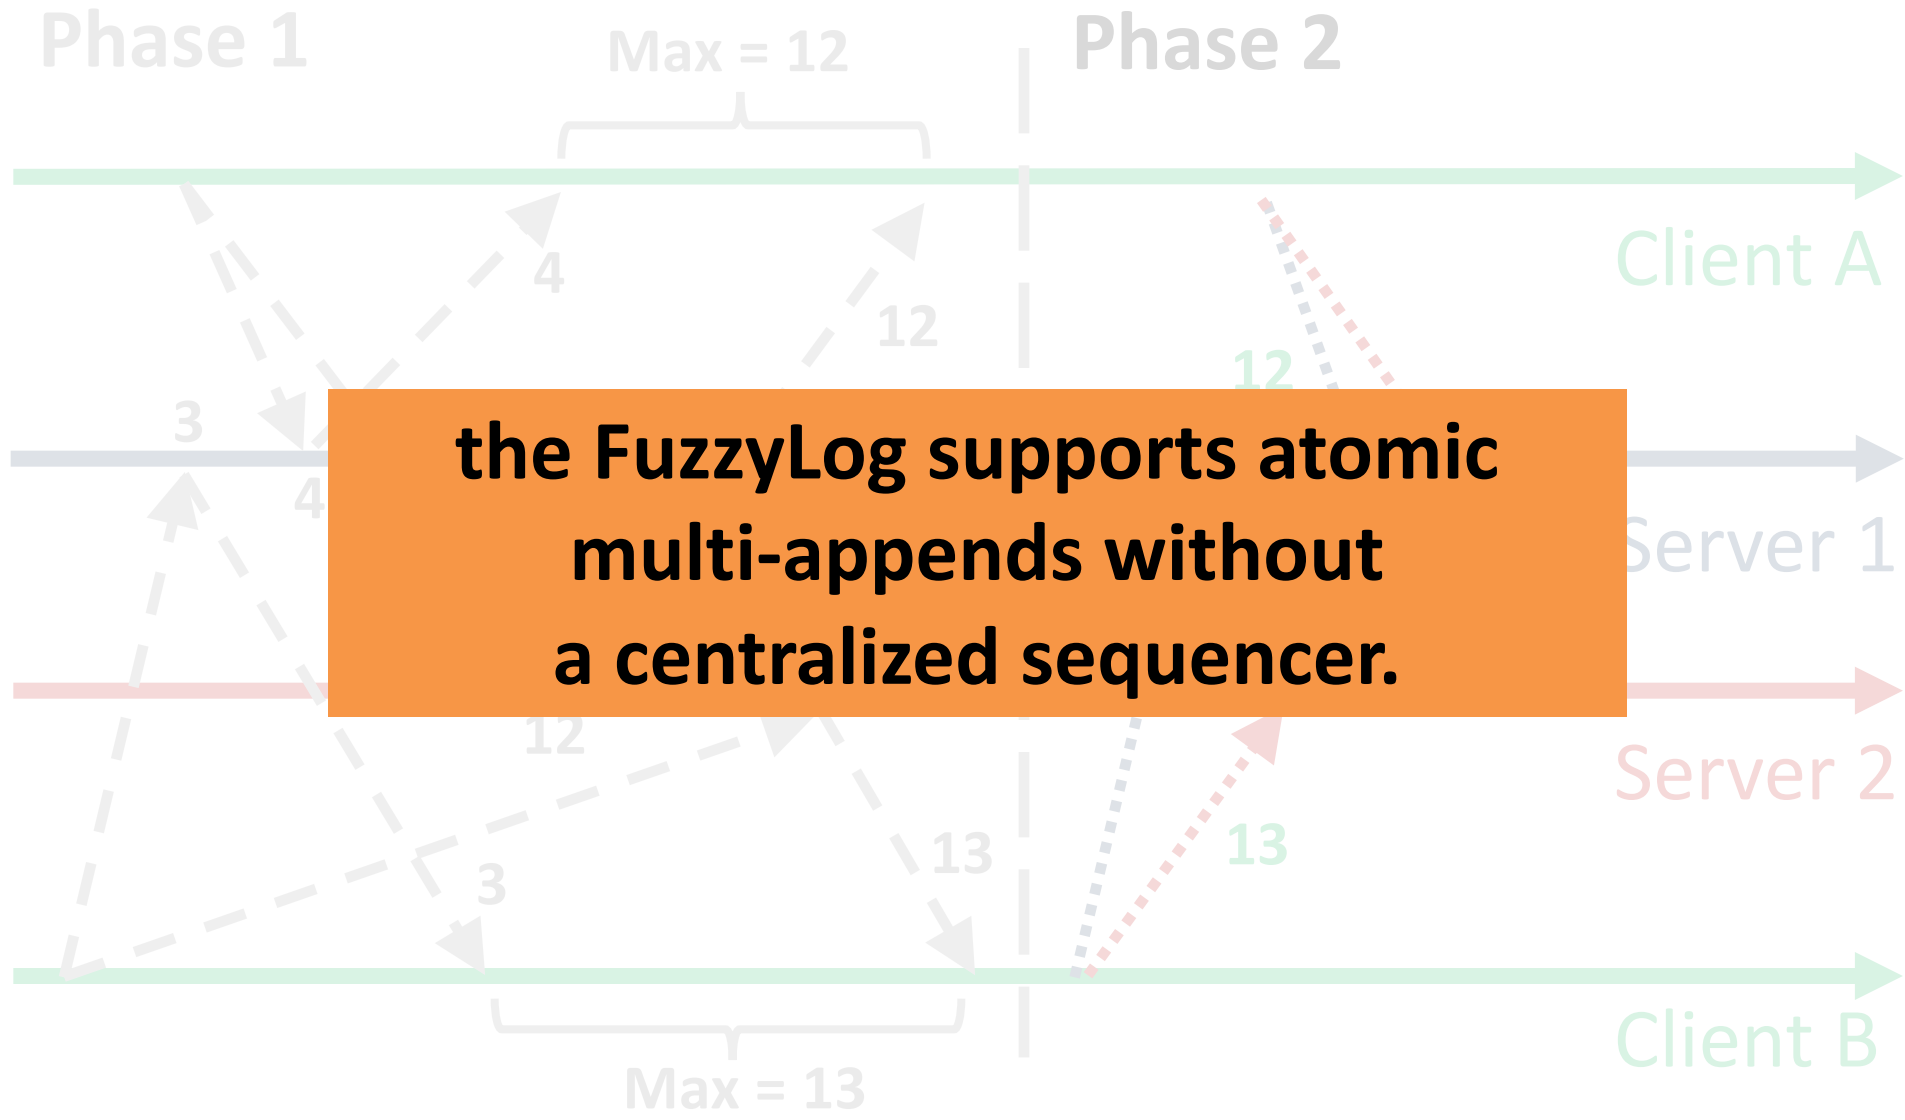
Server 1

Server 2

Client B

Max = 13

