

# Proving Weakly Consistent Applications Correct


Alexey Gotsman

IMDEA Software Institute, Madrid, Spain


*Joint work with*

*Hongseok Yang (Oxford), Carla Ferreira (U Nova Lisboa),  
Mahsa Najafzadeh, Marc Shapiro (INRIA)*

# Eventually consistent databases




- **No synchronisation:** process an update locally, propagate effects to other replicas later
- **Weakens consistency:** deposit seen with a delay




balance = 100

balance  $\geq$  0



balance = 100




balance = 100

withdraw(100) : ✓

balance = 0


balance  $\geq$  0



balance = 100

withdraw(100) : ✓

balance = 0




balance = 100

withdraw(100) : ✓

balance = 0

balance  $\geq$  0





balance = 100

withdraw(100) : ✓

balance = 0

balance = -100






balance = 100


withdraw(100) : ✓

balance = 0



deposit(100)

balance  $\geq$  0





balance = 100

withdraw(100) : ✓

balance = 0

balance = -100






balance = 100

withdraw(100) : ✓

balance = 0



deposit(100)

balance  $\geq$  0



balance = 100

withdraw(100) : ✓

balance = 0

balance = -100

### Tune consistency:

- Withdrawals strongly consistent
- Deposits eventually consistent

# Consistency choices

- Databases with multiple consistency levels:
  - ▶ Commercial: Amazon DynamoDB, Basho Riak, Microsoft DocumentDB
  - ▶ Research: Li<sup>+</sup> OSDI'12; Terry<sup>+</sup> SOSP'13; Balegas<sup>+</sup> EuroSys'15...
- Pay for stronger semantics with latency, possible unavailability and money
- Hard to figure out the minimum consistency necessary to maintain correctness -  
proof rule and tool



# Consistency model

- **Generic model** - not implemented, but can encode many existing models that are:  
*RedBlue consistency [Li<sup>+</sup> 2012],  
reservation locks [Balegas<sup>+</sup> 2015],  
parallel snapshot isolation [Sovran<sup>+</sup> 2011], ...*
- **Causal consistency** as a baseline: observe an update → observe the updates it depends on
- A construct for **strengthening consistency** on demand

# Operation semantics



$\sigma$




$\llbracket op \rrbracket_{val}$

Replica states:  $\sigma \in \text{State}$

Return value:  $\llbracket op \rrbracket_{val} \in \text{State} \rightarrow \text{Value}$

# Operation semantics




Replica states:  $\sigma \in \text{State}$

Return value:  $[[op]]_{val} \in \text{State} \rightarrow \text{Value}$

Effector:  $[[op]]_{eff} \in \text{State} \rightarrow (\text{State} \rightarrow \text{State})$

# Operation semantics




Replica states:  $\sigma \in \text{State}$

Return value:  $[[op]]_{val} \in \text{State} \rightarrow \text{Value}$

Effector:  $[[op]]_{eff} \in \text{State} \rightarrow (\text{State} \rightarrow \text{State})$

# Operation semantics




Replica states:  $\sigma \in \text{State}$

Return value:  $[[op]]_{val} \in \text{State} \rightarrow \text{Value}$

Effector:  $[[op]]_{eff} \in \text{State} \rightarrow (\text{State} \rightarrow \text{State})$

# Operation semantics




Replica states:  $\sigma \in \text{State}$

Return value:  $[[op]]_{val} \in \text{State} \rightarrow \text{Value}$

Effector:  $[[op]]_{eff} \in \text{State} \rightarrow (\text{State} \rightarrow \text{State})$

