



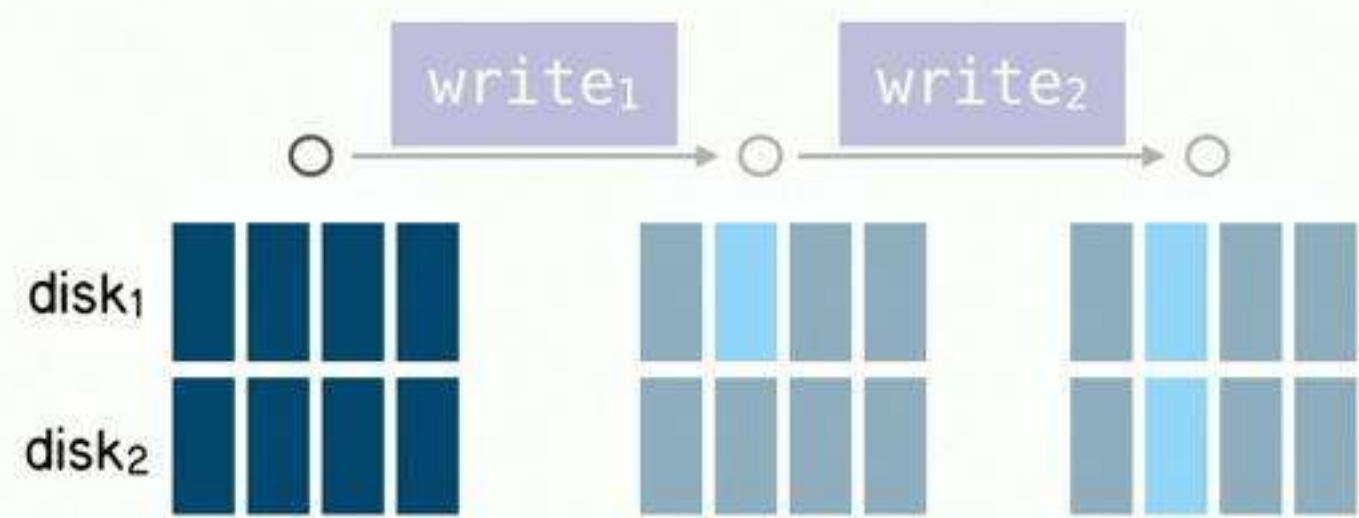
# Argosy: Verifying layered storage systems with recovery refinement

