efficient synchronization of state-based CRDTs

Vitor Enes, Paulo Sérgio Almeida, Carlos Baquero, João Leitão

9 Apr. 2019 @ ICDE'19





rich abstraction - register, counter, set, map, ...



rich abstraction - register, counter, set, map, ...

why replicate?quality of service





rich abstraction - register, counter, set, map, ...

why replicate?quality of service





rich abstraction - register, counter, set, map, ...

why replicate? - quality of service





rich abstraction - register, counter, set, map, ...

why replicate? - quality of service

no conflicts?!

- eventual consistency
- automatic conflict resolution







- operation-based
- state-based
 - delta-based



- operation-based
- state-based
 - delta-based
- problem with delta-based: naive delta propagation - filter



- operation-based
- state-based
 - delta-based
- problem with delta-based: naive delta propagation
 - filter
- solution
 - decomposition; filter; join



- operation-based
- state-based
 - delta-based

- problem with delta-based: naive delta propagation

- filter
- solution
 - decomposition; filter; join







- operation-based
- state-based
 - delta-based

- problem with delta-based: naive delta propagation

- filter
- solution
 - decomposition; filter; join









DECOMPOSITION





- operation-based
- state-based
 - delta-based

- problem with delta-based: naive delta propagation

- filter
- solution
 - decomposition; filter; join







- operation-based
- state-based
 - delta-based

- problem with delta-based: naive delta propagation

- filter
- solution
 - decomposition; filter; join





- operation-based
- state-based
 - delta-based
- problem with delta-based: naive delta propagation
 - filter
- solution
 - decomposition; filter; join
- decomposition of state-based CRDTs





- operation-based
- state-based
 - delta-based
- problem with delta-based: naive delta propagation
 - filter
- solution
 - decomposition; filter; join
- decomposition of state-based CRDTs
- results





- operation-based
- state-based
 - delta-based
- problem with delta-based: naive delta propagation
 - filter
- solution
 - decomposition; filter; join
- decomposition of state-based CRDTs
- results
- summary





	operation-based	state-based	delta-based
middleware guarantees	exactly-once causal		
payload	operation	full state	delta



	operation-based	state-based	delta-based
middleware guarantees	exactly-once causal		
payload	operation	full state	delta



	operation-based	state-based	delta-based
middleware guarantees	exactly-once causal		
payload	operation 📀	full state	delta



	operation-based	state-based	delta-based
middleware guarantees	exactly-once causal	—	
payload	operation 📀	full state	delta



	operation-based	state-based	delta-based
middleware guarantees	exactly-once causal	—	
payload	operation 📀	full state 😑	delta



	operation-based	state-based	delta-based
middleware guarantees	exactly-once causal	—	—
payload	operation 📀	full state 🖨	delta



	operation-bas
middleware guarantees	exactly-once causal
payload	operation





	operation-bas
middleware guarantees	exactly-once causal
payload	operation









BANDWIDTH (as bad as state-based)



in reality...



BANDWIDTH (as bad as state-based)



in reality... CPU (worse than state-based)





main contribution





- simplest example: set



- simplest example: set S, \subseteq, \cup subset set union



- simplest example: set









- state-based: mutators (m) inflate local state
- delta-based: delta-mutators (dm) return deltas

flate local state (**dm**) return deltas







- state-based: mutators (m) inflate local state
- delta-based: delta-mutators (dm) return deltas



nflate local state 5 (**dm**) return deltas







- state-based: mutators (m) inflate local state
- delta-based: delta-mutators (dm) return deltas

> add c to the set {a, b} m({a, b}) = {a, b, c}

nflate local state 5 (**dm**) return deltas







- state-based: mutators (m) inflate local state
- delta-based: delta-mutators (dm) return deltas

> add c to the set {a, b} m({a, b}) = {a, b, c} dm({a, b}) = {c}

nflate local state 5 (**dm**) return deltas


state-based CRDTs





- state-based: mutators (m) inflate local state
- delta-based: delta-mutators (dm) return deltas
 - 1. joined with local state
 - 2. added to a delta-buffer (to be propagated)

