

Sequential Programming for Replicated Data Stores

Nicholas V. Lewchenko¹ Arjun Radhakrishna²
Akash Gaonkar¹ Pavol Černý¹

¹University of Colorado Boulder

²Microsoft

ICFP 2019

WHY DISTRIBUTE?

Distributed architectures are required for software services that people rely on.

WHY DISTRIBUTE?

Distributed architectures are required for software services that people rely on.

Distributed applications can survive change.

Centralized services cannot.

WHY DISTRIBUTE?


Distributed architectures are required for software services that people rely on.

Distributed applications can survive change.


Centralized services cannot.




REPLICATED DATA STORES




REPLICATED DATA STORES




REPLICATED DATA STORES



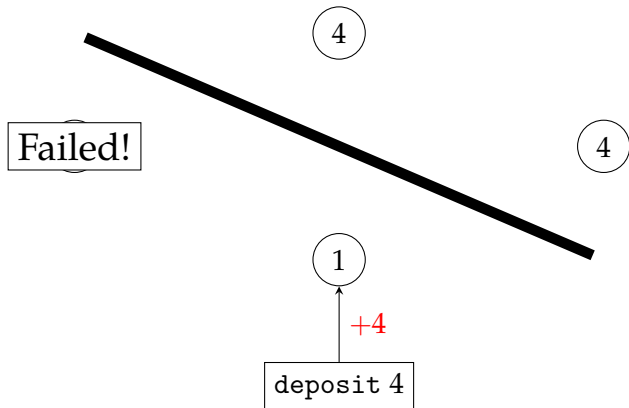
REPLICATED DATA STORES




REPLICATED DATA STORES



REPLICATED DATA STORES



REPLICATED DATA STORES



Availability

PRE- AND POST-CONDITION REASONING

Given some input to a program, what is its output?

PRE- AND POST-CONDITION REASONING

Given some input to a program, what is its output?

A formal system: Dependent Refinement Types

$$\text{sort} : (xs : \text{List}) \rightarrow \{xs' : \text{List} \mid \text{length } xs = \text{length } xs'\}$$

PRE- AND POST-CONDITION REASONING

Given some input to a program, what is its output?

A formal system: Dependent Refinement Types

$$\text{sort} : (xs : \text{List}) \rightarrow \{xs' : \text{List} \mid \text{length } xs = \text{length } xs'\}$$

How could we extend this to a replicated store operation?

Conditions on **pre-store** \rightarrow Conditions on **(Effect \times Return)**

PRE- AND POST-CONDITION REASONING

Given some input to a program, what is its output?

A formal system: Dependent Refinement Types


$\text{sort} : (xs : \text{List}) \rightarrow \{xs' : \text{List} \mid \text{length } xs = \text{length } xs'\}$

How could we extend this to a replicated store operation?


Conditions on **pre-store** \rightarrow Conditions on **(Effect \times Return)**

Consistency


REPLICATED DATA STORES



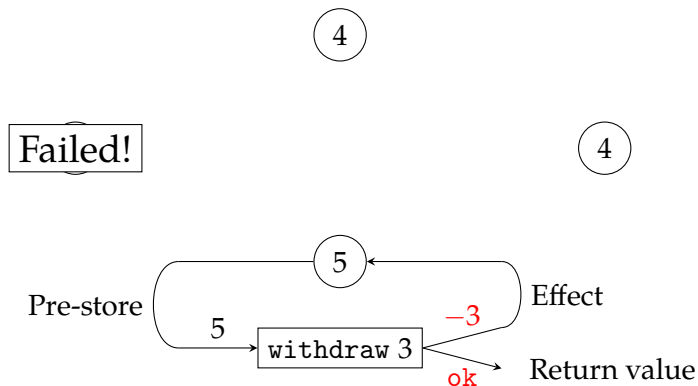
REPLICATED DATA STORES



REPLICATED DATA STORES



REPLICATED DATA STORES



Pre-stores of withdraw 3: 5✓

REPLICATED DATA STORES

Failed!