**Tej Chajed, Joseph Tassarotti, Frans Kaashoek, Nickolai Zeldovich**

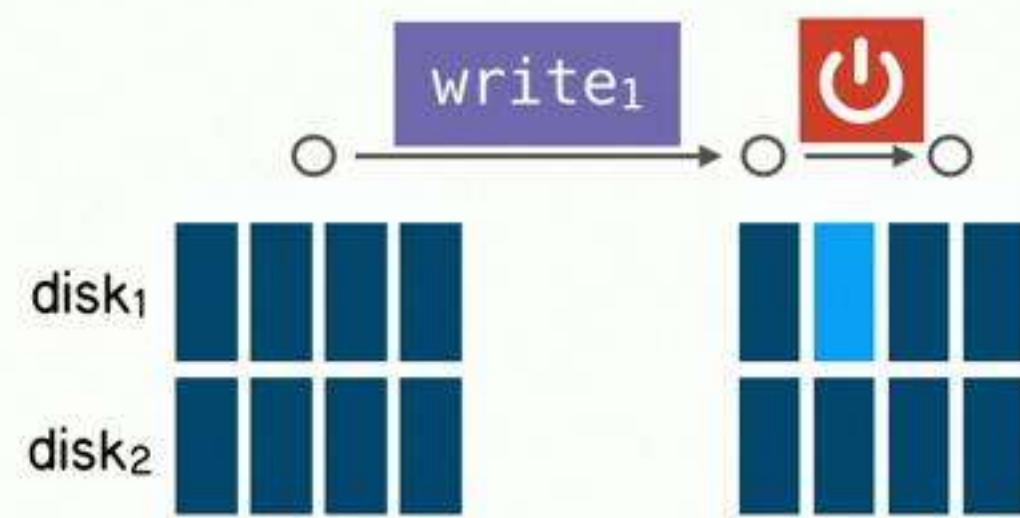
MIT



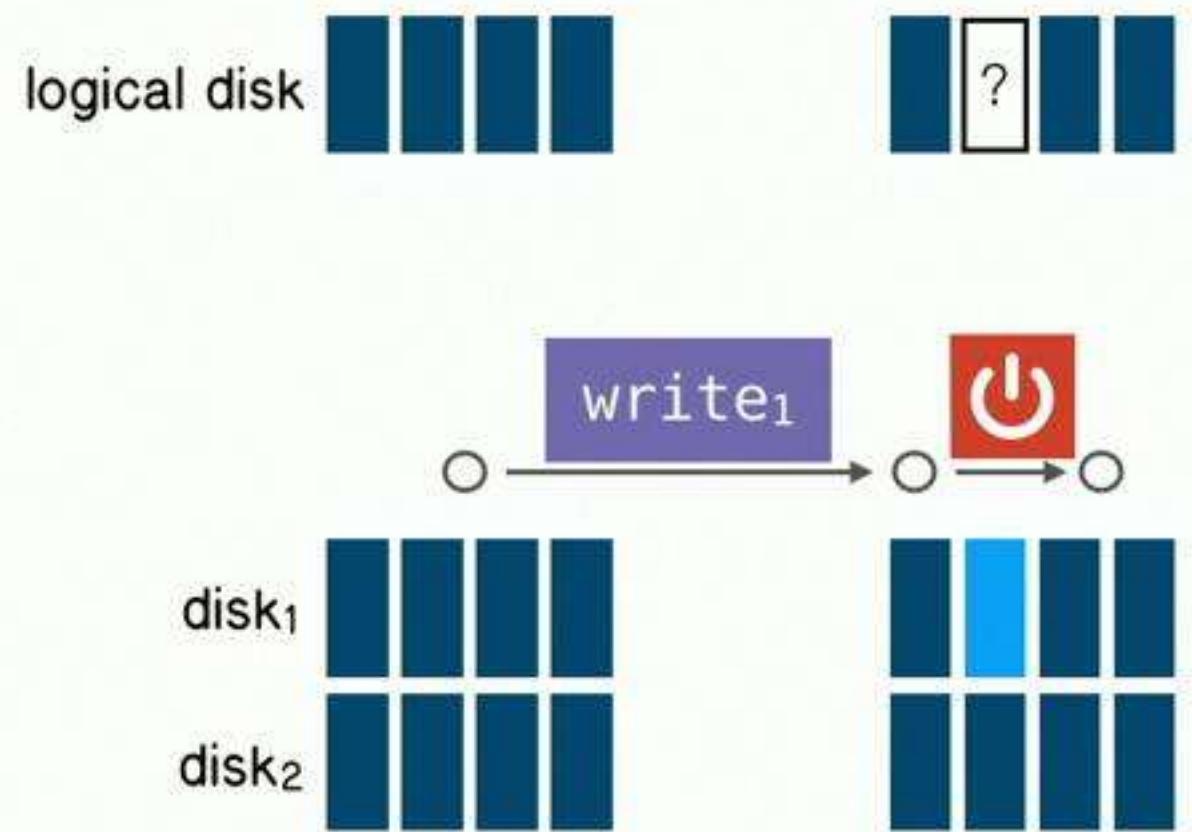
Bob writes a replication system



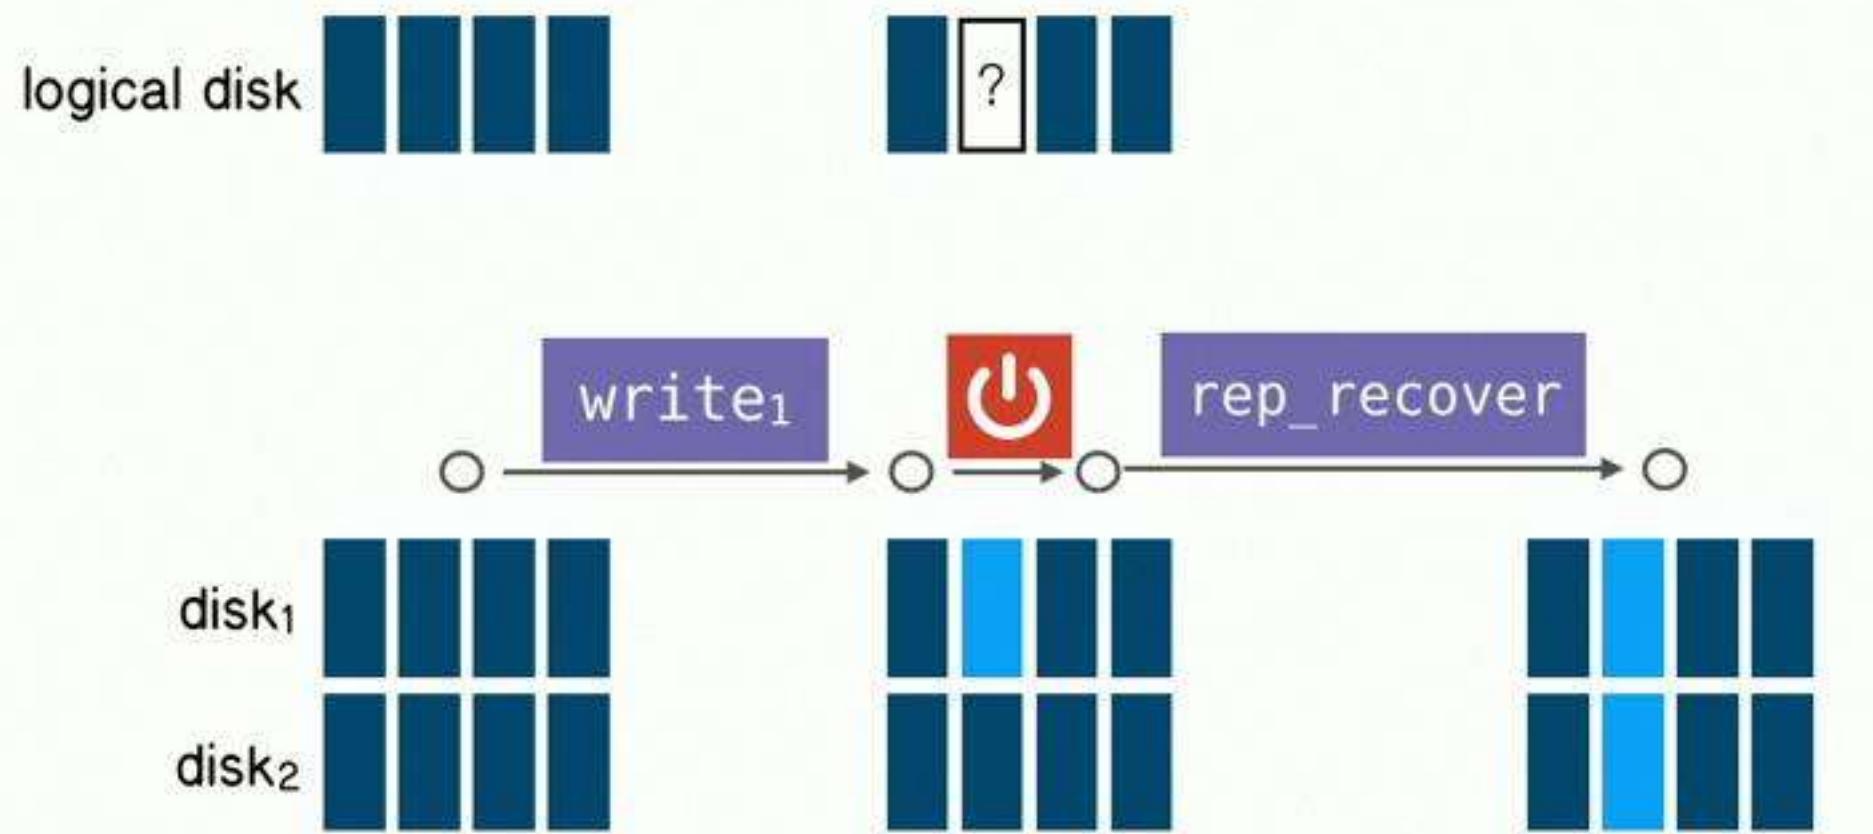
Bob writes a replication system



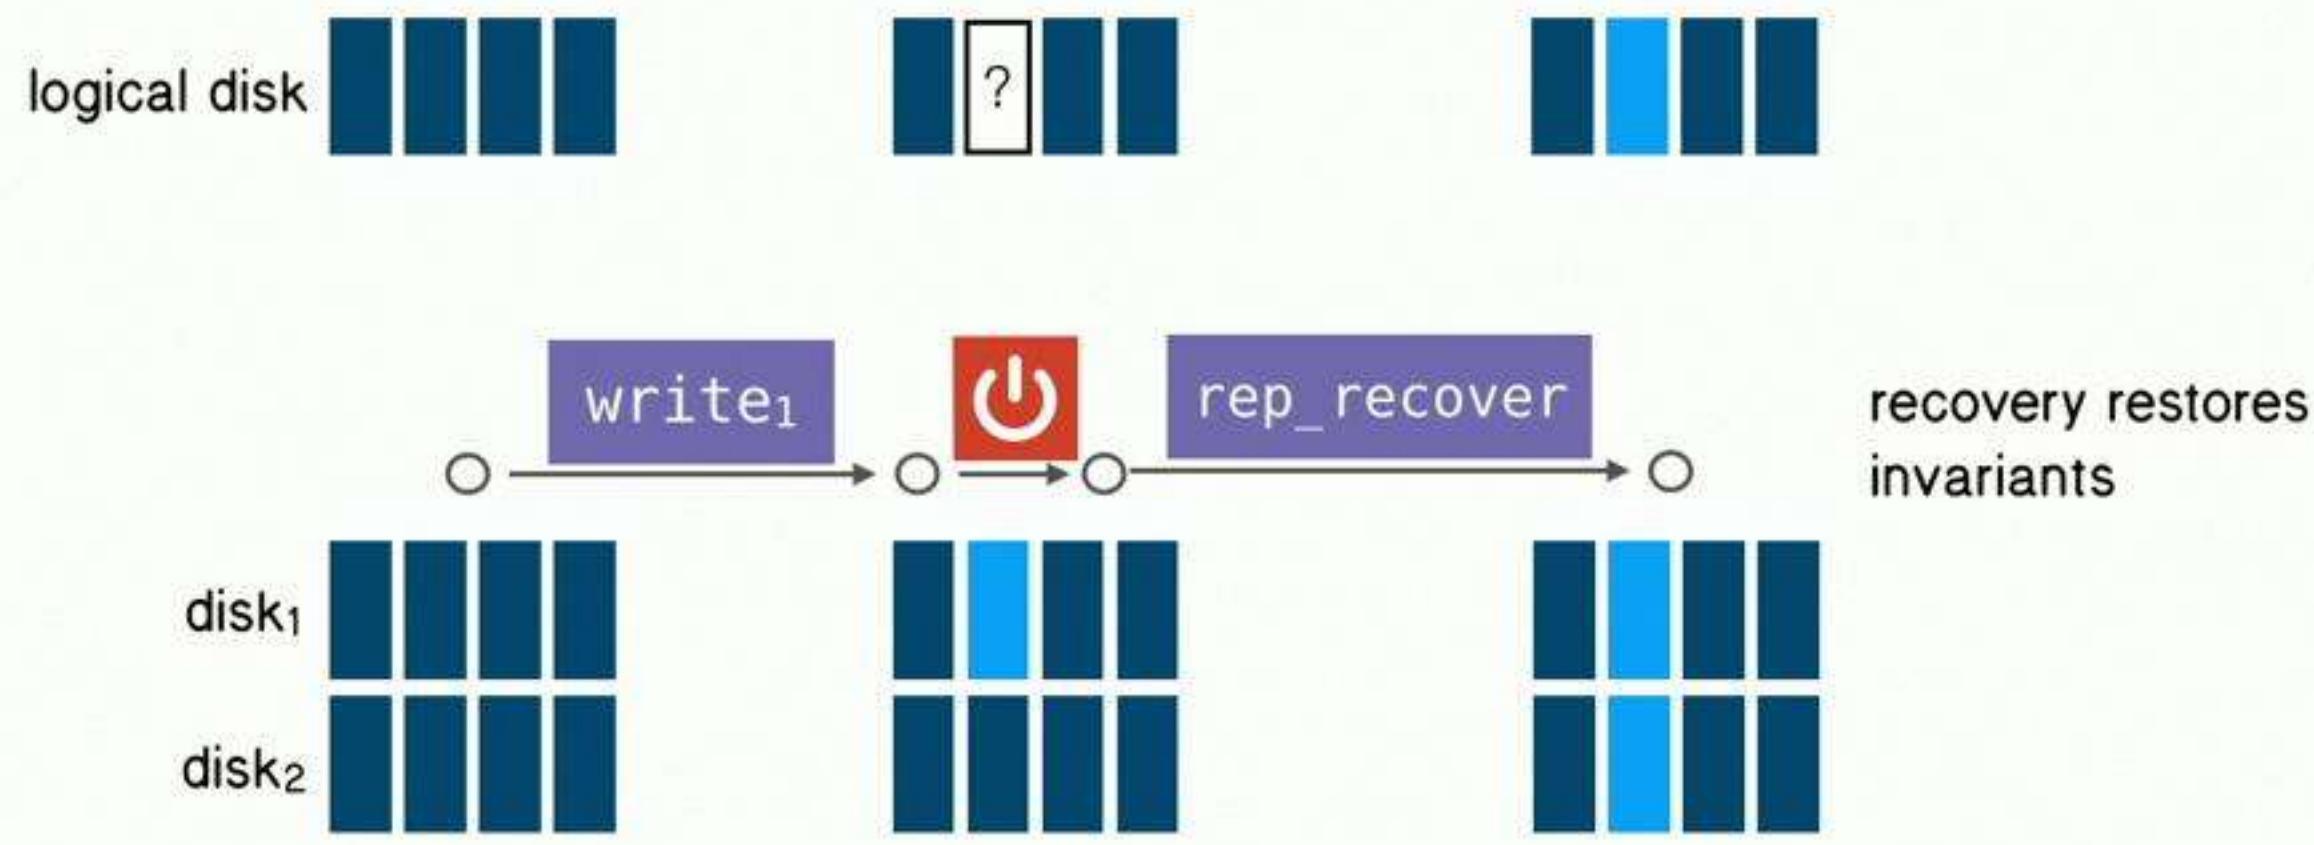
Bob writes a replication system



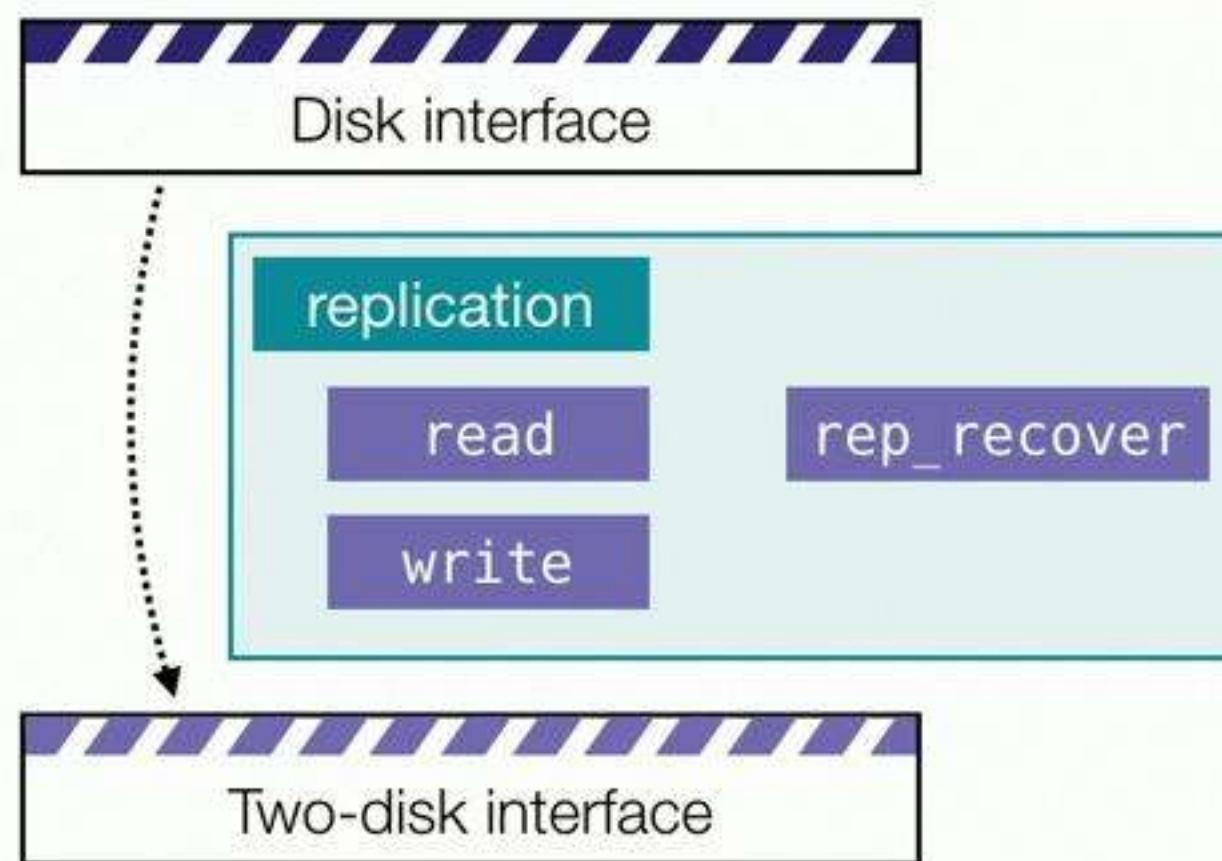
Bob writes a replication system



Bob writes a replication system and implements its recovery procedure

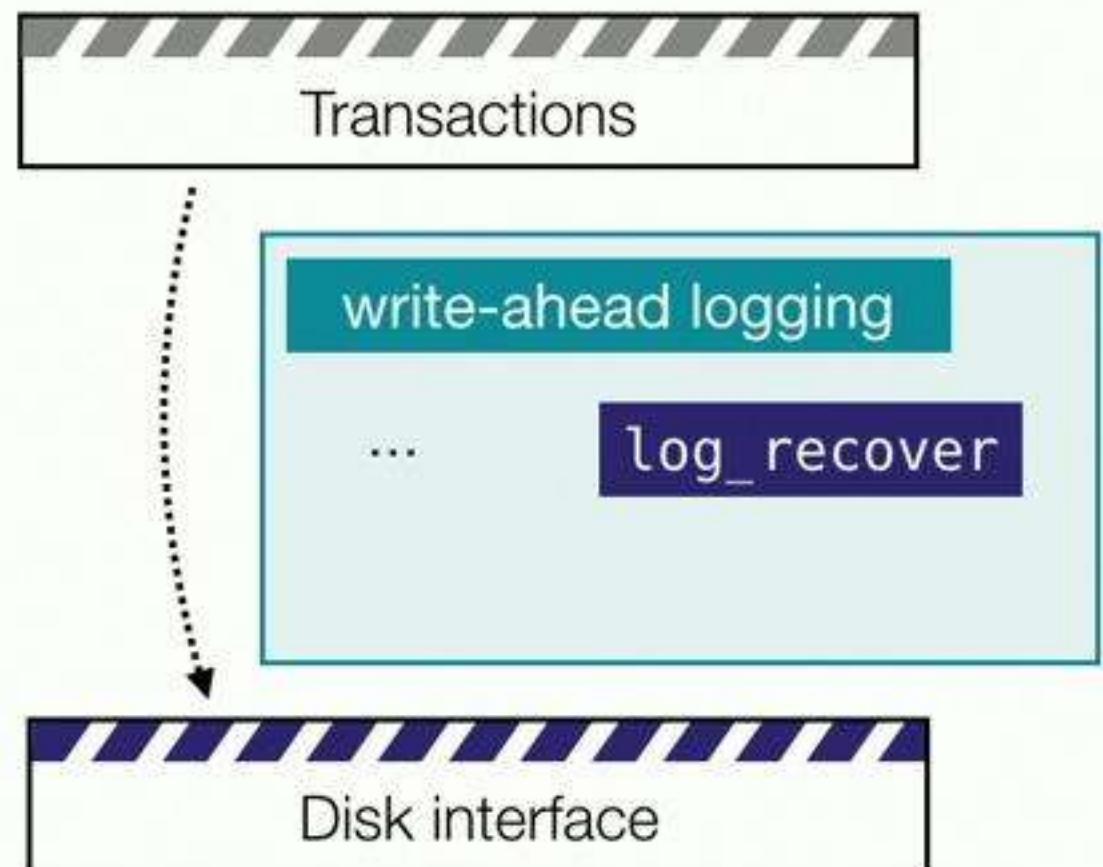


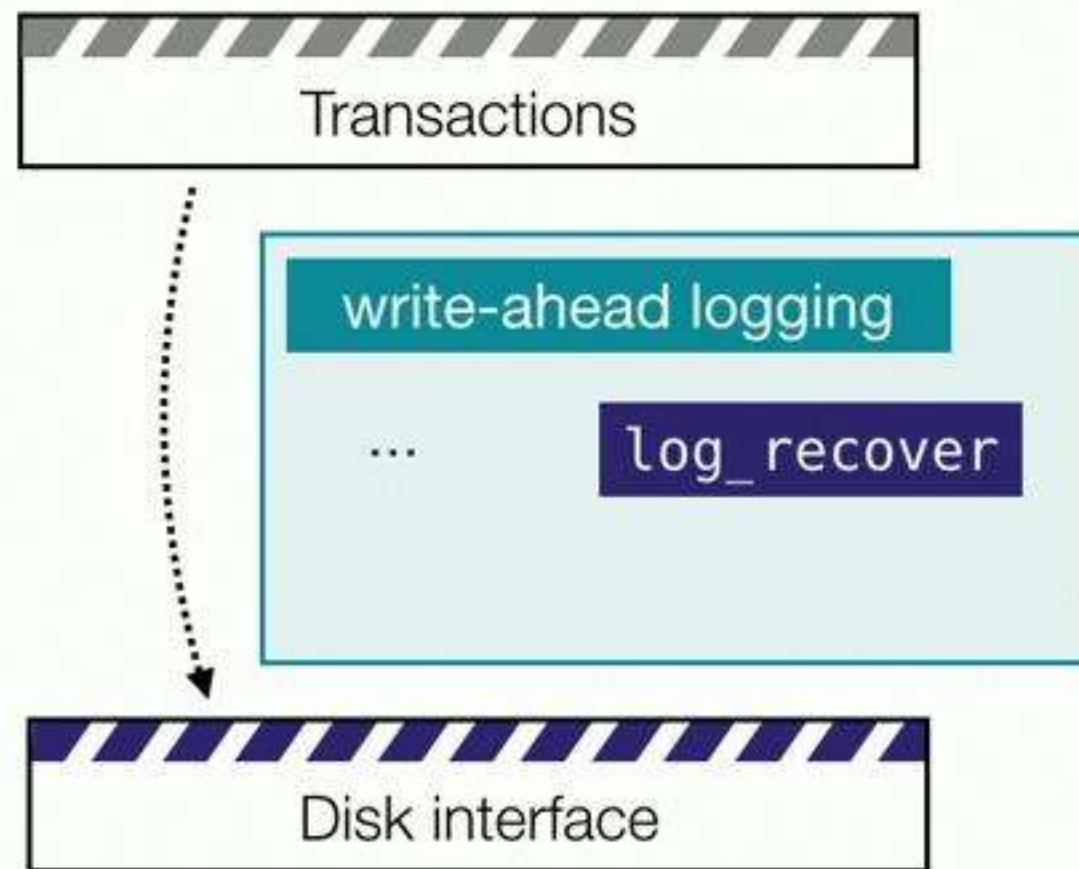
Bob writes a replication system and implements its recovery procedure



