



# Efficient Synchronization of State-based CRDTs

---

Vitor Enes

Carlos Baquero

João Leitão

(Universidade do Minho)

(Universidade Nova de Lisboa)

Tuesday 3<sup>rd</sup> October, 2017

1. Data Replication
2. Conflict-free Replicated Data Types
  - State-based CRDTs
  - Efficient Synchronization
  - Results

# Data Replication

---

# Data Replication

- Improves **availability**
  - available even when some replicas are unreachable
- Improves **performance**
  - *reduced latency*: geo-distributed replicas give us low-latency links
  - *increased throughput*: multiple replicas serving data

# Data Replication

- Improves **availability**
  - available even when some replicas are unreachable
- Improves **performance**
  - *reduced latency*: geo-distributed replicas give us low-latency links
  - *increased throughput*: multiple replicas serving data
- Several consistency models, e.g.:
  - **Strong Consistency**:
    - illusion of a single-copy by synchronizing replicas on each update
    - very successful within the data center but
    - wide-area networks increase latency of requests and decrease system's throughput
  - **Eventual Consistency**:
    - synchronize replicas in the background
    - replicas may diverge

# Conflict-free Replicated Data Types

---



TRIFORK.  
...think software



bet365



riak




Tapjoy




LEAGUE  
of  
LEGENDS

## Operation-based







## Operation-based




## Operation-based




## Operation-based




## State-based




## Operation-based




## State-based




## Operation-based




## State-based



## Operation-based



## State-based



**\$\$ full state transmission \$\$**

## Efficient Synchronization of State-based CRDTs

How can we avoid replica A sending the full state to replica B?

How can we avoid replica A sending the full state to replica B?

## Two fundamental problems

1. A knows something about the state of B:
  - **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B



How can we avoid replica A sending the full state to replica B?

## Two fundamental problems

### 1. A knows something about the state of B:

- **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B
- Techniques:
  - *avoid sending acknowledged entries in  $\delta$ -buffer*  
(Classic Delta-based CRDTs)
  - ...
  - ...

## How can we avoid replica A sending the full state to replica B?

### Two fundamental problems

#### 1. A knows something about the state of B:

- **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B
- Techniques:
  - *avoid sending acknowledged entries in  $\delta$ -buffer*  
(Classic Delta-based CRDTs)
  - ...
  - ...

#### 2. A knows nothing:

- **Goal:** design a protocol that minimizes state transmission

## How can we avoid replica A sending the full state to replica B?

### Two fundamental problems

#### 1. A knows something about the state of B:

- **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B
- Techniques:
  - *avoid sending acknowledged entries in  $\delta$ -buffer*  
(Classic Delta-based CRDTs)
  - ...
  - ...

#### 2. A knows nothing:


- **Goal:** design a protocol that minimizes state transmission
- Techniques:
  - ...
  - ...

## Solving the first problem


$$\mathbf{A} \quad \{\} \xrightarrow{\text{add } a} \{a\}$$

$$\mathbf{B} \quad \{\} \xrightarrow{\text{add } b} \{b\}$$


## Solving the first problem




## Solving the first problem




## Solving the first problem



# Solving the first problem



## Avoid back-propagation of $\delta$ -groups





## How can we avoid replica A sending the full state to replica B?

### Two fundamental problems

#### 1. A knows something about the state of B:

- **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B
- Techniques:
  - *avoid sending acknowledged entries in  $\delta$ -buffer*  
(Classic Delta-based CRDTs)
  - **avoid back-propagation of  $\delta$ -groups (BP)**
  - ...

#### 2. A knows nothing:

- **Goal:** design a protocol that minimizes state transmission
- Techniques:
  - ...
  - ...


## Solving the first problem

**A**  $\{\} \xrightarrow{\text{add } a} \{a\}$


**B**  $\{\} \xrightarrow{\text{add } b} \{b\}$

**C**  $\{\} \xrightarrow{\text{add } c} \{c\}$


## Solving the first problem




## Solving the first problem




## Solving the first problem



# Solving the first problem



## Remove redundant state in $\delta$ -groups



## How can we avoid replica A sending the full state to replica B?







### Two fundamental problems

#### 1. A knows something about the state of B:




- **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B
- Techniques:
  - *avoid sending acknowledged entries in  $\delta$ -buffer*  
(Classic Delta-based CRDTs)
  - *avoid back-propagation of  $\delta$ -groups* (BP)
  - **remove redundant state in  $\delta$ -groups** (RR)

#### 2. A knows nothing:

- **Goal:** design a protocol that minimizes state transmission
- Techniques:
  - ...
  - ...



- State-Based
- Delta-Based
- Delta-Based BP
- Delta-Based RR
- Delta-Based BP+RR






### State-driven synchronization

$$\mathbf{A} \quad \{a\} \xrightarrow{\text{add } x,y} \{a, x, y\}$$

$$\mathbf{B} \quad \{a\} \xrightarrow{\text{add } z} \{a, z\}$$


# Solving the second problem

## State-driven synchronization




# Solving the second problem

## State-driven synchronization



# Solving the second problem

## State-driven synchronization




## Digest-driven synchronization

**A**  $\{a\} \xrightarrow{\text{add } x,y} \{a, x, y\}$


**B**  $\{a\} \xrightarrow{\text{add } z} \{a, z\}$

# Solving the second problem

## State-driven synchronization




## Digest-driven synchronization




# Solving the second problem

## State-driven synchronization




## Digest-driven synchronization




# Solving the second problem

## State-driven synchronization



## Digest-driven synchronization



## How can we avoid replica A sending the full state to replica B?

### Two fundamental problems

#### 1. A knows something about the state of B:


- **Goal:** A to compute a delta ( $\delta$ ) of minimum size to be sent to B
- Techniques:
  - *avoid sending acknowledged entries in  $\delta$ -buffer* (Classic Delta-based CRDTs)
  - *avoid back-propagation of  $\delta$ -groups* (BP)
  - *remove redundant state in  $\delta$ -groups* (RR)

#### 2. A knows nothing:

- **Goal:** design a protocol that minimizes state transmission
- Techniques:
  - **state-driven synchronization**
  - **digest-driven synchronization**




**PCounter - Line**




Transmission (KB)

**PCounter - Ring**




Transmission (KB)

**PCounter - HyParView**




Transmission (KB)

**GSet - Line**




Transmission (KB)

**GSet - Ring**




Transmission (KB)

**GSet - HyParView**




Transmission (KB)

**AWSet - Line**




Transmission (KB)

**AWSet - Ring**



Transmission (KB)

**AWSet - HyParView**



Transmission (KB)

- State-Based
- State-Driven
- Digest-Driven

[bit.ly/msc-presentation](https://bit.ly/msc-presentation)



# Efficient Synchronization of State-based CRDTs

---

Vitor Enes

Carlos Baquero

João Leitão

(Universidade do Minho)

(Universidade Nova de Lisboa)

Tuesday 3<sup>rd</sup> October, 2017