4

2


4

-3

Also a pre-store
of withdraw 3


Pre-stores of withdraw 3: $5\checkmark, 4\checkmark$

REPLICATED DATA STORES



Pre-stores of withdraw 3: $5\checkmark, 4\checkmark, 4?$

REPLICATED DATA STORES



Pre-stores of withdraw 3: $5\checkmark, 4\checkmark, 4?, 1?$

AN UNFORTUNATE CONFLICT

So how do we maintain both consistency and availability?

AN UNFORTUNATE CONFLICT

So how do we maintain both consistency and availability?

We don't...

AN UNFORTUNATE CONFLICT

So how do we maintain both consistency and availability?

We don't...

- ▶ **C**onsistency: pre/post logic can be enforced
- ▶ **A**vailability: a called operation always returns a response
- ▶ **P**artitions: the network may drop arbitrary messages

CAP Theorem: You can only have two.

AN UNFORTUNATE CONFLICT

So how do we maintain both consistency and availability?

We don't...

- ▶ Consistency: **output** depends on complete **input**
- ▶ Availability: **output** must eventually be returned
- ▶ Partitions: complete **input** might never arrive

CAP Theorem: You can only have two.

BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

Consistency and **Availability** can be balanced as needed.

BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

Consistency and **Availability** can be balanced as needed.

- ▶ **C**onsistency: **output** depends on complete **input**

BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

Consistency and **Availability** can be balanced as needed.

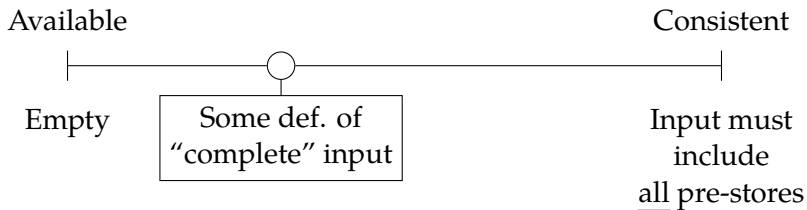
- ▶ **C**onsistency: output depends on **complete**(?) input

BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

Consistency and **Availability** can be balanced as needed.

- **C**onsistency: output depends on **complete**(?) input

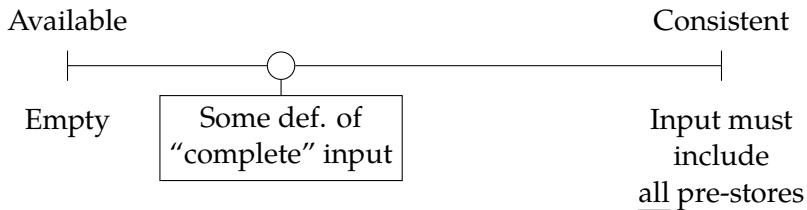


BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

Consistency and **Availability** can be balanced as needed.

- ▶ **C**onsistency: output depends on **complete**(?) input



No balance is universal!

BUT MY APP NEEDS ALL THREE!

Partitions are unavoidable for a distributed system.

Consistency and **Availability** can be balanced as needed.

- ▶ **C**onsistency: output depends on **complete**(?) input



No balance is universal!

The Special Tasks of Replicated Store Programming:

The Special Tasks of Replicated Store Programming:

1. Configure segments of application to enforce particular consistency levels.

The Special Tasks of Replicated Store Programming:

1. Configure segments of application to enforce particular consistency levels.
2. Verify that chosen consistency levels preserve desired application properties (pre/post).

The Special Tasks of Replicated Store Programming:

0. Invent a domain of useful consistency levels.
1. Configure segments of application to enforce particular consistency levels.
2. Verify that chosen consistency levels preserve desired application properties (pre/post).

Carol is a programming language that simplifies the Special Tasks.

Carol is a programming language that simplifies the Special Tasks.

- ▶ Requires only local, sequential reasoning from the programmer

Carol is a programming language that simplifies the Special Tasks.