# Operation semantics




State =  $Z$

$$[[balance()]]_{val}(\sigma) = \sigma$$

$$[[balance()]]_{eff}(\sigma) = \lambda\sigma. \sigma$$


# Operation semantics



$$[[deposit(100)]]_{eff}(\sigma) = \lambda\sigma'. (\sigma' + 100)$$

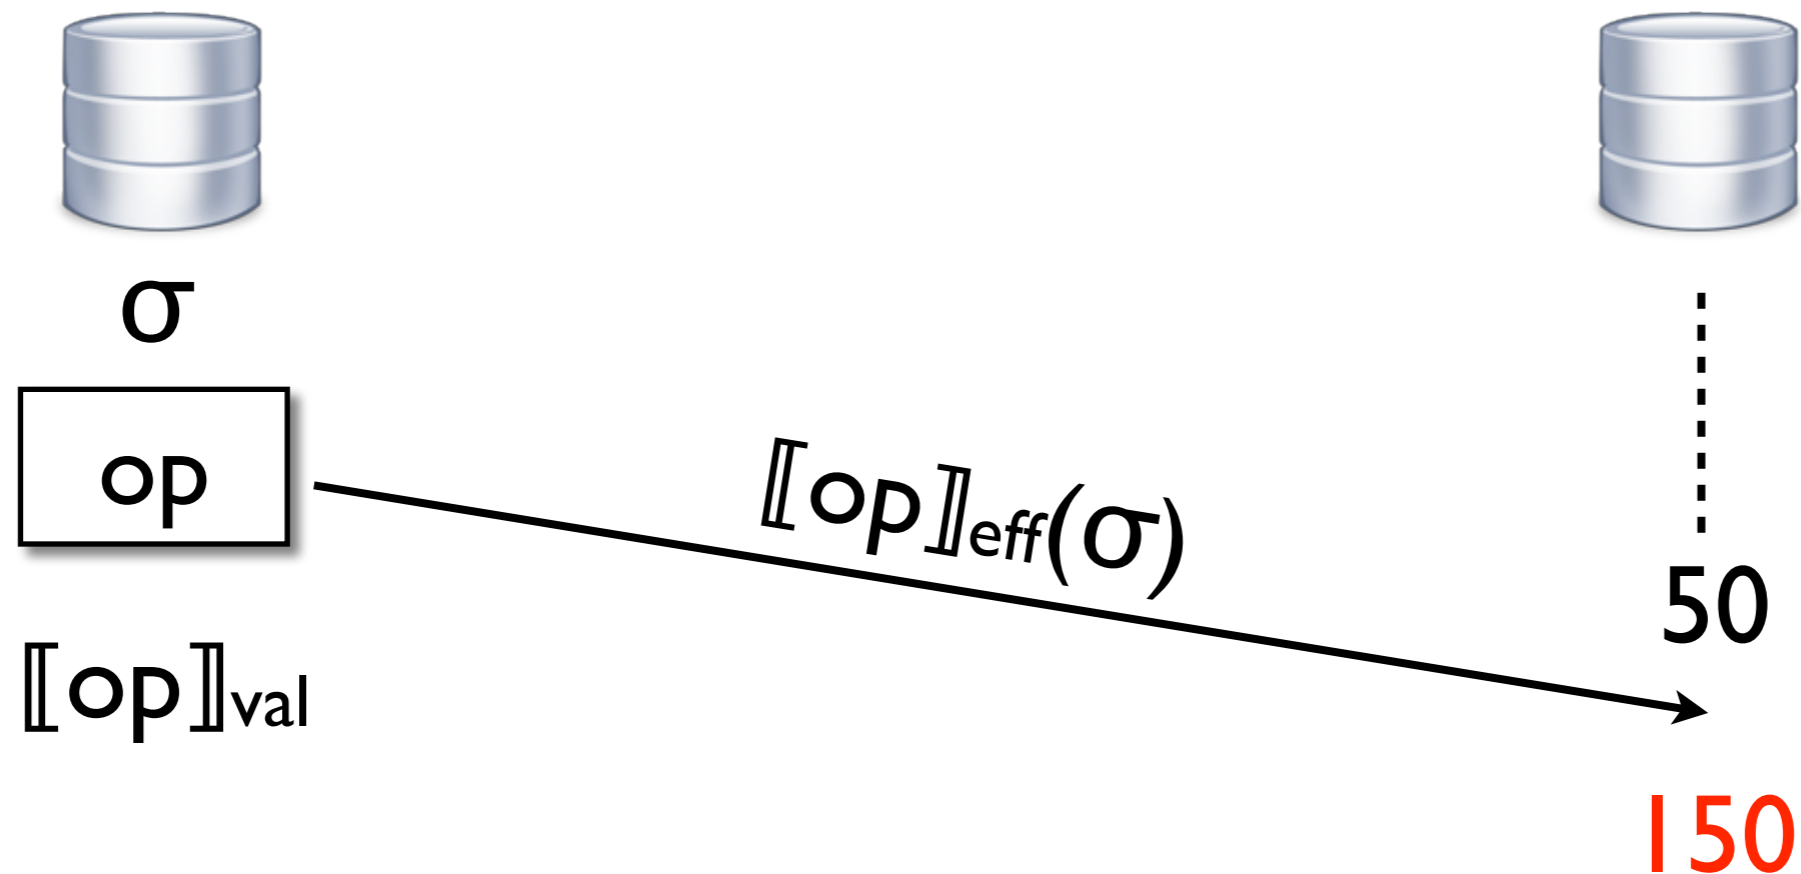


# Operation semantics



$$\llbracket deposit(100) \rrbracket_{eff}(\sigma) = \lambda \sigma'. (\sigma' + 100)$$

