## **Just-Right Consistency**

As available as possible As consistent as necessary Correct by design

Marc Shapiro, Sorbonne-Université—LIP6 & Inria
Annette Bieniusa, U. Kaiserslautern
Nuno Preguiça, U. Nova Lisboa
Christopher Meiklejohn, U. Catholique de Louvain
Valter Balegas, U. Nova Lisboa

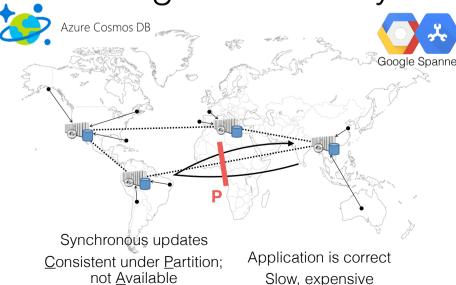


[ Just-Right Consistency ]

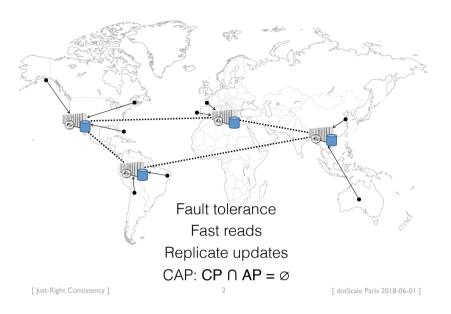


[ dotScale Paris 2018-06-01 ]

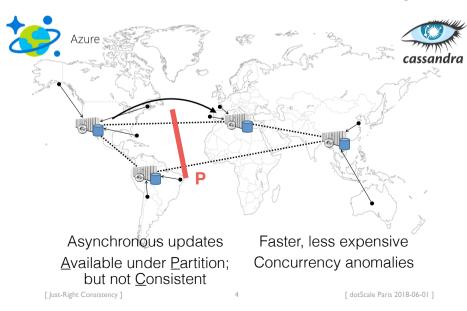
## Strong Consistency



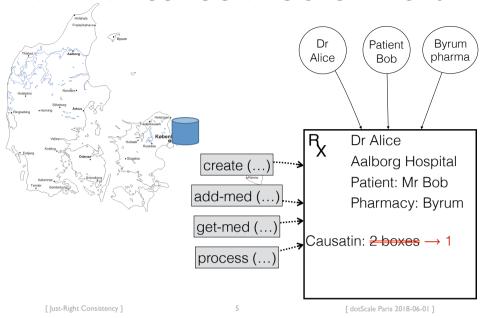
#### Geo-distributed DB



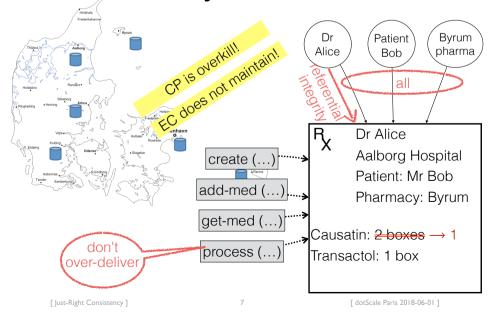
## **Eventual Consistency**



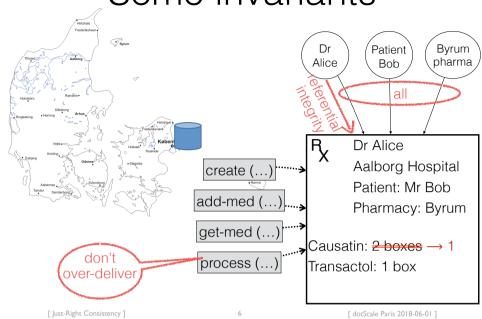
#### FMK Fælles Medicinkort



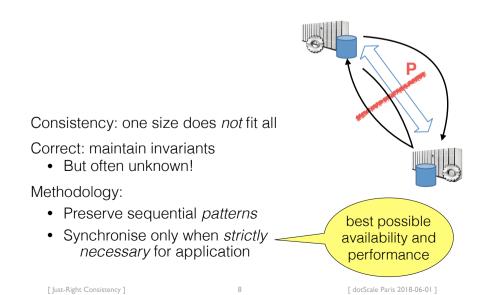
## Availability + invariants?



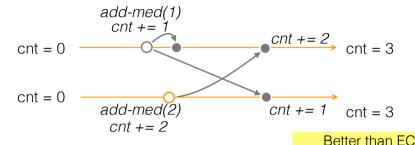
#### Some invariants



## Approach



#### **CRDTs**



CP is overkill!

1=RHS!

2=LHS!

Concurrent, asynchronous updates

- Standard register model: assignments ⇒ CP
- AP ⇒ concurrent updates + merge

CRDT: register, counter, set, map, sequence

• Plug-in replacement for sequential type

Rx.Patient: write\_once\_reg. Rx.Meds: set

[ Just-Right Consistency ] 9 [ dotScale Paris 2018-06-01 ]

## Causal consistency



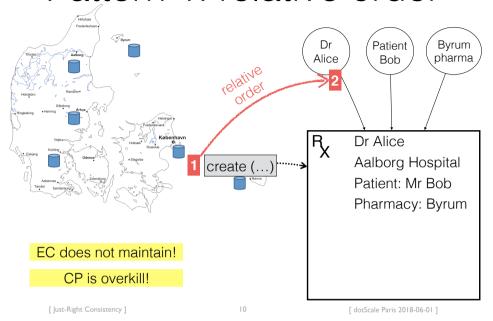
- "Bob points to Rx ⇒ Rx valid"
- General case: LHS ⇒ RHS
- pattern: RHS!; LHS!