- ▶ Requires only local, sequential reasoning from the programmer
- ▶ Supports dependent refinement type system (based on Liquid Types)

Carol is a programming language that simplifies the Special Tasks.

- ▶ Requires only local, sequential reasoning from the programmer
- ▶ Supports dependent refinement type system (based on Liquid Types)

Made possible by a novel replicated store runtime.

LET'S WRITE AN ATM APPLICATION!

First, let's deposit money.

`deposit := λn . issue (Add n) in n`

LET'S WRITE AN ATM APPLICATION!

$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$

First, let's deposit money.

`deposit := $\lambda n.$ issue (Add n) in n`

LET'S WRITE AN ATM APPLICATION!

$$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$$

First, let's deposit money.

$$\text{deposit} := \lambda n. \mathbf{issue} (\text{Add } n) \mathbf{in } n$$

Now we check our balance.

LET'S WRITE AN ATM APPLICATION!

$$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$$

First, let's deposit money.

$$\text{deposit} := \lambda n. \mathbf{issue} (\text{Add } n) \mathbf{in} n$$

Now we check our balance.

$$\text{balance} := \mathbf{query} x \mathbf{in} x$$

$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$

How do we safely withdraw money?

$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$ $\llbracket \text{Sub } n \rrbracket := \lambda x. x - n$

How do we safely withdraw money?

$\text{withdraw} := \lambda n. \text{query } x \text{ in}$
 if $n \leq x$ **then** (issue **Sub** n in n) **else** 0

$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$ $\llbracket \text{Sub } n \rrbracket := \lambda x. x - n$

$\llbracket x : \text{LEQ} \rrbracket := x \leq \text{pre-stores}$

How do we safely withdraw money?

$\text{withdraw} := \lambda n. \text{query } x : \text{LEQ in}$
 $\quad \text{if } n \leq x \text{ then (issue Sub } n \text{ in } n) \text{ else } 0$

A **query** term can be annotated with a **consistency guard**, which the runtime enforces until termination of the operation.

$\llbracket \text{Add } n \rrbracket := \lambda x. x + n$ $\llbracket \text{Sub } n \rrbracket := \lambda x. x - n$

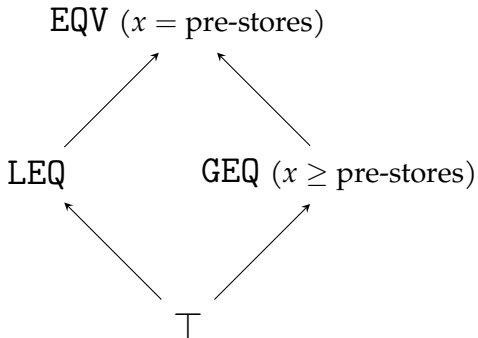
$\llbracket x : \text{LEQ} \rrbracket := x \leq \text{pre-stores}$

How do we safely withdraw money?


`withdraw := $\lambda n.$ query $x : \text{LEQ}$ in
 if $n \leq x$ then (issue Sub n in n) else 0`

A **query** term can be annotated with $\llbracket x : \text{LEQ} \rrbracket$ which the runtime enforces until term

Special Task 1 ✓




Guards: consistency level domain based on **data refinements**.



Guards: consistency level domain based on **data refinements**.


1. Provides immediately clear **data-based guarantees**.



Guards: consistency level domain based on **data refinements**.

1. Provides immediately clear **data-based guarantees**.
2. Enables **local reasoning**.

Meaning of " $x : \text{LEQ}$ " does not depend on what other operations exist.



Guards: consistency level domain based on **data refinements**.

1. Provides immediately clear **data-based guarantees**.
2. Enables **local reasoning**.

Meaning of “ $x : \text{LEQ}$ ” does not depend on what other operations exist.