# Operation semantics




$$[[deposit(100)]]_{eff}(\sigma) = \lambda\sigma'. (\sigma' + 100)$$

# Ensuring eventual consistency

- Effectors have to **commute**
- **Eventual consistency**: replicas receiving the same messages in different orders end up in the same state
- **Replicated data types** [Shapiro<sup>+</sup> 2011]: ready-made commutative implementations

$$\llbracket \text{deposit}(100) \rrbracket_{\text{eff}}(\sigma) = \lambda \sigma'. (\sigma' + 100)$$


# Operation semantics



$[[withdraw(100)]]_{eff}(\sigma) =$

if  $\sigma \geq 100$  then  $(\lambda\sigma'. \sigma' - 100)$  else  $(\lambda\sigma'. \sigma')$


# Operation semantics



$[[withdraw(100)]]_{eff}(\sigma) =$

if  $\sigma \geq 100$  then  $(\lambda\sigma'. \sigma' - 100)$  else  $(\lambda\sigma'. \sigma')$


# Operation semantics



$\llbracket \text{withdraw}(100) \rrbracket_{\text{eff}}(\sigma) =$

if  $\sigma \geq 100$  then  $(\lambda \sigma'. \sigma' - 100)$  else  $(\lambda \sigma'. \sigma')$

# Operation semantics



$[[withdraw(100)]]_{eff}(\sigma) =$

if  $\sigma \geq 100$  then  $(\lambda\sigma'. \sigma' - 100)$  else  $(\lambda\sigma'. \sigma')$



balance = 100

withdraw(100) : ✓

balance = 0

$\lambda\sigma'.\sigma' - 100$



balance = 100

withdraw(100) : ✓

balance = 0

$\llbracket \text{withdraw}(100) \rrbracket_{\text{eff}}(\sigma) =$

if  $\sigma \geq 100$  then  $(\lambda\sigma'.\sigma' - 100)$  else  $(\lambda\sigma'.\sigma')$





balance = 100

withdraw(100) : ✓

balance = 0

$\lambda\sigma'.\sigma' - 100$



balance = 100

withdraw(100) : ✓

balance = 0

balance = -100

$\llbracket \text{withdraw}(100) \rrbracket_{\text{eff}}(\sigma) =$

if  $\sigma \geq 100$  then  $(\lambda\sigma'.\sigma' - 100)$  else  $(\lambda\sigma'.\sigma')$

# Strengthening consistency

Token system  $\approx$  locks on steroids:

- Token =  $\{\tau_1, \tau_2, \dots\}$
- Symmetric conflict relation  $\bowtie \subseteq \text{Token} \times \text{Token}$

# Strengthening consistency

Token system  $\approx$  locks on steroids:

- $\text{Token} = \{\tau_1, \tau_2, \dots\}$
- Symmetric conflict relation  $\bowtie \subseteq \text{Token} \times \text{Token}$

Example - mutual exclusion lock:

$\text{Token} = \{\tau\}; \tau \bowtie \tau$

# Strengthening consistency

Token system  $\approx$  locks on steroids:

- $\text{Token} = \{\tau_1, \tau_2, \dots\}$
- Symmetric conflict relation  $\bowtie \subseteq \text{Token} \times \text{Token}$

Example - mutual exclusion lock:

$\text{Token} = \{\tau\}; \tau \bowtie \tau$

Each operation associated with a set of tokens:

$\llbracket \text{op} \rrbracket_{\text{tok}} \in \text{State} \rightarrow \mathcal{P}(\text{Token})$

# Operations associated with conflicting tokens cannot be unaware of each other



balance = 100

$T \times T$



balance = 100

withdraw(100) : ✓

{T}

# Operations associated with conflicting tokens cannot be unaware of each other



balance = 100

withdraw(100) : ✓

{T}

T ✕ T



balance = 100



withdraw(100) : ?

{T}

Anything I don't know about?

# Operations associated with conflicting tokens cannot be unaware of each other



balance = 100


$T \times T$



balance = 100

withdraw(100) : ✓

{T}



balance = 0



withdraw(100) : ?