> add c to the set {a, b} m({a, b}) = {a, b, c} dm({a, b}) = {c}



delta-buffer notation







delta-buffer notation







delta-buffer notation













B



delta-buffer

·····









































































{x, y}

[{x}, {x, y}]

{x, y}

{x, y}

① C filters received {x, y} = is there something new? $= \{x, y\} \not\subseteq \{x\}$



{x, y}



1) C filters received {x, y} = is there something new? = $\{x, y\} \not\subseteq \{x\}$

SOLUTION: decomposition; filter; join

{x, y} [{x}, {x, y}]



[{x, y}]

{**x**, **y**}

{x, y}

<u>ר</u> '



{x, y}



1) C filters received {x, y} = is there something new? = $\{x, y\} \not\subseteq \{x\}$

SOLUTION: decomposition; filter; join = select what's new

[{x}, {x, y}]

{x, y}

1



{x, y}



1) C filters received {x, y} = is there something new? = $\{x, y\} \not\subseteq \{x\}$

SOLUTION: decomposition; filter; join = select what's new $= \cdots = \{y\}$

{x, y}

[{x}, {x, y}]

{x} [{x}]

[{x, y}]

{**x**, **y**}

{x, y}



{x, y}

- example: $decomposition(\{a, b, c\}) = \{\{a\}, \{b\}, \{c\}\}\}$

decomposition



- example: decomposition $(\{a, b, c\}) = \{\{a\}, \{b\}, \{c\}\}\}$

- S is a decomposition of x if:

decomposition



- example: decomposition $(\{a, b, c\}) = \{\{a\}, \{b\}, \{c\}\}\}$
- S is a decomposition of x if: 1. join of S gives x

 $X S_1 = \{\{b\}, \{c\}\}$



- example: decomposition($\{a, b, c\}$) = { $\{a\}, \{b\}, \{c\}\}$
- S is a decomposition of x if: 1. join of S gives x
 - 2. no element is redundant

 $X S_1 = \{\{b\}, \{c\}\}$ $X S_2 = \{\{a, b\}, \{b\}, \{c\}\}\}$



- example: decomposition($\{a, b, c\}$) = { $\{a\}, \{b\}, \{c\}\}$
- S is a decomposition of x if: 1. join of S gives x
 - 2. no element is redundant
 - 3. every element is irreducible

 $X S_1 = \{\{b\}, \{c\}\}$ $X S_2 = \{\{a, b\}, \{b\}, \{c\}\}\}$ $X S_3 = \{\{a, b\}, \{c\}\}$



- example: decomposition($\{a, b, c\}$) = { $\{a\}, \{b\}, \{c\}\}$
- S is a decomposition of x if: 1. join of S gives x
 - 2. no element is redundant
 - 3. every element is irreducible

 $X S_1 = \{\{b\}, \{c\}\}$ $X S_2 = \{\{a, b\}, \{b\}, \{c\}\}\}$ $X S_3 = \{\{a, b\}, \{c\}\}$ $\checkmark S_4 = \{\{a\}, \{b\}, \{c\}\}\}$



- example:
- S is a decomposition of x if: 1. join of S gives x
 - 2. no element is redundant
 - 3. every element is irreducible

 $\mathsf{decomposition}(\{a, b, c\}) = \{\{a\}, \{b\}, \{c\}\} = \Downarrow \{a, b, c\}$

 $X S_1 = \{\{b\}, \{c\}\}$ $X S_2 = \{\{a, b\}, \{b\}, \{c\}\}\}$ $X S_3 = \{\{a, b\}, \{c\}\}$ $\checkmark S_4 = \{\{a\}, \{b\}, \{c\}\}\}$



decomposition; filter; join = difference

decomposition; filter; join = difference

- example:

decomposition; filter; join = difference

- example: difference($\{x, y\}, \{x\}$) = $\{y\}$
decomposition; filter; join - example: difference($\{x, y\}, \{x\}$) = $\{y\} = \Delta(\{x, y\}, \{x\})$

decomposition; filter; join - example: difference $(\{x, y\}, \{x\}) = \{y\} = \Delta(\{x, y\}, \{x\})$