Deliver in the right order: Causal

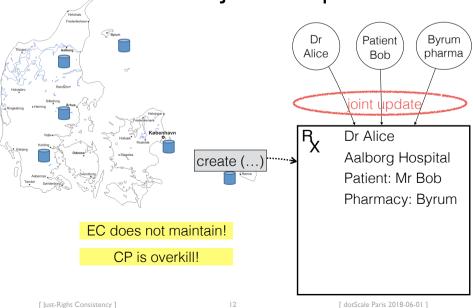
Consistency

#### AP-compatible

#### Pattern 1: relative order



### Pattern 2: joint update



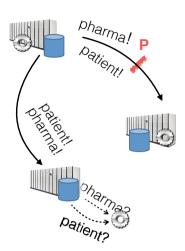
## All-or-nothing

*create-p* updates doctor, patient & pharmacy record

Transmit joint updates together

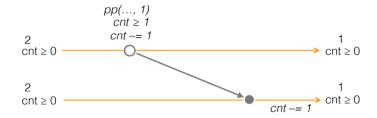
+ Read from same snapshot

AP-compatible



[ Just-Right Consistency ] 13 [ dotScale Paris 2018-06-01 ]

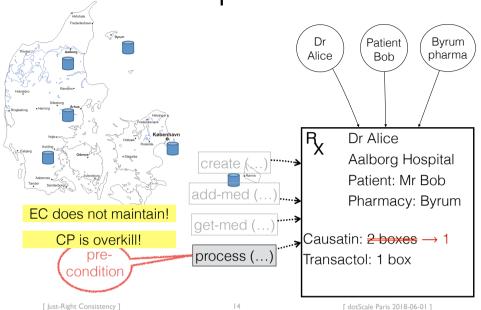
#### **CAP-sensitive** invariants



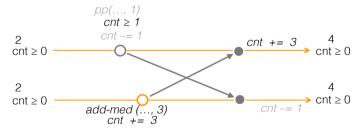
```
process-p (..., nb) {
  if cnt \ge nb  // precondition at source
  cnt = nb  // at every replica
  } // cnt \ge 0
```

Precondition *stable* w.r.t. concurrent *add-med* Concurrency OK

#### Pattern 3: precondition



## Stable precondition



```
process-p (..., nb) {

if cnt \ge nb  // precondition at source

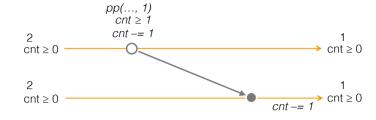
cnt = nb  // at every replica

} // cnt \ge 0
```

Precondition *stable* w.r.t. concurrent *add-med* Concurrency OK

[Just-Right Consistency] | 15 [dotScale Paris 2018-06-01] [Just-Right Consistency] | 16 [dotScale Paris 2018-06-01]

#### CAP-sensitive invariants



```
process-p (..., nb) {
  if cnt \ge nb  // precondition at source
  cnt = nb  // at every replica
  } // cnt \ge 0
```

Precondition not stable w.r.t. concurrent process-p

- Forbid concurrency? Synchro, CP.
- Or remove invariant? AP, degraded semantics

[ Just-Right Consistency ] 17 [ dotScale Paris 2018-06-01 ]

# Summary: Just-Right Consistency

Tailor consistency to application invariants

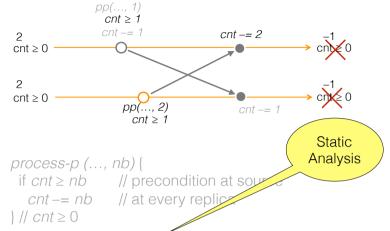
• (possibly unknown)

Baseline: Correct app, 1-copy, one op. at a time

Three patterns:

- Ordered updates ⇒ Causal Consistency, AP
- Joint updates ⇒ All-or-nothing, AP
- CAP-sensitive: precondition
  - Stable ⇒ concurrent OK. AP.
  - Otherwise, concurrency control. CP

### Not stable precondition



Precondition not stable w.r.t. concurrent process-p

- Forbid concurrency
- Or, give up on invariant

[ Just-Right Consistency ] 18 [ dotScale Paris 2018-06-01 ]



#### CRDT data model

- Register, counter, set, map, sequence
- Extends sequential semantics

Transactional Causal Consistency (TCC)

- Strongest AP model
- Supports Joint Updates, Relative Order

CISE static analysis tools

Open source, well engineered







#### © Creative Commons Attribution-ShareAlike 4.0 Intl. License

#### You are free to:

- Share copy and redistribute the material in any medium or format
- Adapt remix, transform, and build upon the material

for any purpose, even commercially, under the following terms:



Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your



 $\label{eq:ShareAlike} Share Alike — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.$ 

[ Just-Right Consistency ]

21

[ dotScale Paris 2018-06-01 ]