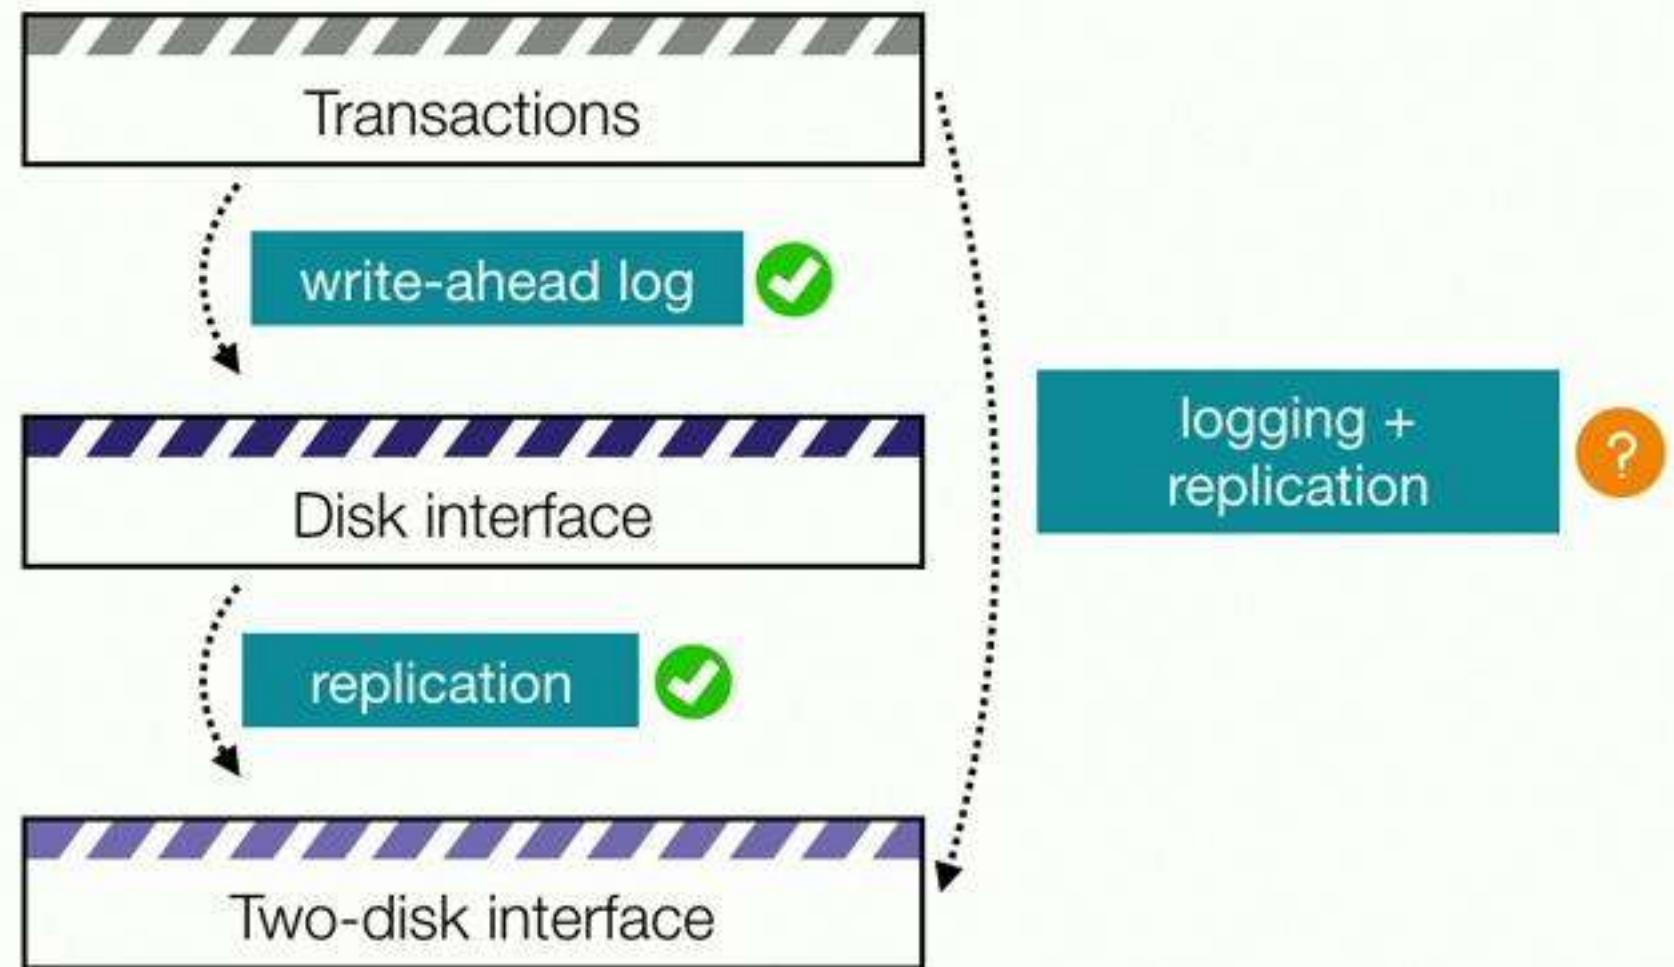
Bob is careful and writes a  
machine-checked proof of correctness

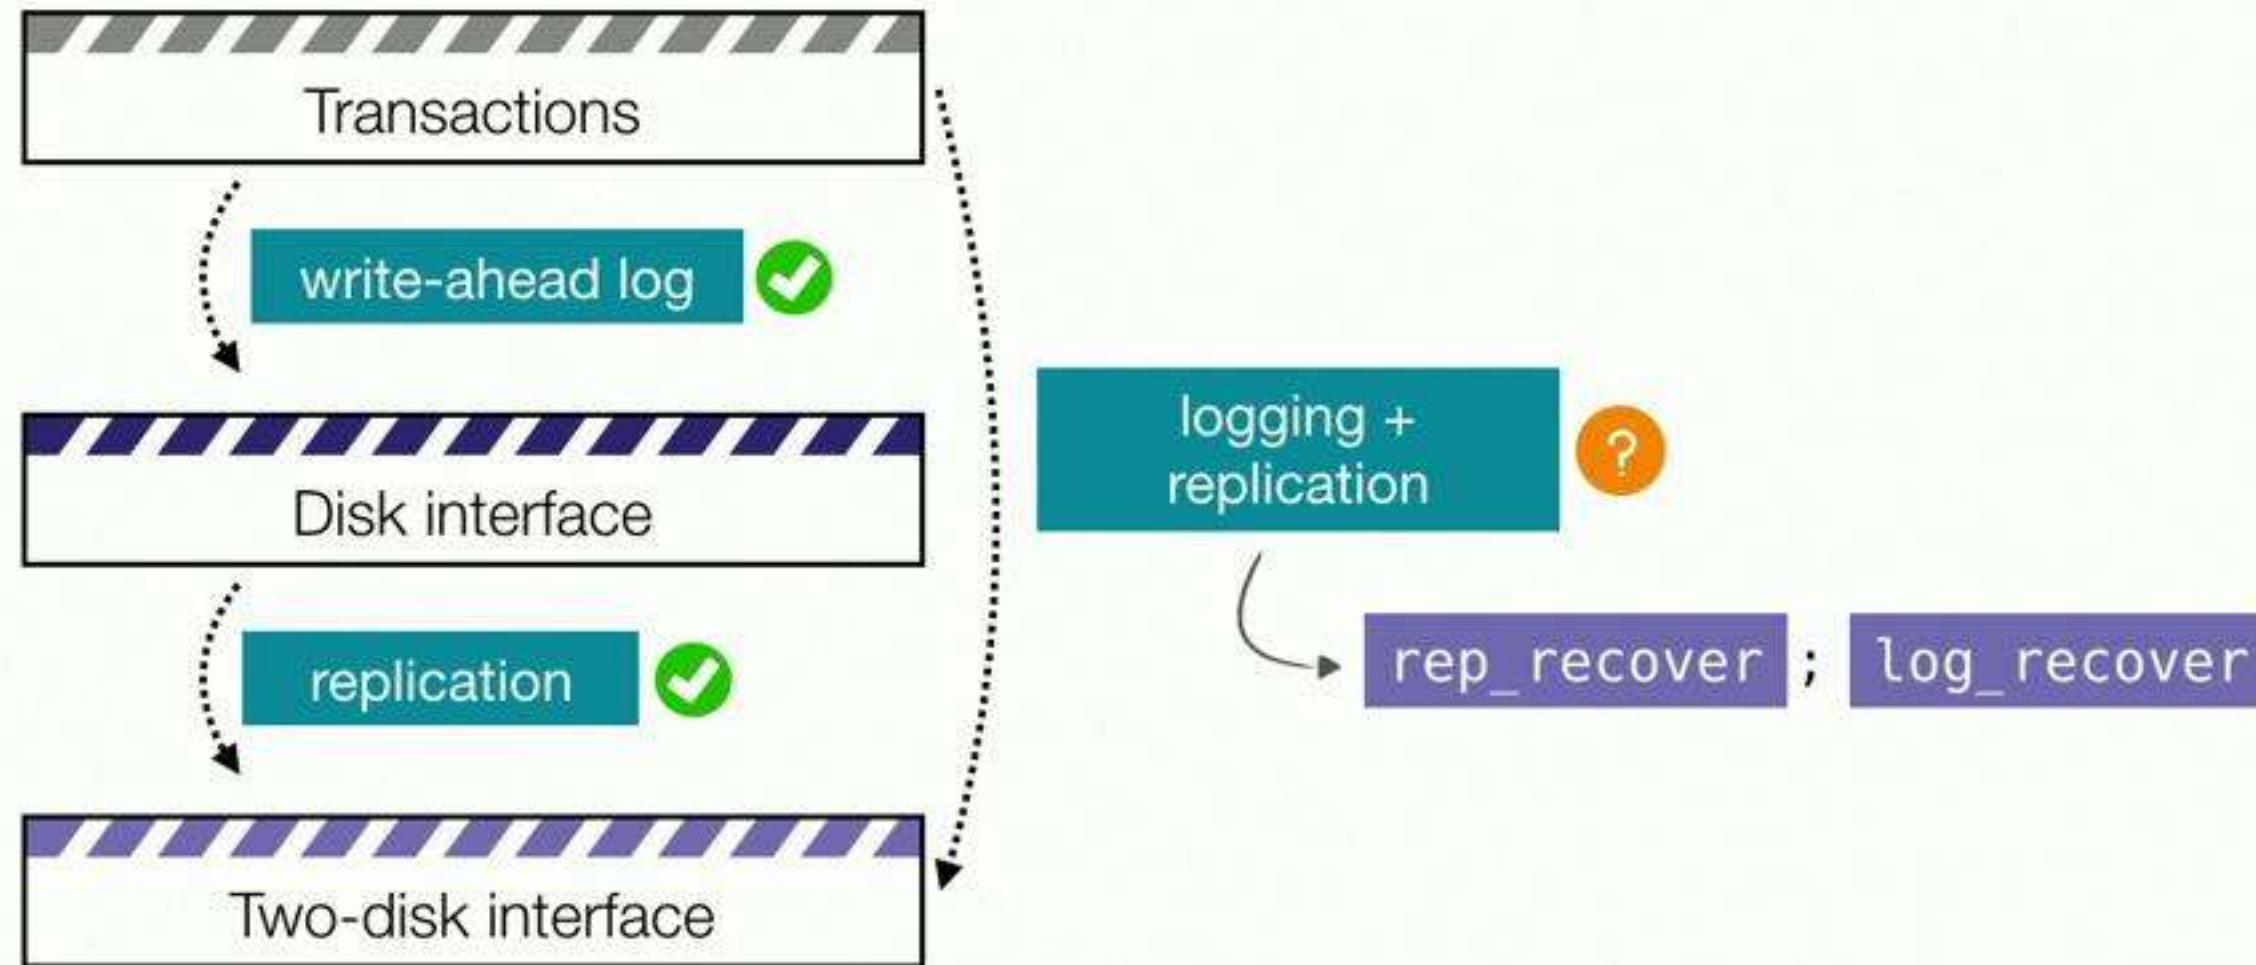
✓ read and write are atomic if you run  
rep\_recover after every crash





- ✓ ops are atomic if you run `log_recover` after every crash





# Challenge: crashes during composed recovery

rep\_recover  under crashes

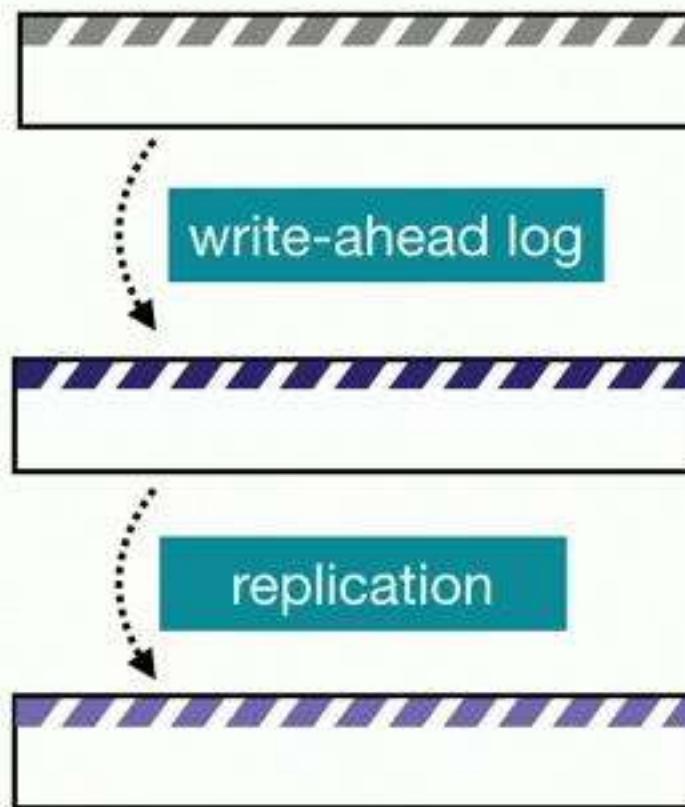
log\_recover  under crashes

---

rep\_recover ; log\_recover

how do we prove correctness  
under crashes using the existing proofs?