Special Task 0 ✓

$$\Gamma \vdash t : \{ \mathbf{Op} \ D \ A \mid \varphi_{s,e,r} \}$$

D is a **Conflict-Aware Replicated Datatype (CARD)** that defines the effects and guards of a store.

$$\Gamma \vdash t : \{ \mathbf{Op} \ D \ A \mid \varphi_{s,e,r} \}$$

D is a **Conflict-Aware Replicated Datatype (CARD)** that defines the effects and guards of a store.

$$\begin{aligned} \vdash \text{deposit} & : (n : \text{Nat}) \rightarrow \{ \mathbf{Op} \ \text{Ctr} \ \text{Nat} \mid \llbracket \mathbf{e} \rrbracket (\mathbf{s}) = \mathbf{s} + n \} \\ \vdash \text{balance} & : \{ \mathbf{Op} \ \text{Ctr} \ \text{Int} \mid \mathbf{e} = \text{id} \} \end{aligned}$$

$$\Gamma \vdash t : \{ \mathbf{Op} \ D \ A \mid \varphi_{s,e,r} \}$$

D is a **Conflict-Aware Replicated Datatype (CARD)** that defines the effects and guards of a store.

$$\begin{aligned} \vdash \text{deposit} & : (n : \text{Nat}) \rightarrow \{ \mathbf{Op} \ \text{Ctr} \ \text{Nat} \mid \llbracket \mathbf{e} \rrbracket (\mathbf{s}) = \mathbf{s} + n \} \\ \vdash \text{balance} & : \{ \mathbf{Op} \ \text{Ctr} \ \text{Int} \mid \mathbf{e} = \text{id} \} \end{aligned}$$

Everything is an operation!

$$\vdash 5 : \{ \mathbf{Op} \ D \ \text{Int} \mid \mathbf{e} = \text{id} \wedge \mathbf{r} = 5 \}$$

VERIFYING WITHDRAW

$$\varphi := (\mathbf{s} \geq 0 \Rightarrow \llbracket \mathbf{e} \rrbracket(\mathbf{s}) \geq 0) \wedge (\mathbf{r} = \mathbf{s} - \llbracket \mathbf{e} \rrbracket(\mathbf{s}))$$

VERIFYING WITHDRAW

$$\varphi := (\mathbf{s} \geq 0 \Rightarrow \llbracket \mathbf{e} \rrbracket(\mathbf{s}) \geq 0) \wedge (\mathbf{r} = \mathbf{s} - \llbracket \mathbf{e} \rrbracket(\mathbf{s}))$$

1. Account never goes below zero

VERIFYING WITHDRAW

$$\varphi := (\mathbf{s} \geq 0 \Rightarrow \llbracket \mathbf{e} \rrbracket(\mathbf{s}) \geq 0) \wedge (\mathbf{r} = \mathbf{s} - \llbracket \mathbf{e} \rrbracket(\mathbf{s}))$$

1. Account never goes below zero
2. Value returned to caller is operation's real effect on store

VERIFYING WITHDRAW

$$\varphi := (\mathbf{s} \geq 0 \Rightarrow \llbracket \mathbf{e} \rrbracket(\mathbf{s}) \geq 0) \wedge (\mathbf{r} = \mathbf{s} - \llbracket \mathbf{e} \rrbracket(\mathbf{s}))$$

1. Account never goes below zero
2. Value returned to caller is operation's real effect on store

`withdraw := λn . query x : LEQ in
 if $n \leq x$ then (issue Sub n in n) else 0`

$$\frac{\Gamma \vdash \text{LEQ} : \text{Guard}(\text{Ctr}) \quad \Gamma, x : \{ \mathbf{Op} \text{ Ctr Int} \mid \mathbf{r} \leq \mathbf{s} \} \vdash \text{if} \dots : \{ \mathbf{Op} \text{ Ctr Nat} \mid \varphi \}}{\Gamma \vdash \text{query } x : \text{LEQ in if} \dots : \{ \mathbf{Op} \text{ Ctr Nat} \mid \varphi \}}$$