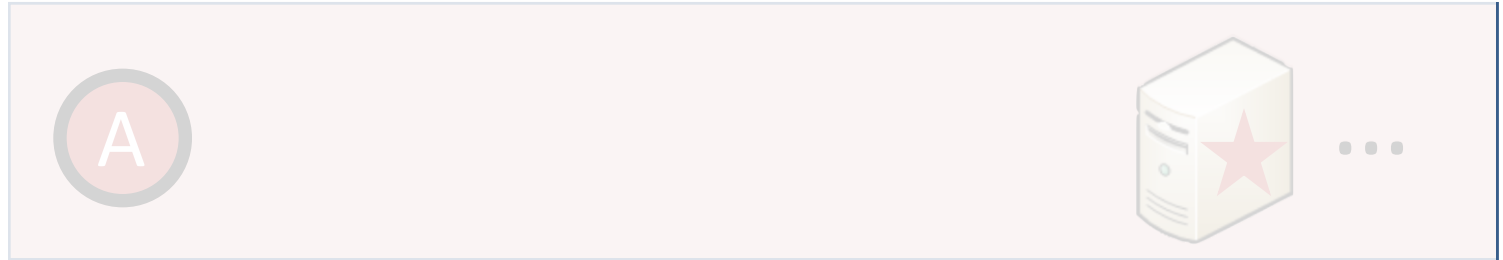
# FuzzyLog mechanism: multi-appends