# Prior work cannot handle multiple recovery procedures



**CHL** [SOSP '15]

not modular

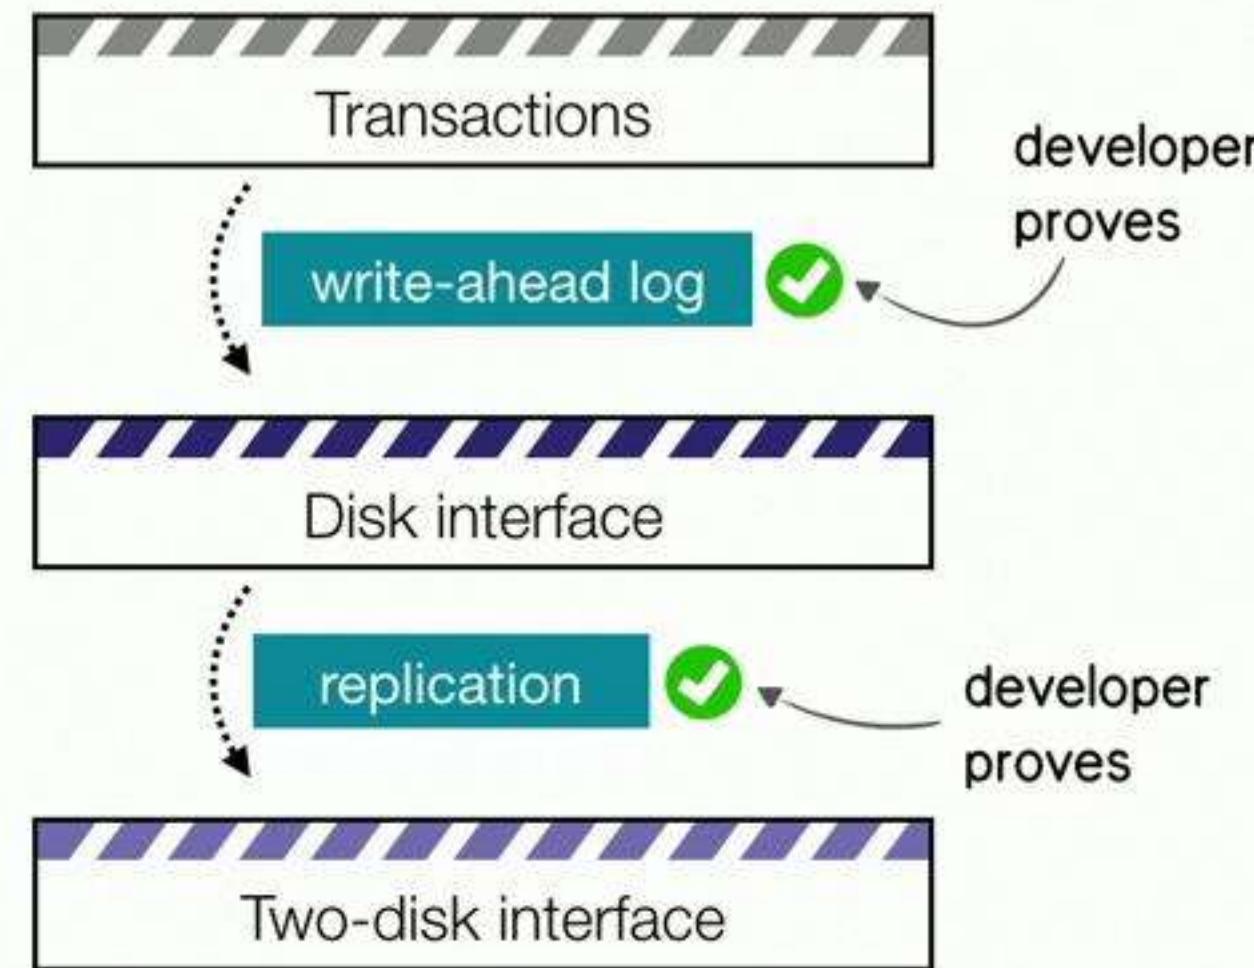
**Yggdrasil** [OSDI '16]

single recovery

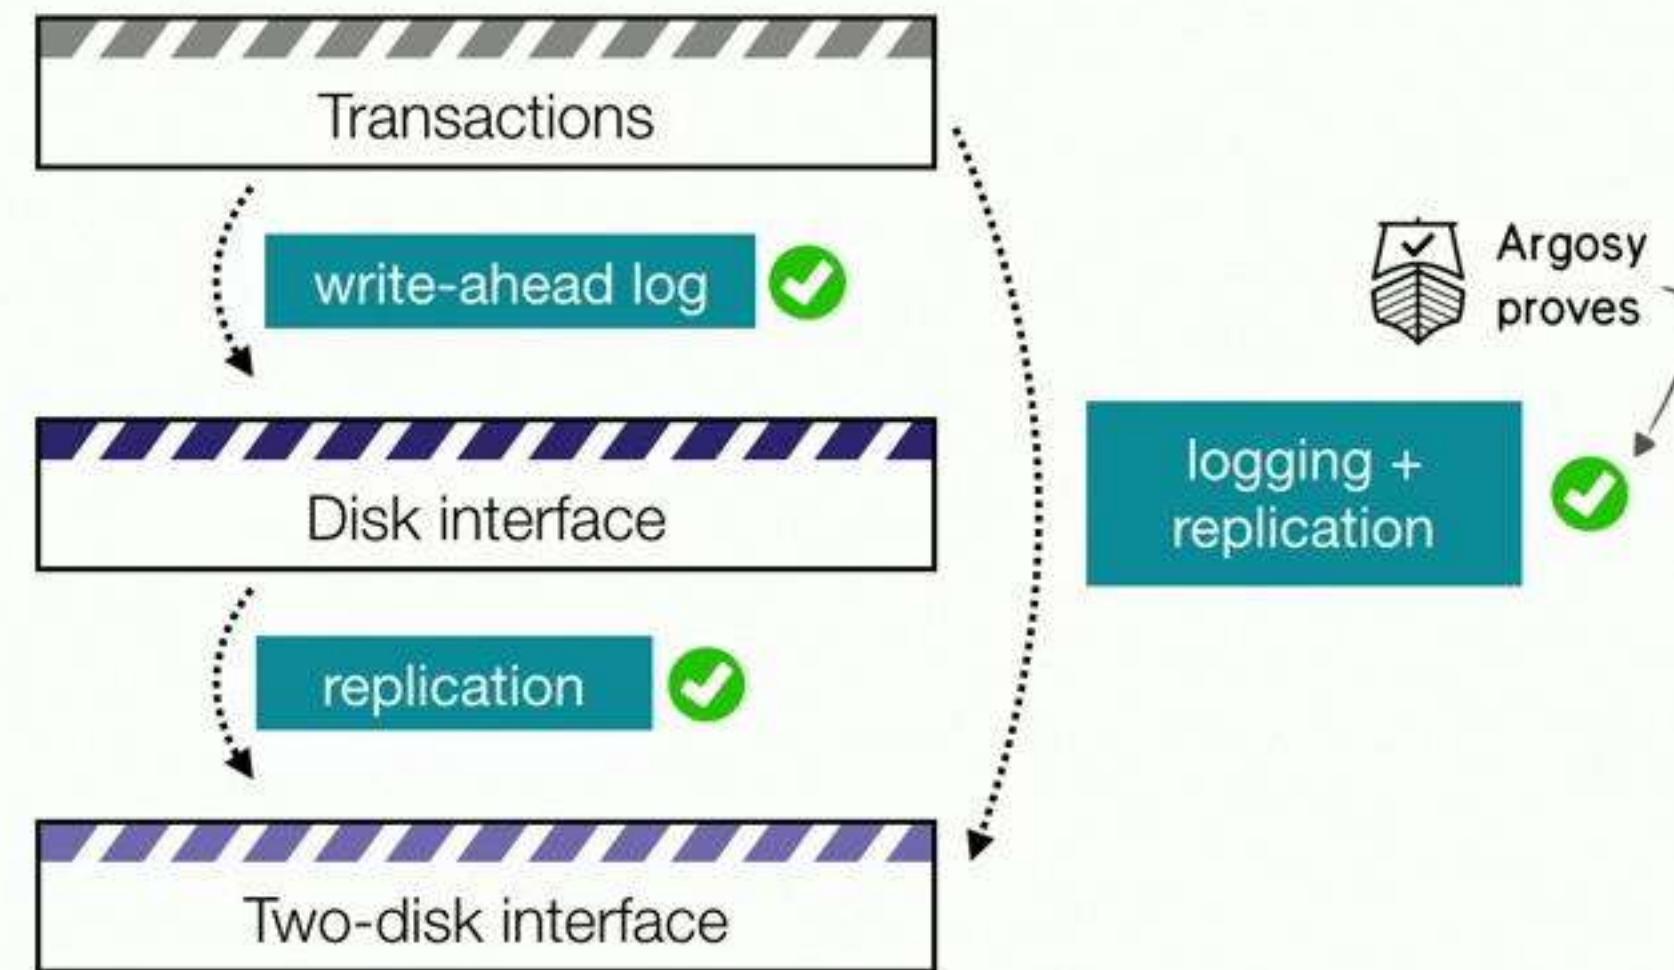
**Flashix** [SCP '16]

restricted recovery  
procedures

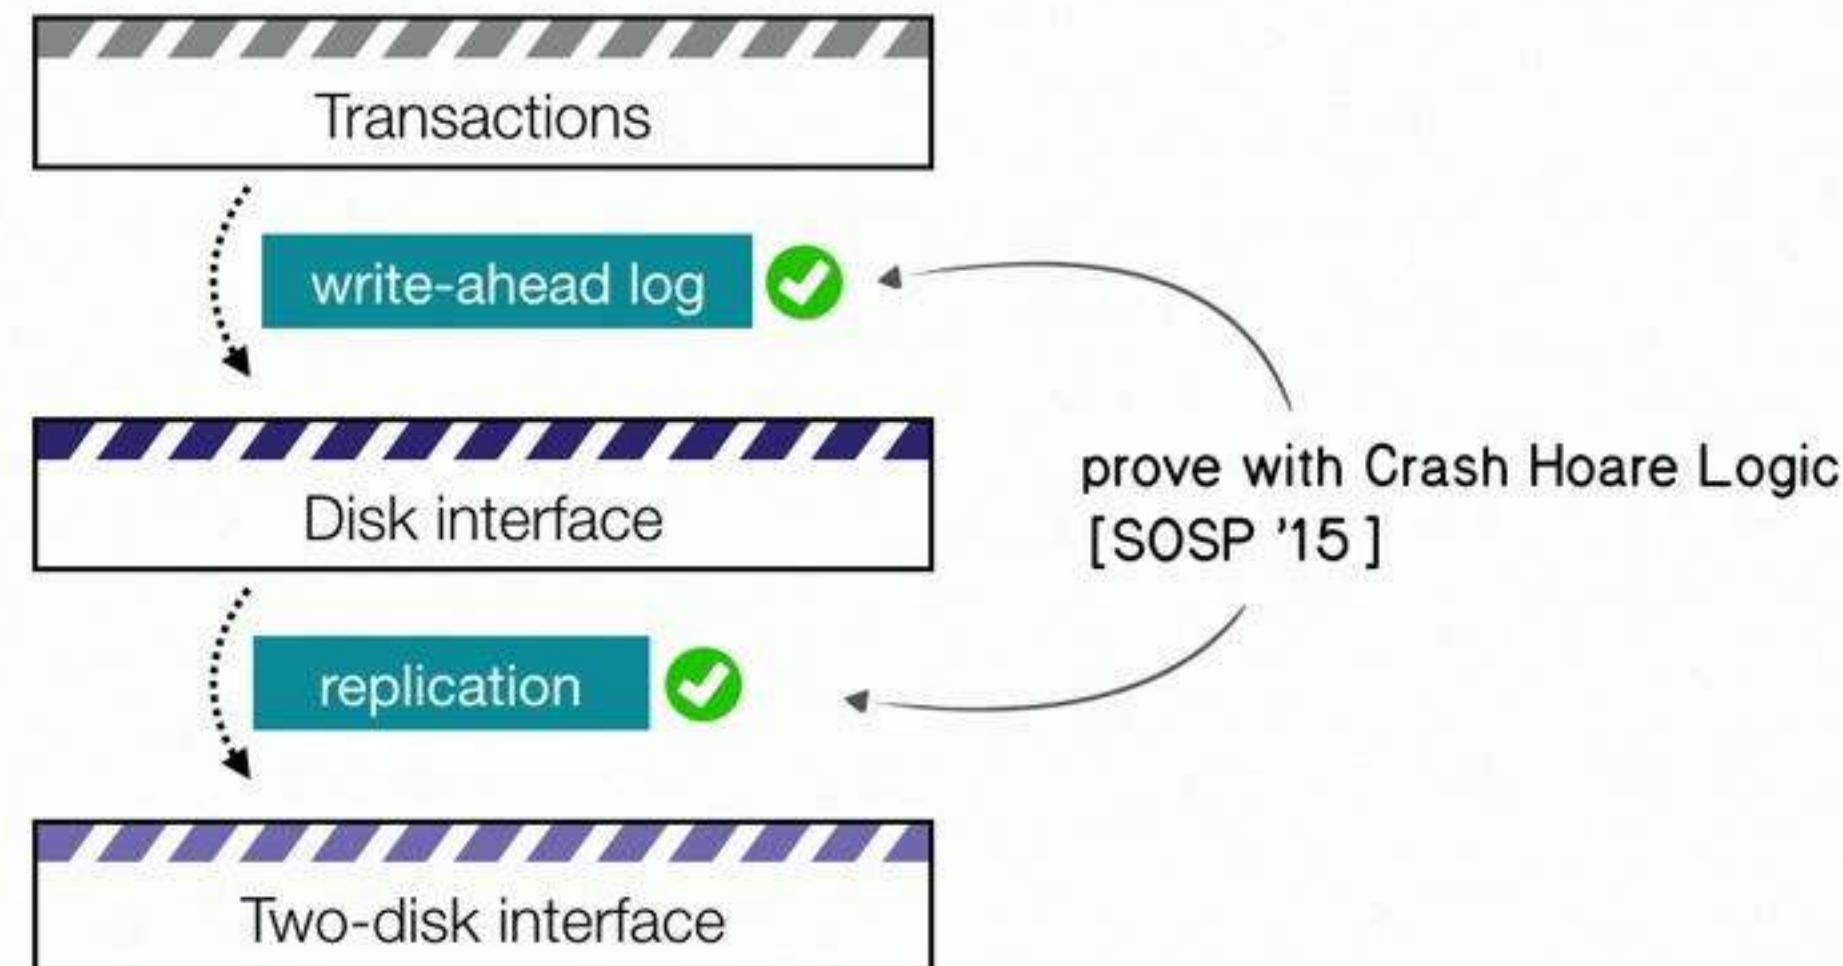
# Argosy supports modular recovery proofs