VERIFYING WITHDRAW

$$\varphi := (\mathbf{s} \geq 0 \Rightarrow \llbracket \mathbf{e} \rrbracket(\mathbf{s}) \geq 0) \wedge (\mathbf{r} = \mathbf{s} - \llbracket \mathbf{e} \rrbracket(\mathbf{s}))$$

1. Account never goes below zero
2. Value returned to caller is operation's real effect on store

`withdraw := λn . query x : LEQ in
if $n \leq x$ then (issue Sub n in n) else 0`

$\Gamma \vdash \text{LEQ} : \text{Guard}(\text{Ctr})$
 $\Gamma, x : \{ \mathbf{Op} \text{ Ctr Int} \mid \mathbf{r} \leq \mathbf{s} \} \vdash \text{if} \dots : \{ \mathbf{Op} \text{ Ctr Nat} \mid \varphi \}$

 $\Gamma \vdash \text{query } x : \text{LEQ in if} \dots :$

Special Task 2 ✓

SO WHO'S PAYING FOR THIS?

Programmer only needs local, sequential reasoning...

SO WHO'S PAYING FOR THIS?


Programmer only needs local, sequential reasoning...

But runtime needs more.

SO WHO'S PAYING FOR THIS?

Programmer only needs local, sequential reasoning...

But runtime needs more.



Accords tell the runtime which effects are safe during a query.


Theorem: If $\{\text{guard}\}$ is in accord with $\{\text{effect}\}$, then a query using $\{\text{guard}\}$ can safely return without including $\{\text{effect}\}$.

WHAT HAVE WE GAINED?

Accords are **more reusable** and involve **less code** than full-operation concurrent verification.

WHAT HAVE WE GAINED?


Accords are **more reusable** and involve **less code** than full-operation concurrent verification.



(conc. ver.)


WHAT HAVE WE GAINED?

Accords are **more reusable** and involve **less code** than full-operation concurrent verification.



(conc. ver.)

vs.




(conc. ver.)

(seq. ver.)

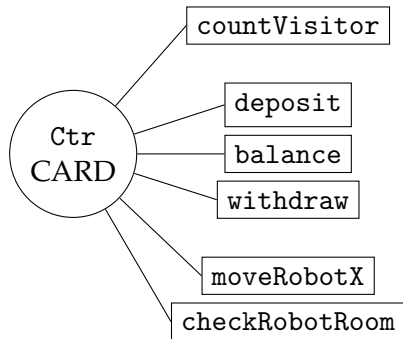
WHAT HAVE WE GAINED?

Accords are **more reusable** and involve **less code** than full-operation concurrent verification.



(conc. ver.)

vs.



(conc. ver.)

(seq. ver.)

FUTURE WORK: ADVANCED RUNTIMES

Preserving semantics

- ▶ Effects or guards—who gets right-of-way?
- ▶ Contention management

Preserving semantics

- ▶ Effects or guards—who gets right-of-way?
- ▶ Contention management

Extending semantics/language

- ▶ Direct messages for safety-preserving side deals.

- ▶ **Consistency guards** isolate programmer from global, concurrent reasoning—operations behave according to local, sequential rules

IN CONCLUSION

- ▶ **Consistency guards** isolate programmer from global, concurrent reasoning—operations behave according to local, sequential rules
- ▶ The local, sequential reasoning is formalized by a **dependent refinement type** system

IN CONCLUSION

- ▶ **Consistency guards** isolate programmer from global, concurrent reasoning—operations behave according to local, sequential rules
- ▶ The local, sequential reasoning is formalized by a **dependent refinement type** system
- ▶ **Accords** statically capture the concurrent knowledge needed to run many not-yet-written applications

IN CONCLUSION

- ▶ **Consistency guards** isolate programmer from global, concurrent reasoning—operations behave according to local, sequential rules
- ▶ The local, sequential reasoning is formalized by a **dependent refinement type** system
- ▶ **Accords** statically capture the concurrent knowledge needed to run many not-yet-written applications

- ▶ Haskell DSL and runtime implementation:
`https://github.com/cuplv/discard`