{T}

# Operations associated with conflicting tokens cannot be unaware of each other



balance = 100


$T \times T$



balance = 100

withdraw(100) : ✓

{T}



balance = 0

withdraw(100) : ✗

{T}



# Operations associated with conflicting tokens cannot be unaware of each other



balance = 100

withdraw(100) : ✓

{T}



deposit(100)

∅

No synchronisation

T ✗ T



balance = 100



balance = 0



withdraw(100) : ✗

{T}

Do we always have  $I = (\text{balance} \geq 0)$ ?



balance = 100

withdraw(100) : ✓

{T}



deposit(100)

∅

No synchronisation

T ✕ T



balance = 100

balance = 0

withdraw(100) : ✗

{T}



$\sigma \in I$

op

$\llbracket \text{op} \rrbracket_{\text{eff}}(\sigma)$




⋮

$\sigma'$

$\llbracket \text{op} \rrbracket_{\text{eff}}(\sigma)(\sigma') \in I ?$


Effect applied in a different state!



$[[\text{op}]]_{\text{eff}}(\sigma) = \text{if } P(\sigma) \text{ then } f(\sigma) \text{ else if...}$


$[[\text{withdraw}(100)]]_{\text{eff}}(\sigma) =$

$\text{if } \sigma \geq 100 \text{ then } (\lambda\sigma'. \sigma' - 100) \text{ else } (\lambda\sigma'. \sigma')$