$\Delta(a,b) = \bigcup \{ x \in \Downarrow a \mid x \not\subseteq b \}$

decomposition; filter; join - example: difference $(\{x, y\}, \{x\}) = \{y\} = \Delta(\{x, y\}, \{x\})$

$\Delta(a,b) = \bigcup \{ x \in \Downarrow a \mid x \not\subseteq b \}$ **DECOMPOSITION**

decomposition; filter; join - example: difference $(\{x, y\}, \{x\}) = \{y\} = \Delta(\{x, y\}, \{x\})$

$\Delta(a,b) = \bigcup \{ x \in \Downarrow a \mid x \not\subseteq b \}$ Decomposition

decomposition; filter; join - example: difference($\{x, y\}, \{x\}$) = $\{y\} = \Delta(\{x, y\}, \{x\})$

JOIN $\Delta(a,b) = \bigcup \{ x \in \Downarrow a \mid x \not\subseteq b \}$ Decomposition







(1) C removes redundant state in received {x, y} = $\Delta(\{x, y\}, \{x\}) = \{y\}$





{x, y}

[{x}, {y}]

(1) C removes redundant state in received {x, y} = $\Delta(\{x, y\}, \{x\}) = \{y\}$





{x, y}

[{x}, {y}]

{x, y}

{x, y]

X

[{x, y}]

(1) C removes redundant state in received {x, y} = $\Delta(\{x, y\}, \{x\}) = \{y\}$



12

{x, y}

[{x}, {y}]





























(1) B should avoid back-propagation of {x}





(1) B should avoid back-propagation of {x}



- test the effect of cycles in the network topology:

- test the effect of cycles in the network topology:



partial-mesh





transmission ratio wrto BP+RR

set - tree

set - mesh

- □ state-based
- delta-based
- delta-based BP
- delta-based RRdelta-based BP+RR





classic delta-based is as bad as state-based



set - mesh



wrto BP+RR

transmissio

80

60

40

20

0

BP is enough if no cycles RR needed in the general case

classic delta-based is as bad as state-based

set - tree

set - mesh







but sets are easy...



but sets are easy...

for each CRDT composition technique, there is a decomposition rule



but sets are easy...

$c \in C: \qquad \qquad \Downarrow c = \{c\}$ $\langle a, b \rangle \in A \times B: \ \Downarrow \langle a, b \rangle = \Downarrow a \times \{\bot\} \cup \{\bot\} \times \Downarrow b$ $\langle c, a \rangle \in C \boxtimes A$: $\Downarrow \langle c, a \rangle = \Downarrow c \times \Downarrow a$ Left $a \in A \oplus B$: \Downarrow Left $a = \{ Left v \mid v \in \Downarrow a \}$ Right $b \in A \oplus B$: \Downarrow Right $b = \{$ Right $v \mid v \in \Downarrow b\}$ $s \in \mathcal{P}(U): \quad \Downarrow s = \{\{e\} \mid e \in s\}$

for each CRDT composition technique, there is a decomposition rule

- $f \in U \hookrightarrow A: \ \Downarrow f = \{k \mapsto v \mid k \in \mathsf{dom}(f) \land v \in \Downarrow f(v)\}$



more results in the paper

- counter and map micro-benchmark
- comparison with:
 - operation-based CRDTs
 - scuttlebutt and scuttlebutt variant
- retwis benchmark



more results in the paper

- counter and map micro-benchmark
- comparison with:
 - operation-based CRDTs
 - scuttlebutt and scuttlebutt variant
- retwis benchmark
- optimal delta-mutators
- state-based CRDTs are distributive lattices - (not simply join-semilattices)



summary



- identified inefficiencies in delta propagation

summary



- identified inefficiencies in delta propagation
- **difference** (Δ) between two states

summary

- introduced concept of decomposition of state-based CRDTs



- identified inefficiencies in delta propagation
- introduced concept of decomposition of state-based CRDTs
 - difference (Δ) between two states
- proposed **BP** and **RR** optimizations
 - BP is enough in acyclic topologies
 - RR is needed in the more general case

summary



efficient synchronization of state-based CRDTs

Vitor Enes, Paulo Sérgio Almeida, Carlos Baquero, João Leitão

m vitorenes.org



9 Apr. 2019 @ ICDE'19