# Argosy supports modular recovery proofs



# Argosy is compatible with existing techniques



# Contributions

Recovery refinement for modular proofs

CHL for proving recovery refinement

see paper

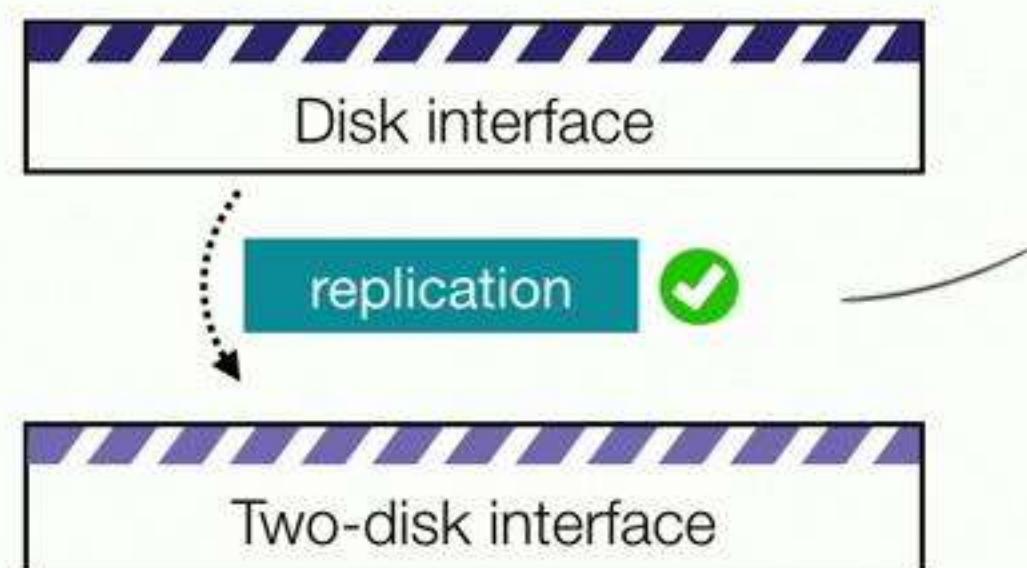
Verified example: logging + replication

see code

Machine-checked proofs in Coq

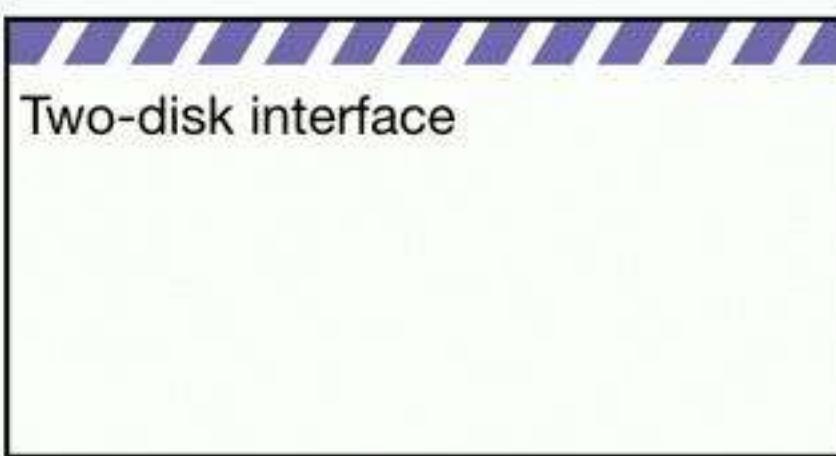
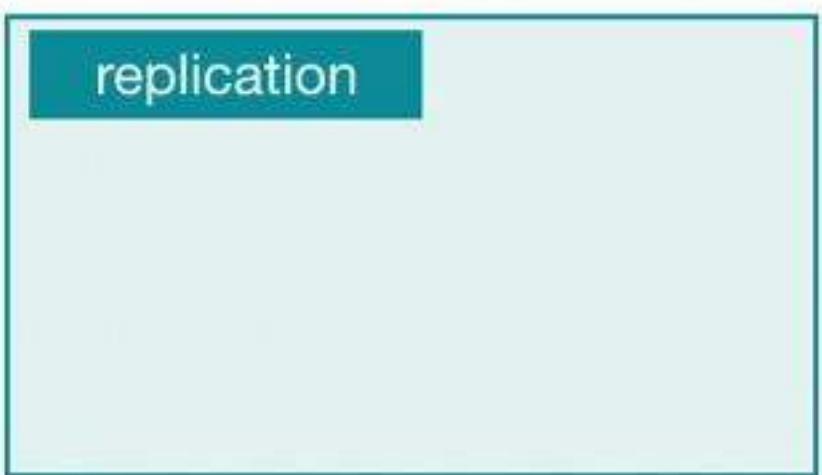