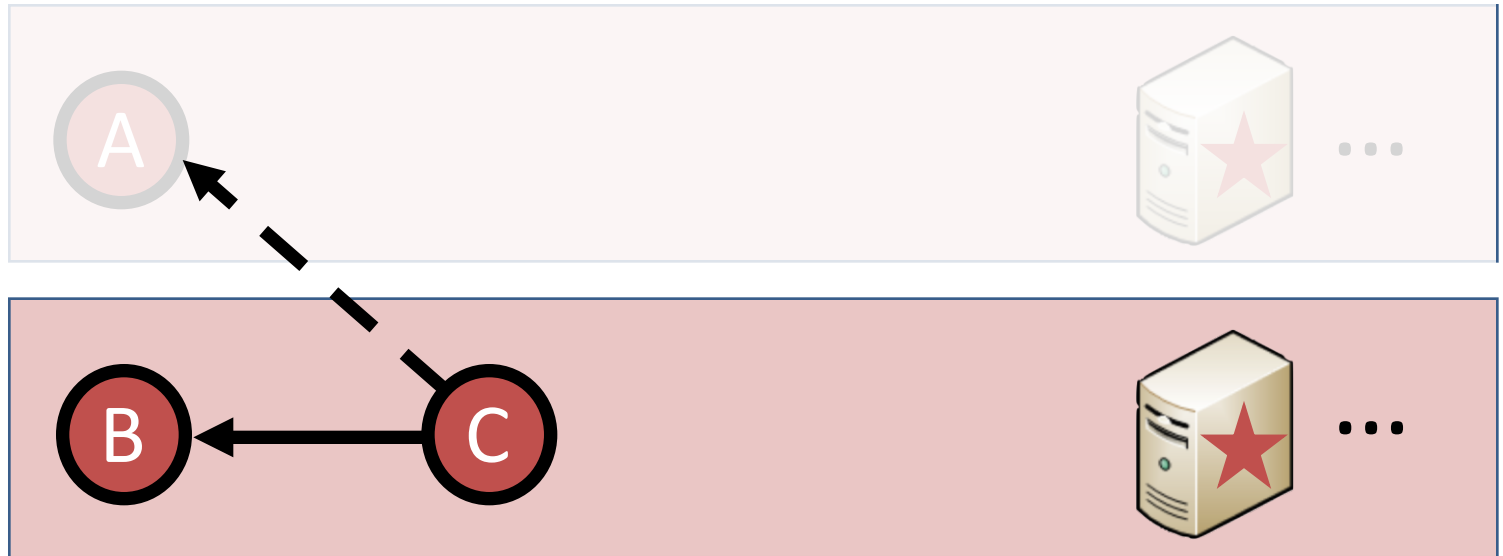
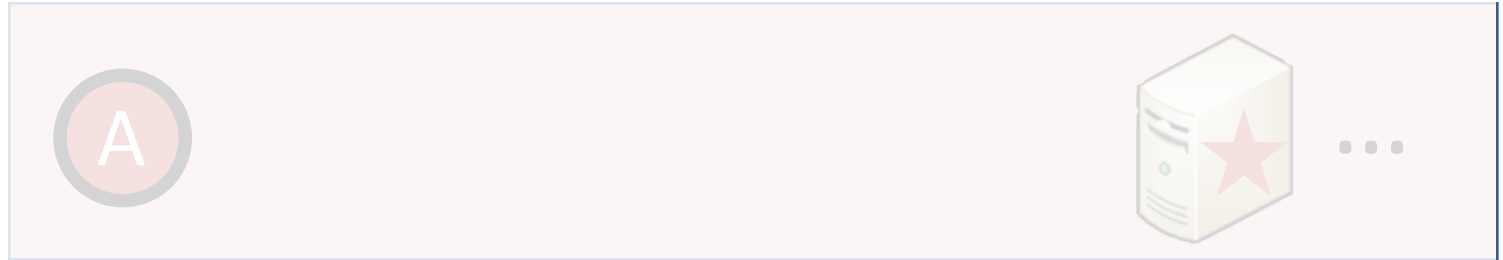
# implementation