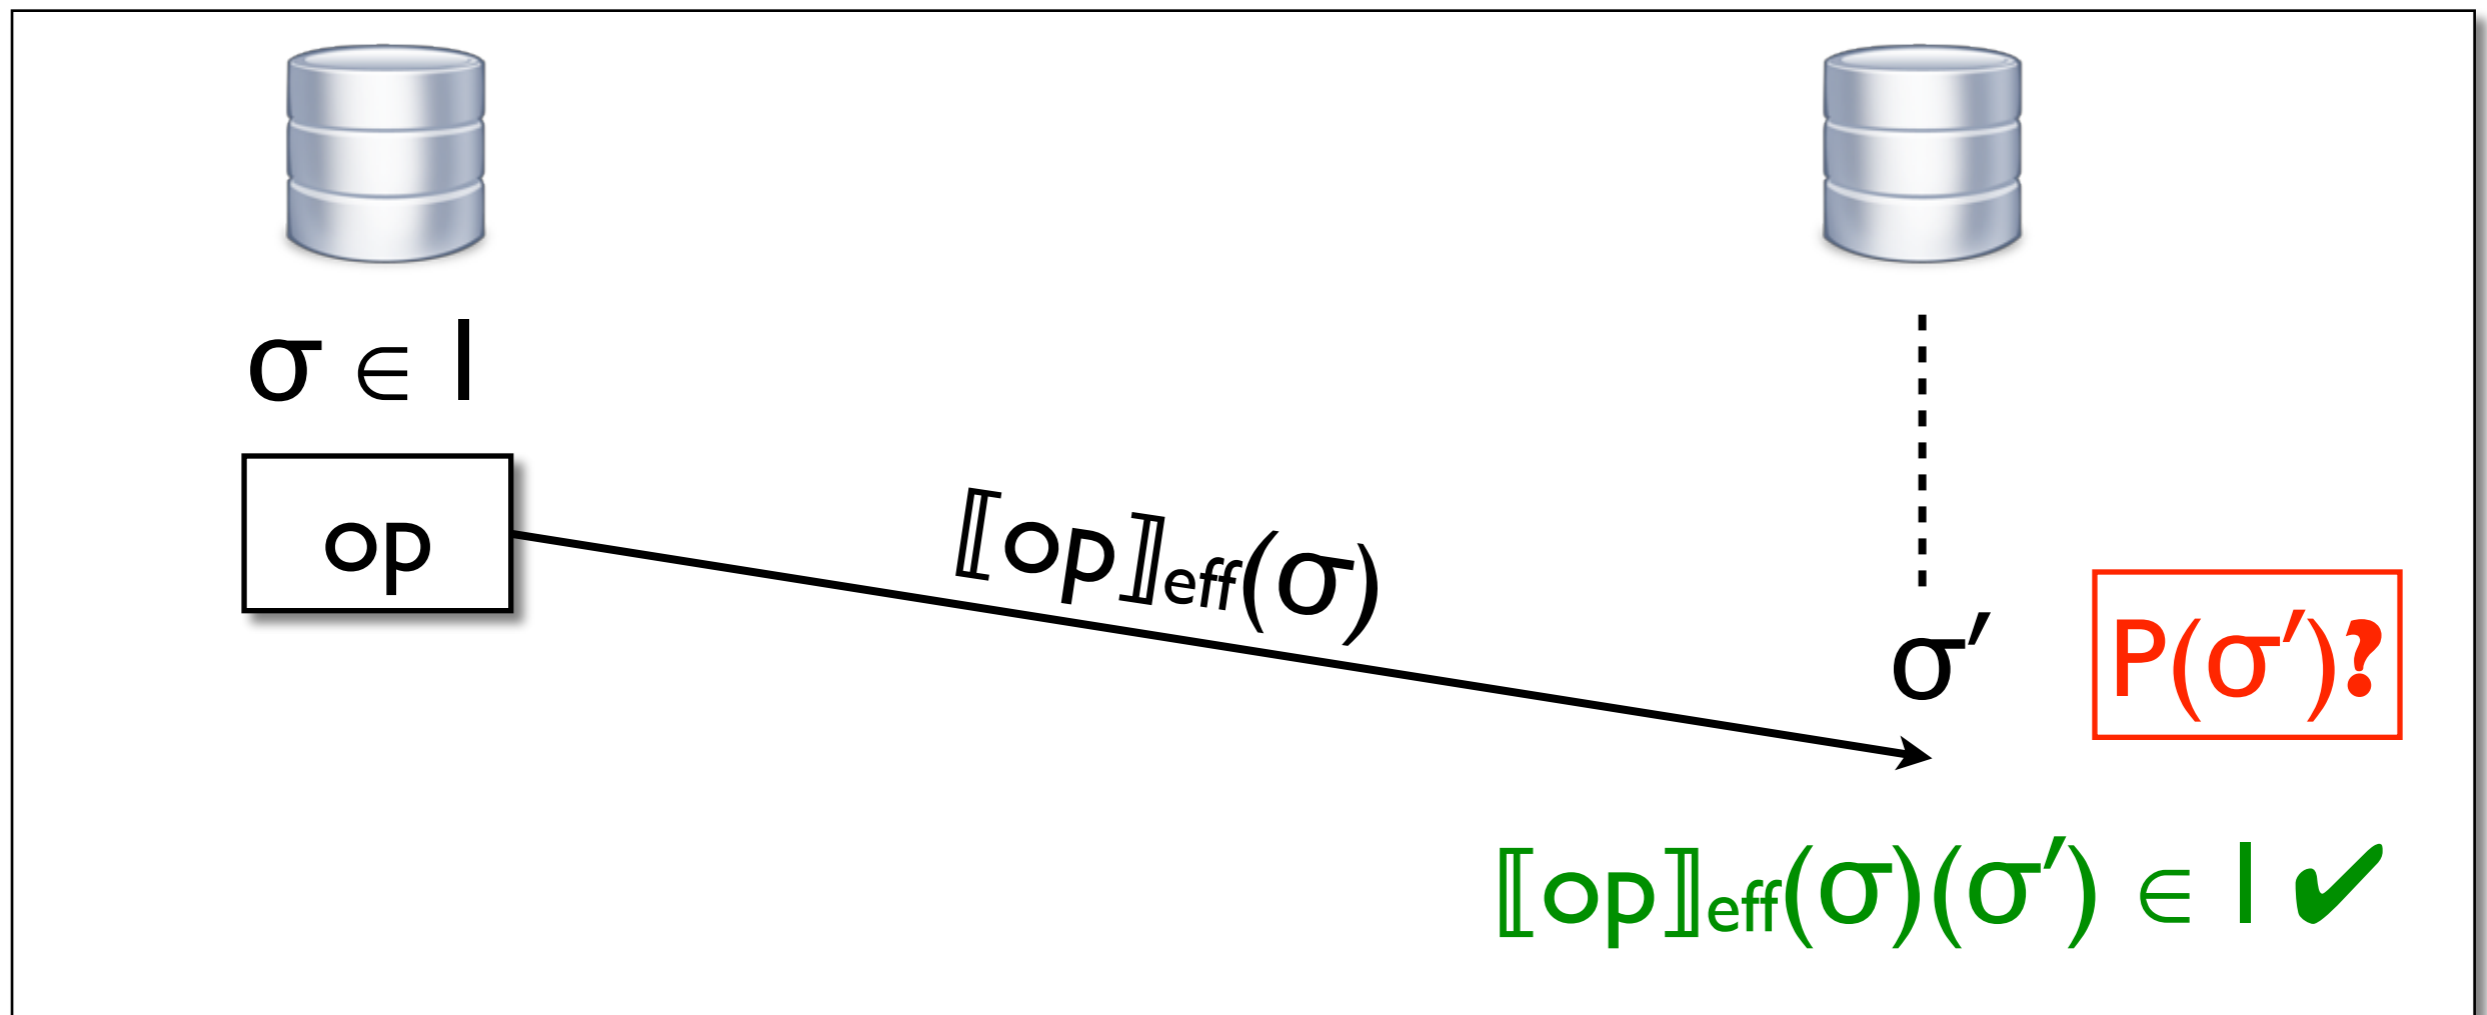
$[[\text{op}]]_{\text{eff}}(\sigma) = \text{if } P(\sigma) \text{ then } f(\sigma) \text{ else if...}$