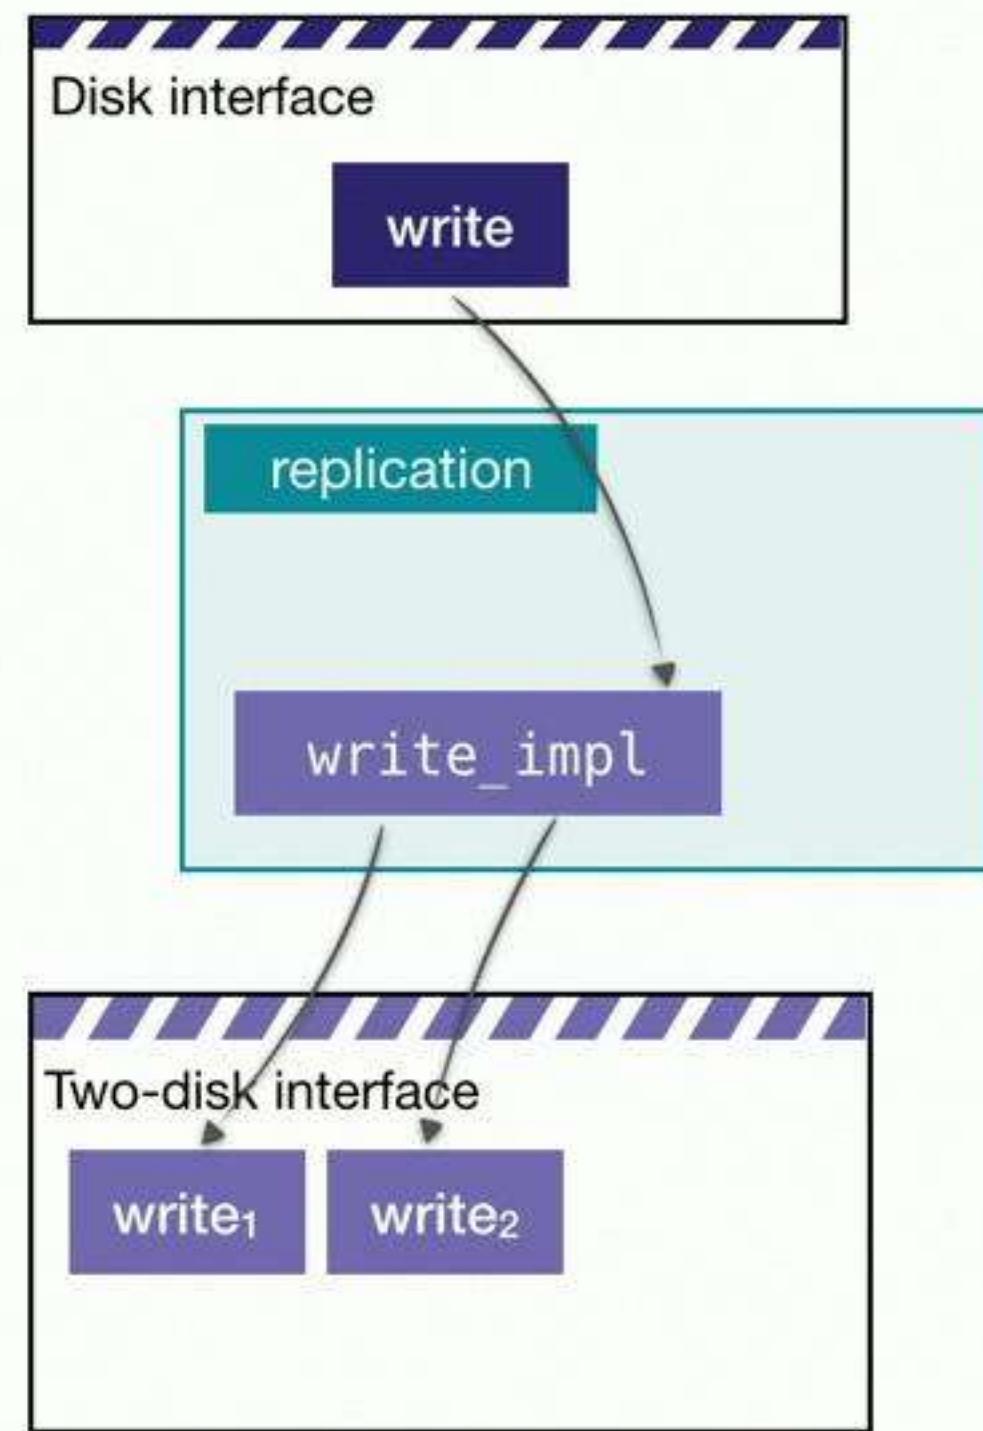
# Preview: recovery refinement

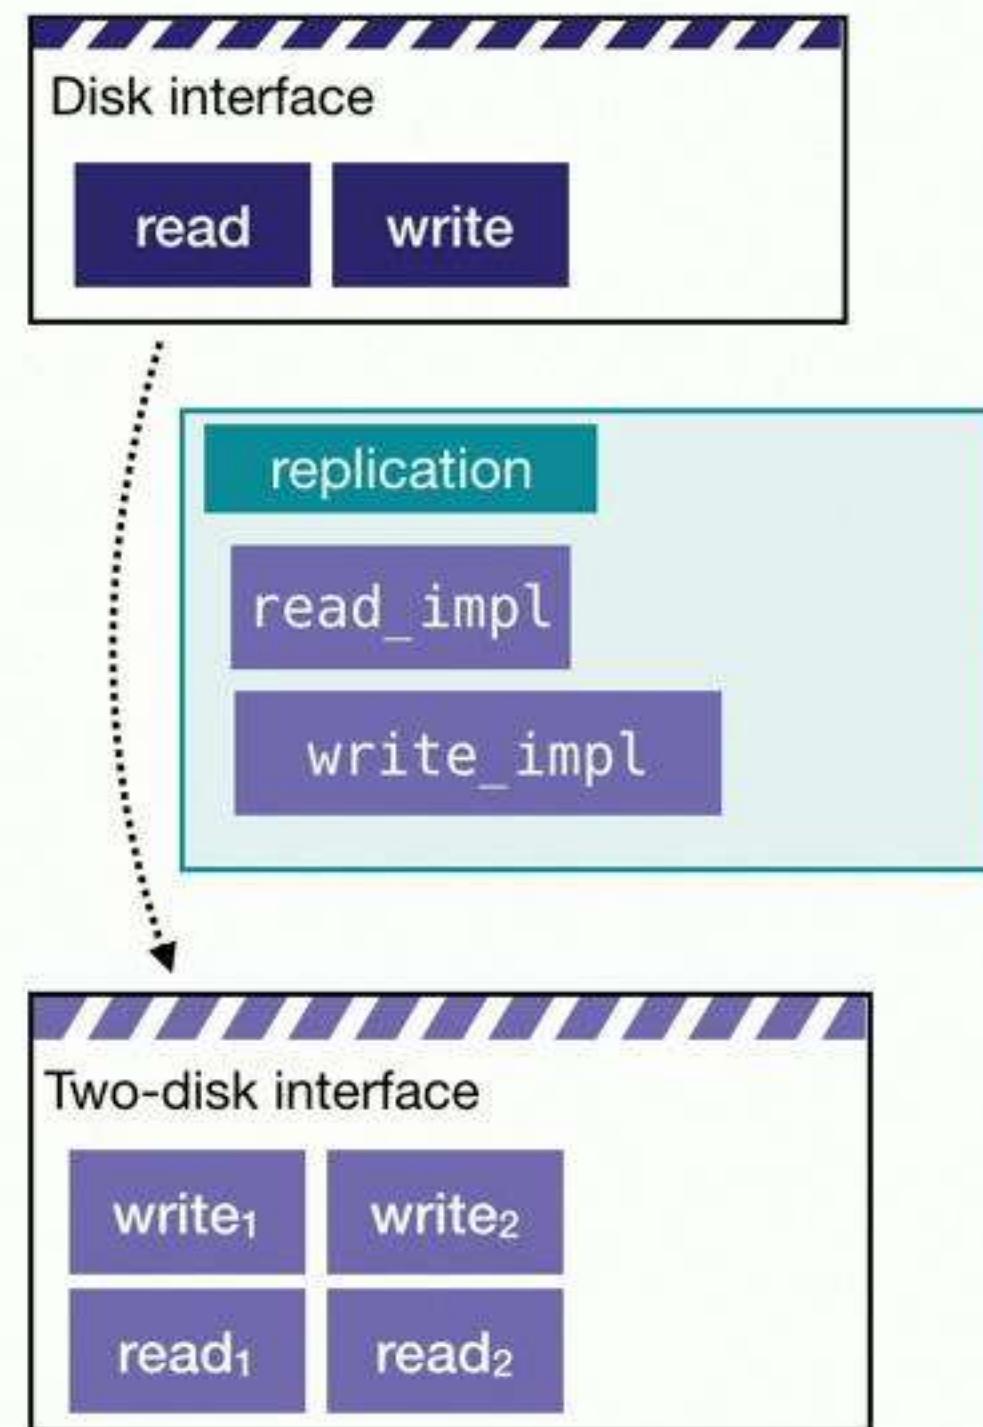


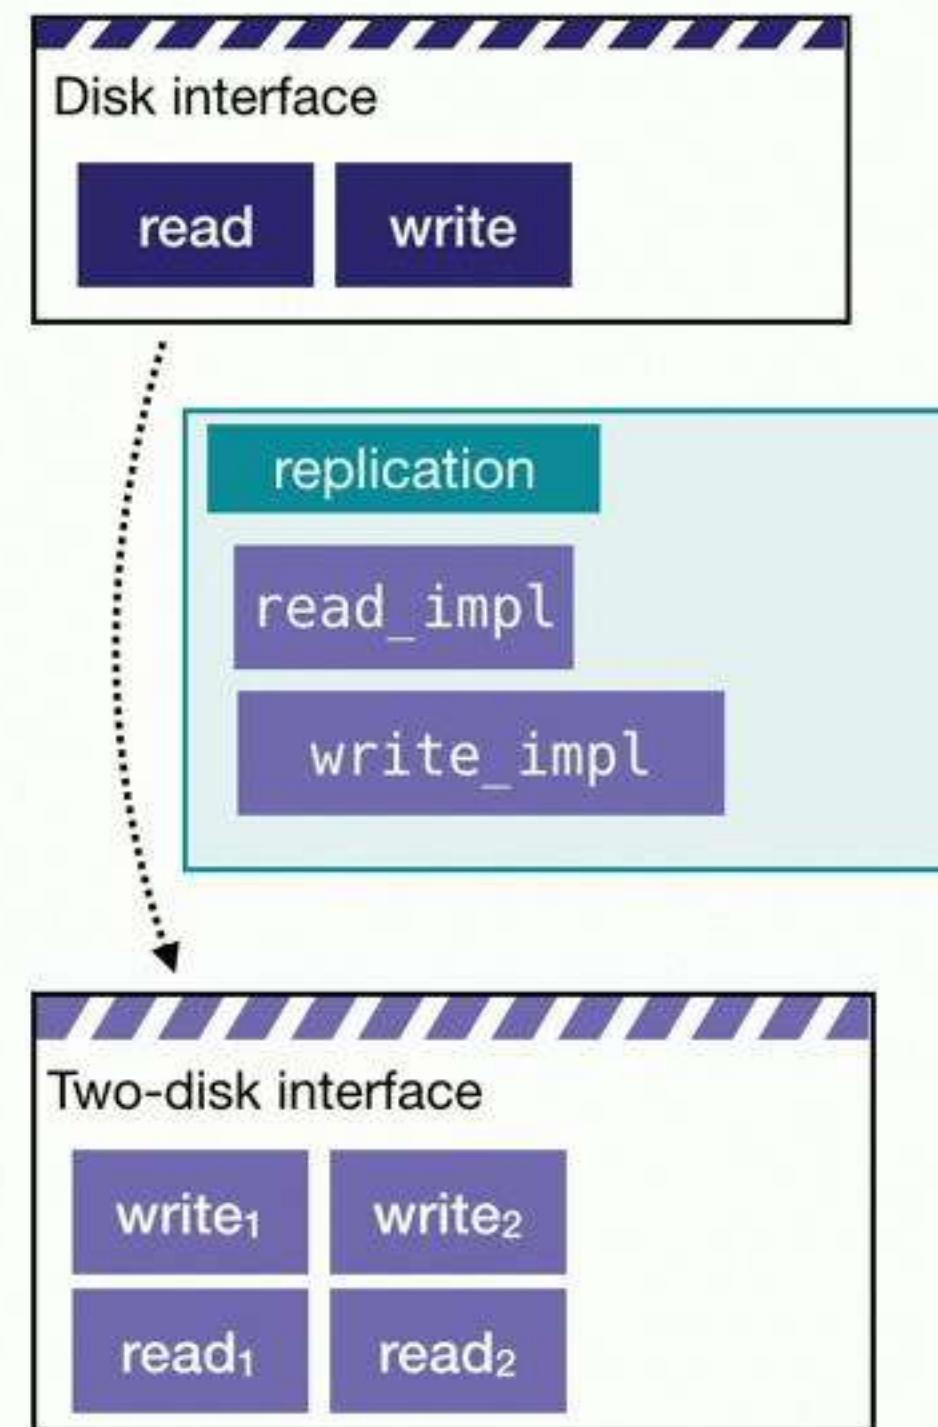
1. Normal execution correctness using *refinement*
2. Crash and recovery correctness using *recovery refinement*

# Refinement

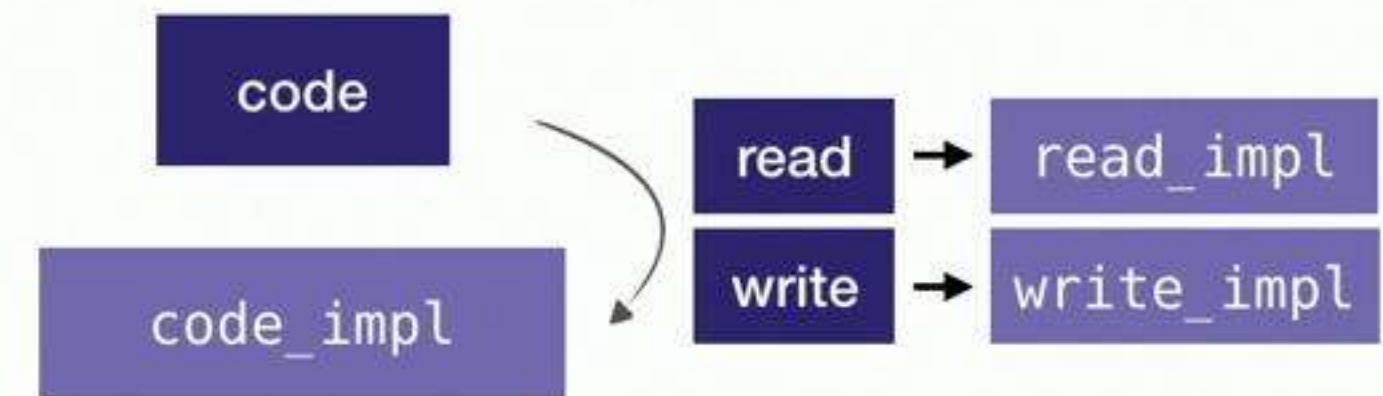




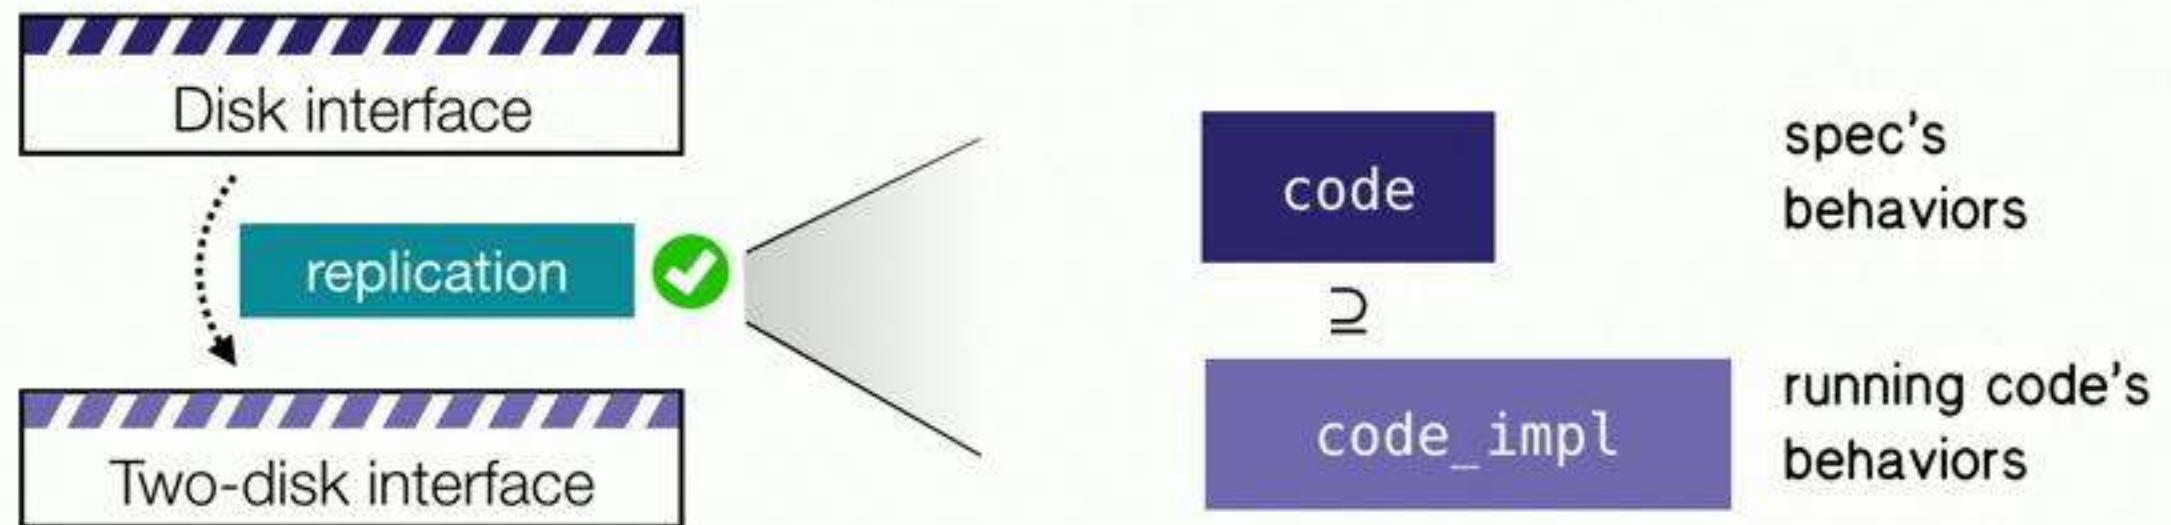




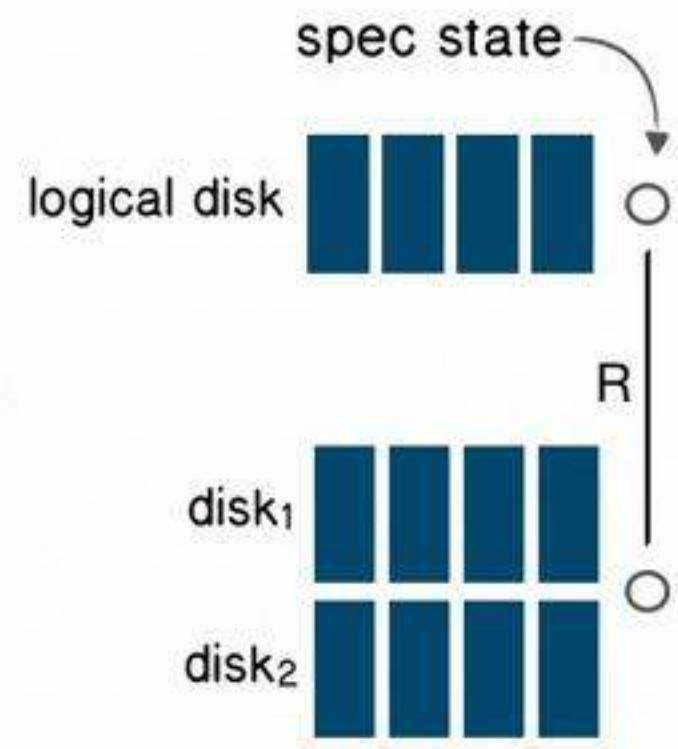
correctness is based on how we use replication:  
run code using Disk interface on top of two disks