- single region
- multiple regions

# FuzzyLog mechanism: shadow logs



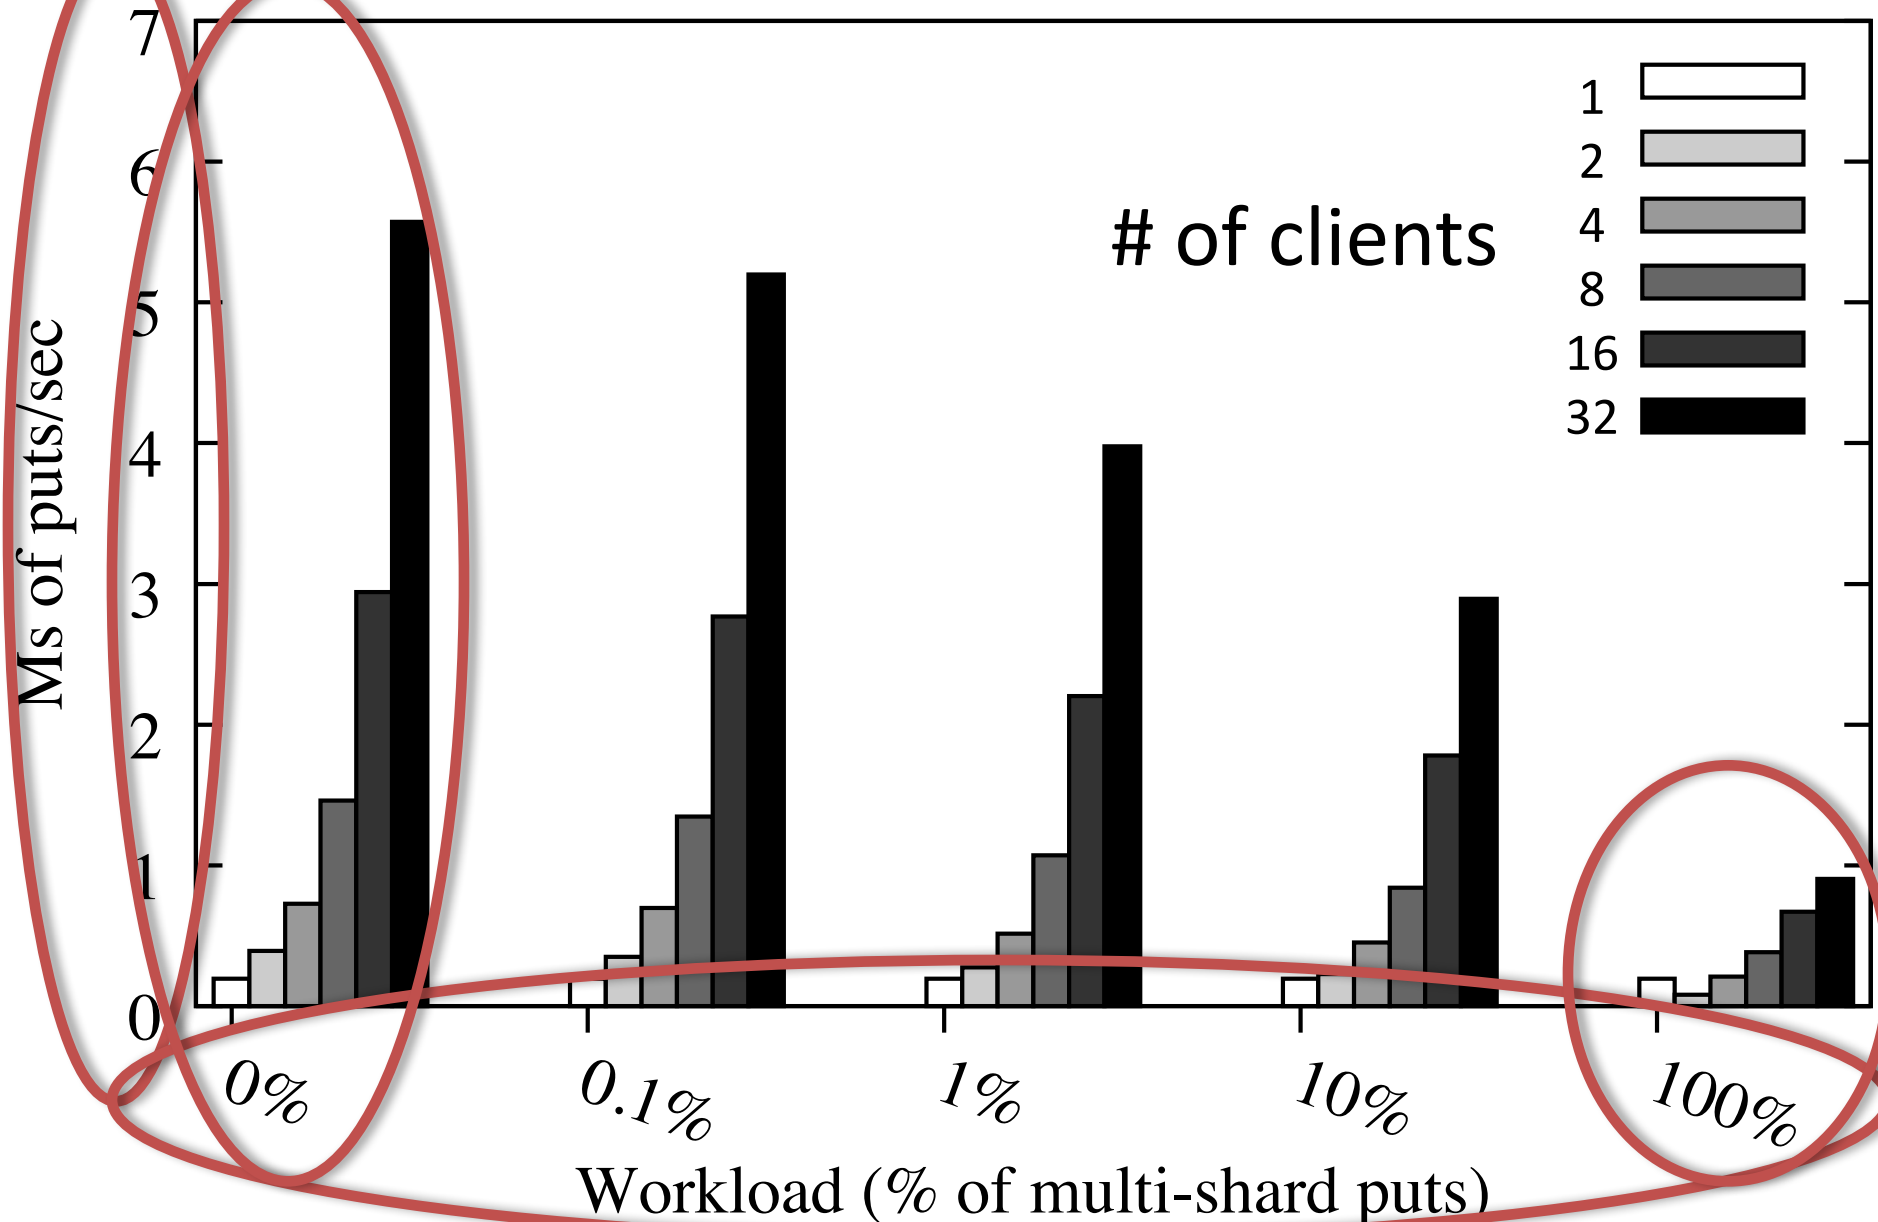
# FuzzyLog mechanism: shadow logs



# outline


- abstraction
- applications
- implementation
- **evaluation**

# AtomicMap performance





# CAPMap performance



# other applications

- transactions: AtomicMap
- best-effort consistency: CAPMap
- RedBlue
- CRDTs
- transactional CRDTs
- FuzzyZookeeper

# conclusion

- a shared log simplifies distributed systems but imposes a total order
- the FuzzyLog is a partially ordered shared log:
  - data sharding
  - geo-replication
- we obtain the simplicity of the shared log approach without its limitations!

A close-up photograph of a weathered log covered in vibrant green moss and lichen. The log is surrounded by a forest floor of brown pine needles and fallen leaves. The text "Thank You!" is centered over the mossy surface.

Thank You!

Globes from:

<https://visualpharm.com/free-icons/globe-595b40b65ba036ed117d2d7e>

<https://visualpharm.com/free-icons/globe%20asia-595b40b65ba036ed117d2d7f>