1. **Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$




$\llbracket op \rrbracket_{eff}(\sigma) = \text{if } P(\sigma) \text{ then } f(\sigma) \text{ else if...}$

- Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$



$\llbracket op \rrbracket_{eff}(\sigma) = \text{if } P(\sigma) \text{ then } f(\sigma) \text{ else if...}$

- Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$



$\llbracket op \rrbracket_{eff}(\sigma) = \text{if } P(\sigma) \text{ then } f(\sigma) \text{ else if...}$

1. **Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$
2. **Precondition stability:**  $P$  will hold when  $f$  is applied at any replica



# CISE tool: 'Cause I'm Strong Enough

Discharges proof obligations using Z3 SMT solver

By Mahsa Najafzadeh (UPMC & INRIA)

$\llbracket \text{op} \rrbracket_{\text{eff}}(\sigma) = \text{if } P(\sigma) \text{ then } f(\sigma) \text{ else if...}$

1. **Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$
2. **Precondition stability:**  $P$  will hold when  $f$  is applied at any replica

```
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance" }, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```

```
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance" }, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```

I. **Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$

```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```



I. **Effector safety:**  $f$  preserves  $I$  when executed in any state satisfying  $P$



```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```

## 2. Precondition stability: $P$ is preserved by concurrent operations

```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance" }, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```

## 2. Precondition stability: $P$ is preserved by concurrent operations

```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance" }, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```



## 2. Precondition stability: $P$ is preserved by concurrent operations

Bug: concurrent withdrawals may violate the invariant

```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = {"Int balance" }, type = XPR.Type.ARGUMENT)
    @XPR(value = "true", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    public static class Withdraw extends AnnotatedOperation { }
}
```



## 2. Precondition stability: $P$ is preserved by concurrent operations

Add a token  
restricting concurrency

```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {

    @XPR(value = "Int balance", type = XPR.Type.ARGUMENT)
    @XPR(value = "Bool token", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance + 100", type = XPR.Type.EFFECT)
    public static class Deposit extends AnnotatedOperation { }

    @XPR(value = {"Int balance", "Bool token"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    @XPR(value = "token := true", type = XPR.Type.Token)
    public static class Withdraw extends AnnotatedOperation { }
}
```

## 2. Precondition stability: $P$ is preserved by concurrent operations



```
Account.java - Static10012 - [~/Documents/workspace-static/Static10012]
Account
SequentialTest
Account.java x
@XPR("Int balance")
@XPR(value = "balance >= 0", type = XPR.Type.INVARIANT)
@Op(Account.Deposit.class)
@Op(Account.Withdraw.class)
public class Account extends AnnotatedSchema {
    @XPR(value = "Int balance", type = XPR.Type.ARGUMENT)
    @XPR(value = "Bool token", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance + 100", type = XPR.Type.EFFECT)
    public class Deposit extends AnnotatedOperation { }
    @XPR(value = {"Int balance", "Bool token"}, type = XPR.Type.ARGUMENT)
    @XPR(value = "balance >= 100", type = XPR.Type.PRECONDITION)
    @XPR(value = "balance := balance - 100", type = XPR.Type.EFFECT)
    @XPR(value = "token := true", type = XPR.Type.Token)
    public static class Withdraw extends AnnotatedOperation { }
}
```

Add a token  
restricting concurrency

"Bool token"

"token := true"

# Conclusion

- **First proof rule and tool** for proving invariants of weakly consistent applications
- **Case studies:** fragments of web applications, replicated file system in progress
- **Future work:** other consistency models, automatic inference of consistency levels