# Correctness: trace inclusion

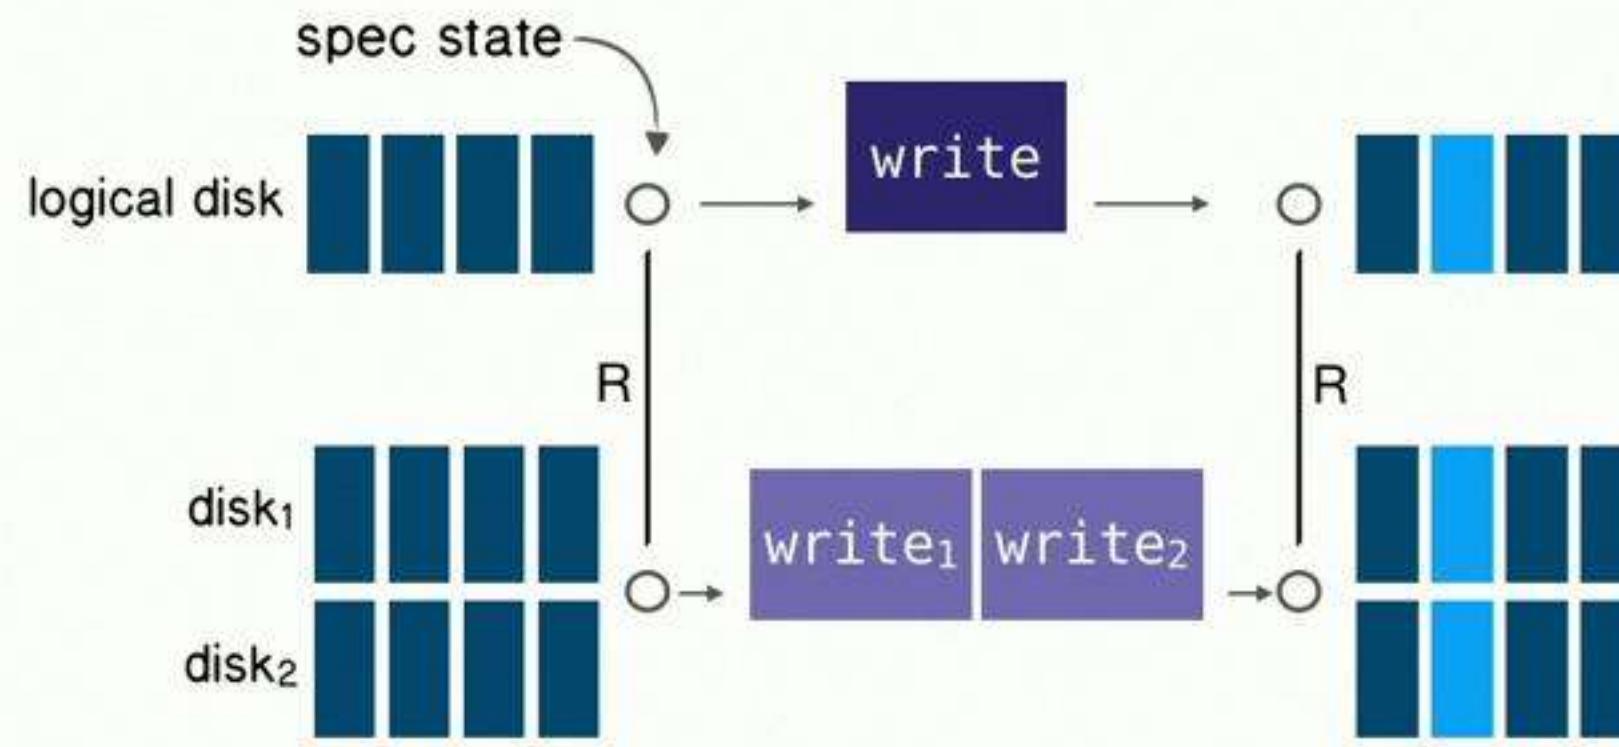


# Proving correctness with an abstraction relation



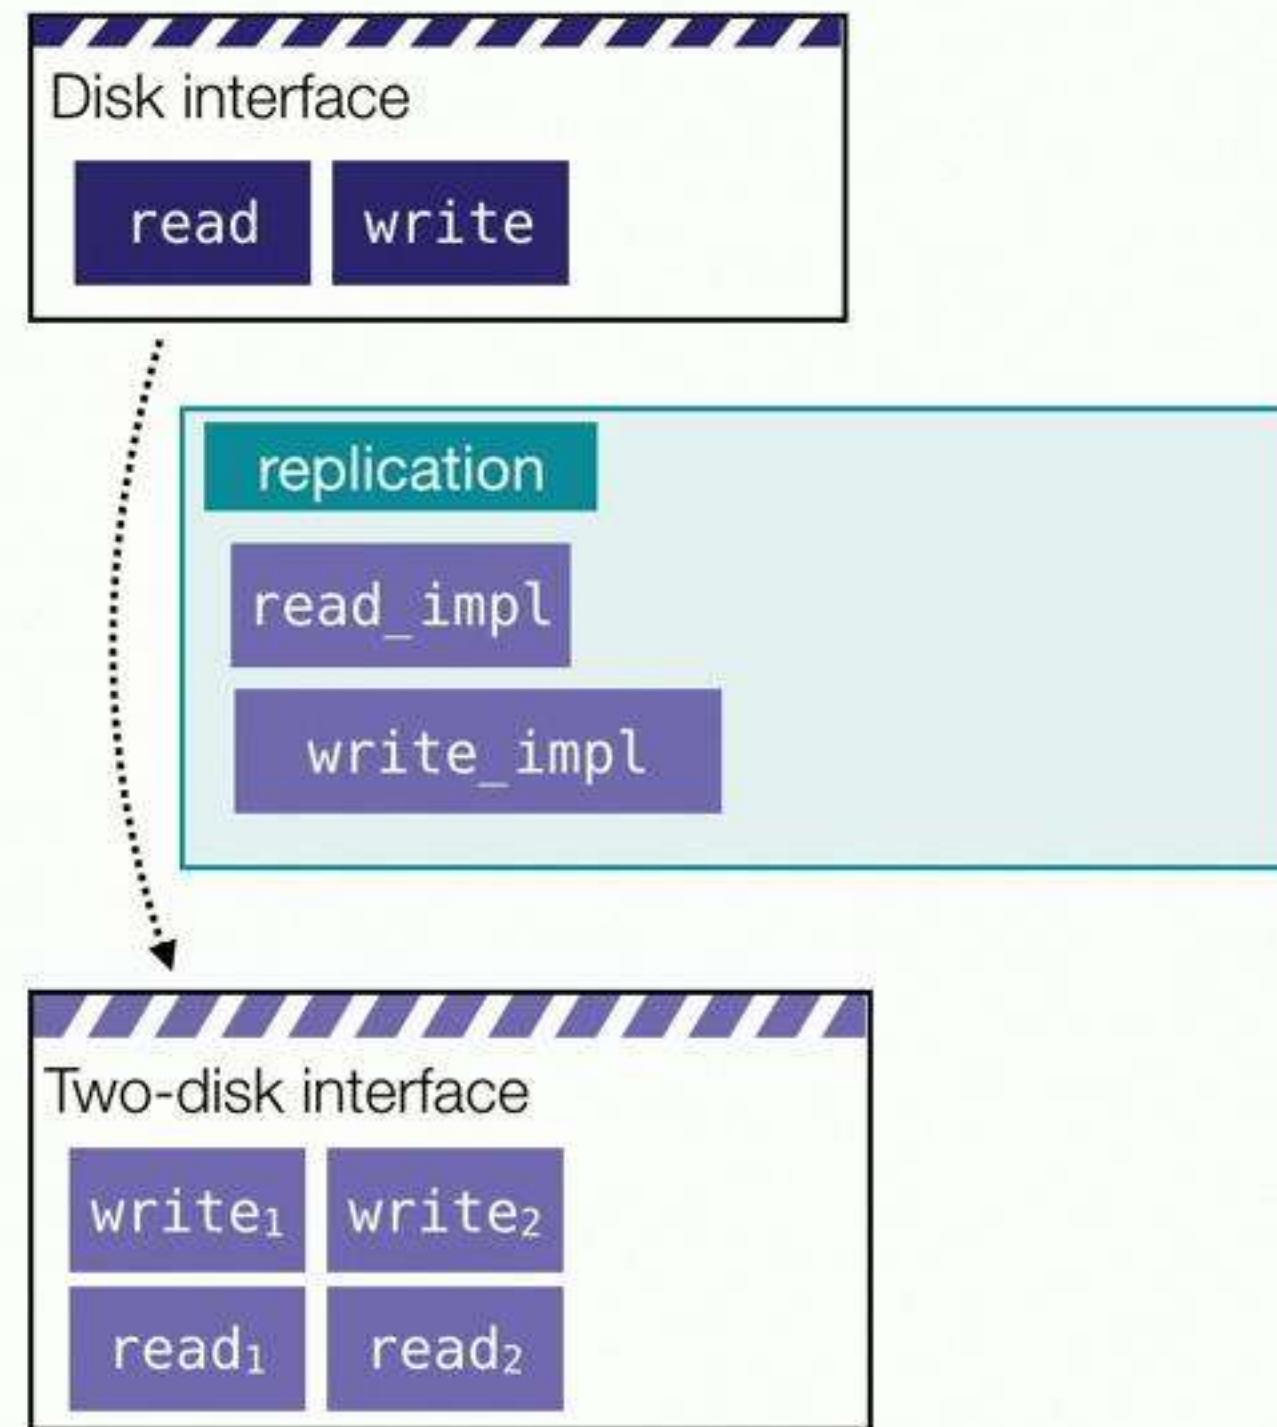
1. developer provides abstraction relation  $R$

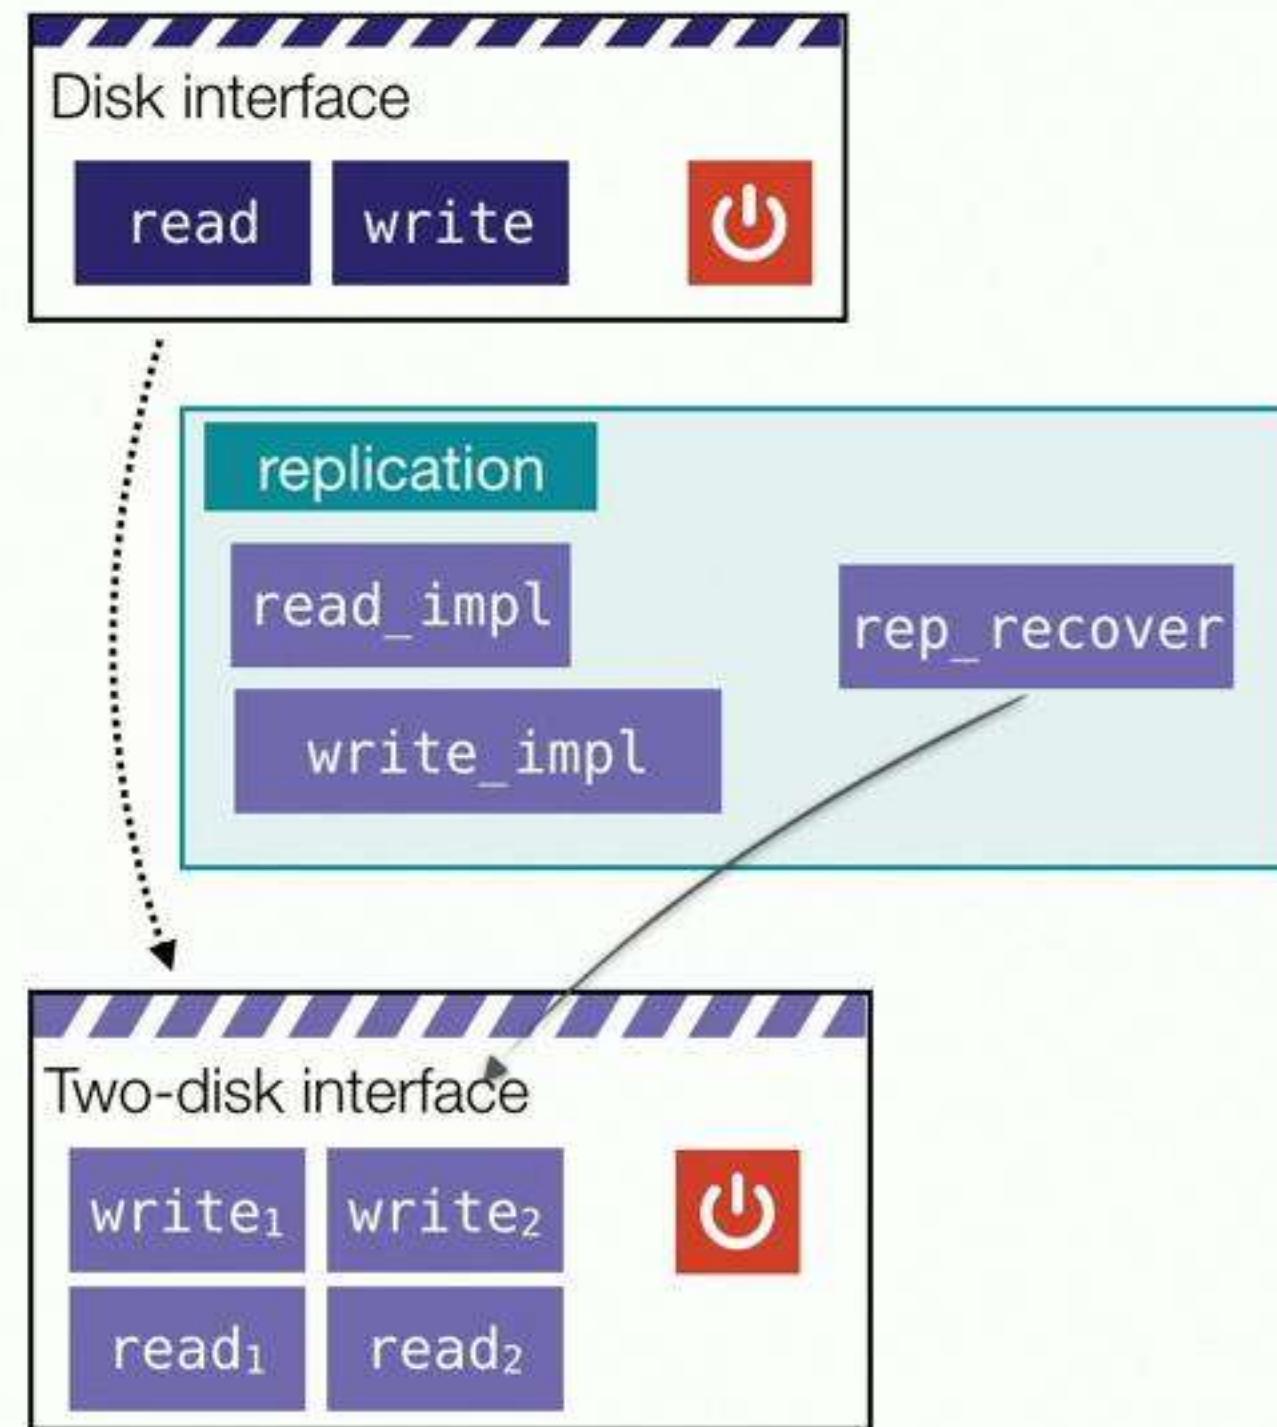
# Proving correctness with an abstraction relation

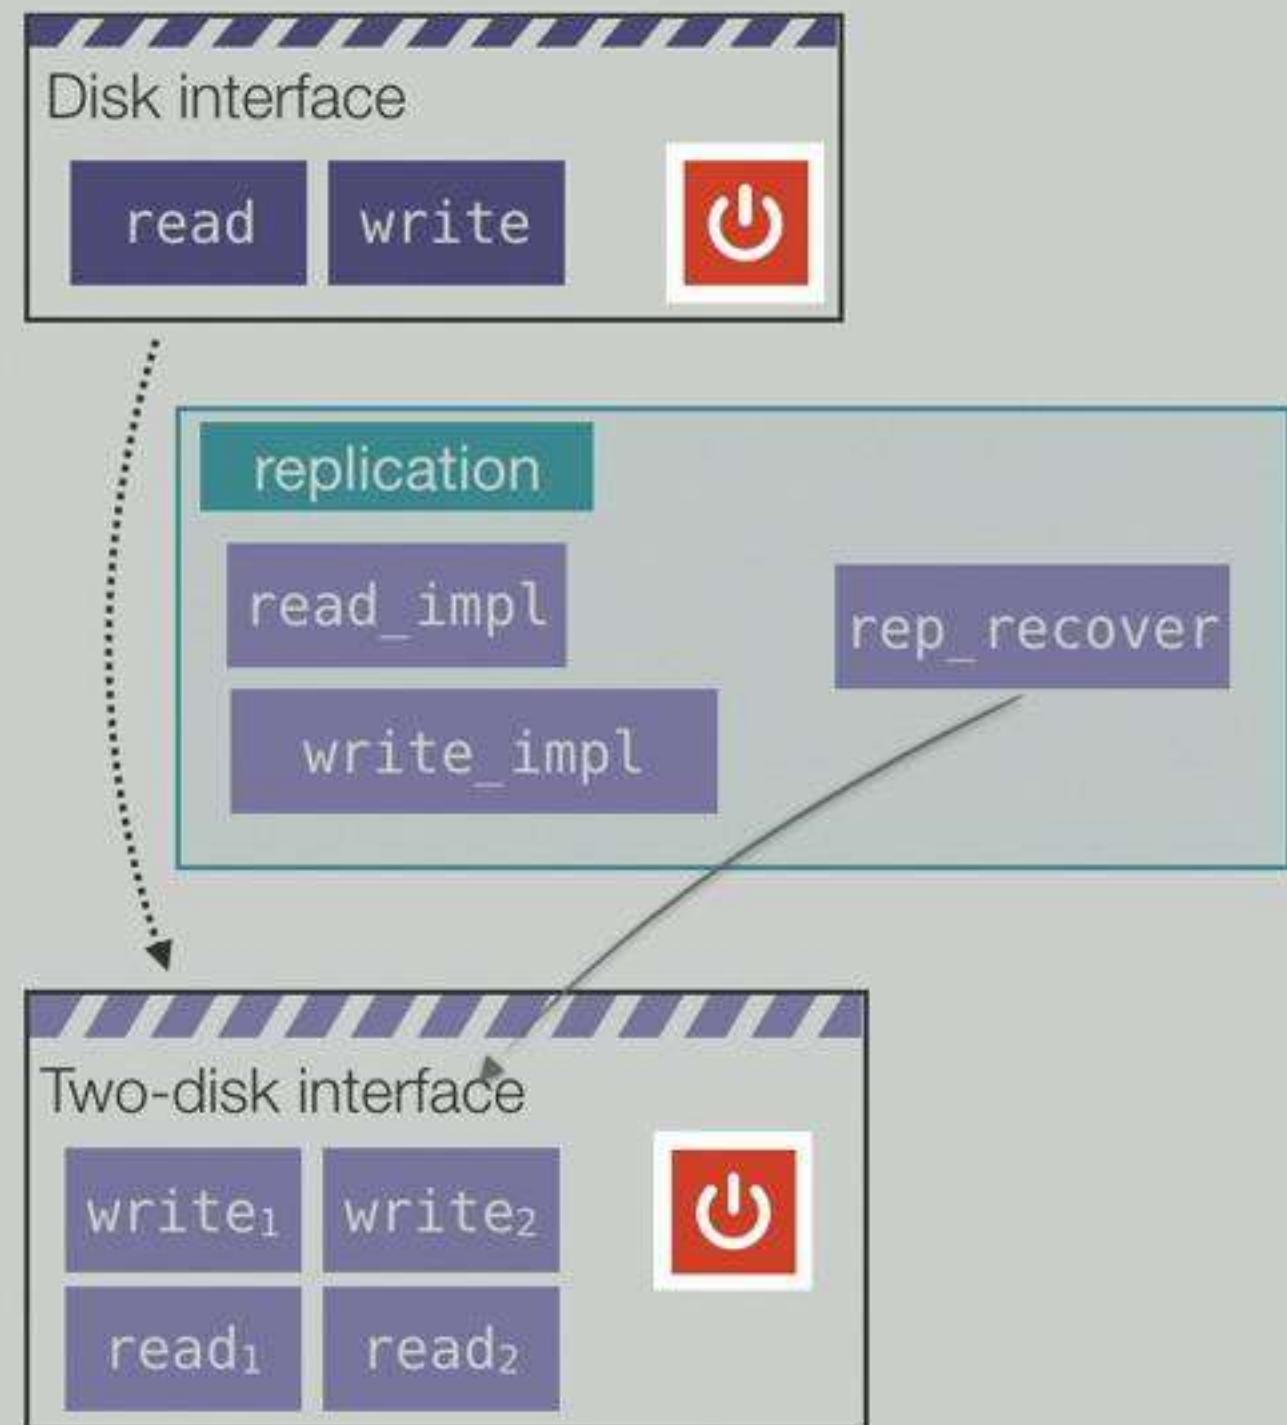


1. developer provides abstraction relation  $R$
2. prove spec execution exists
3. and abstraction relation is preserved

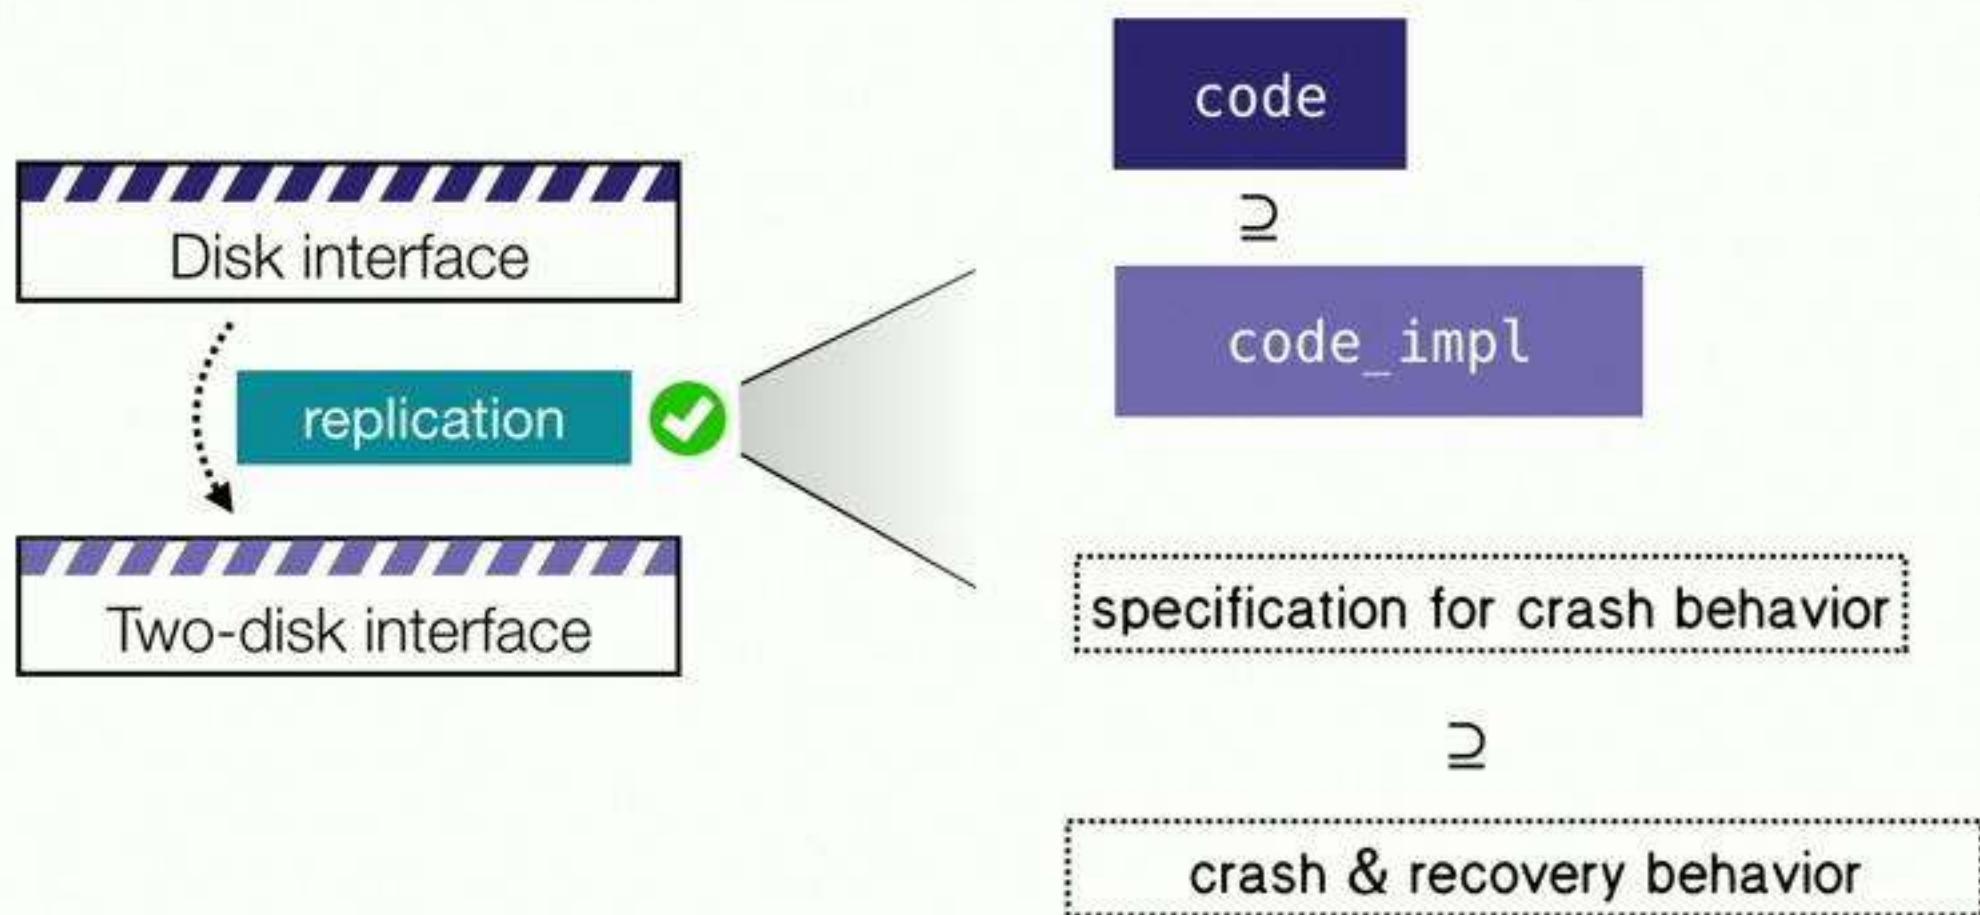
# Recovery refinement



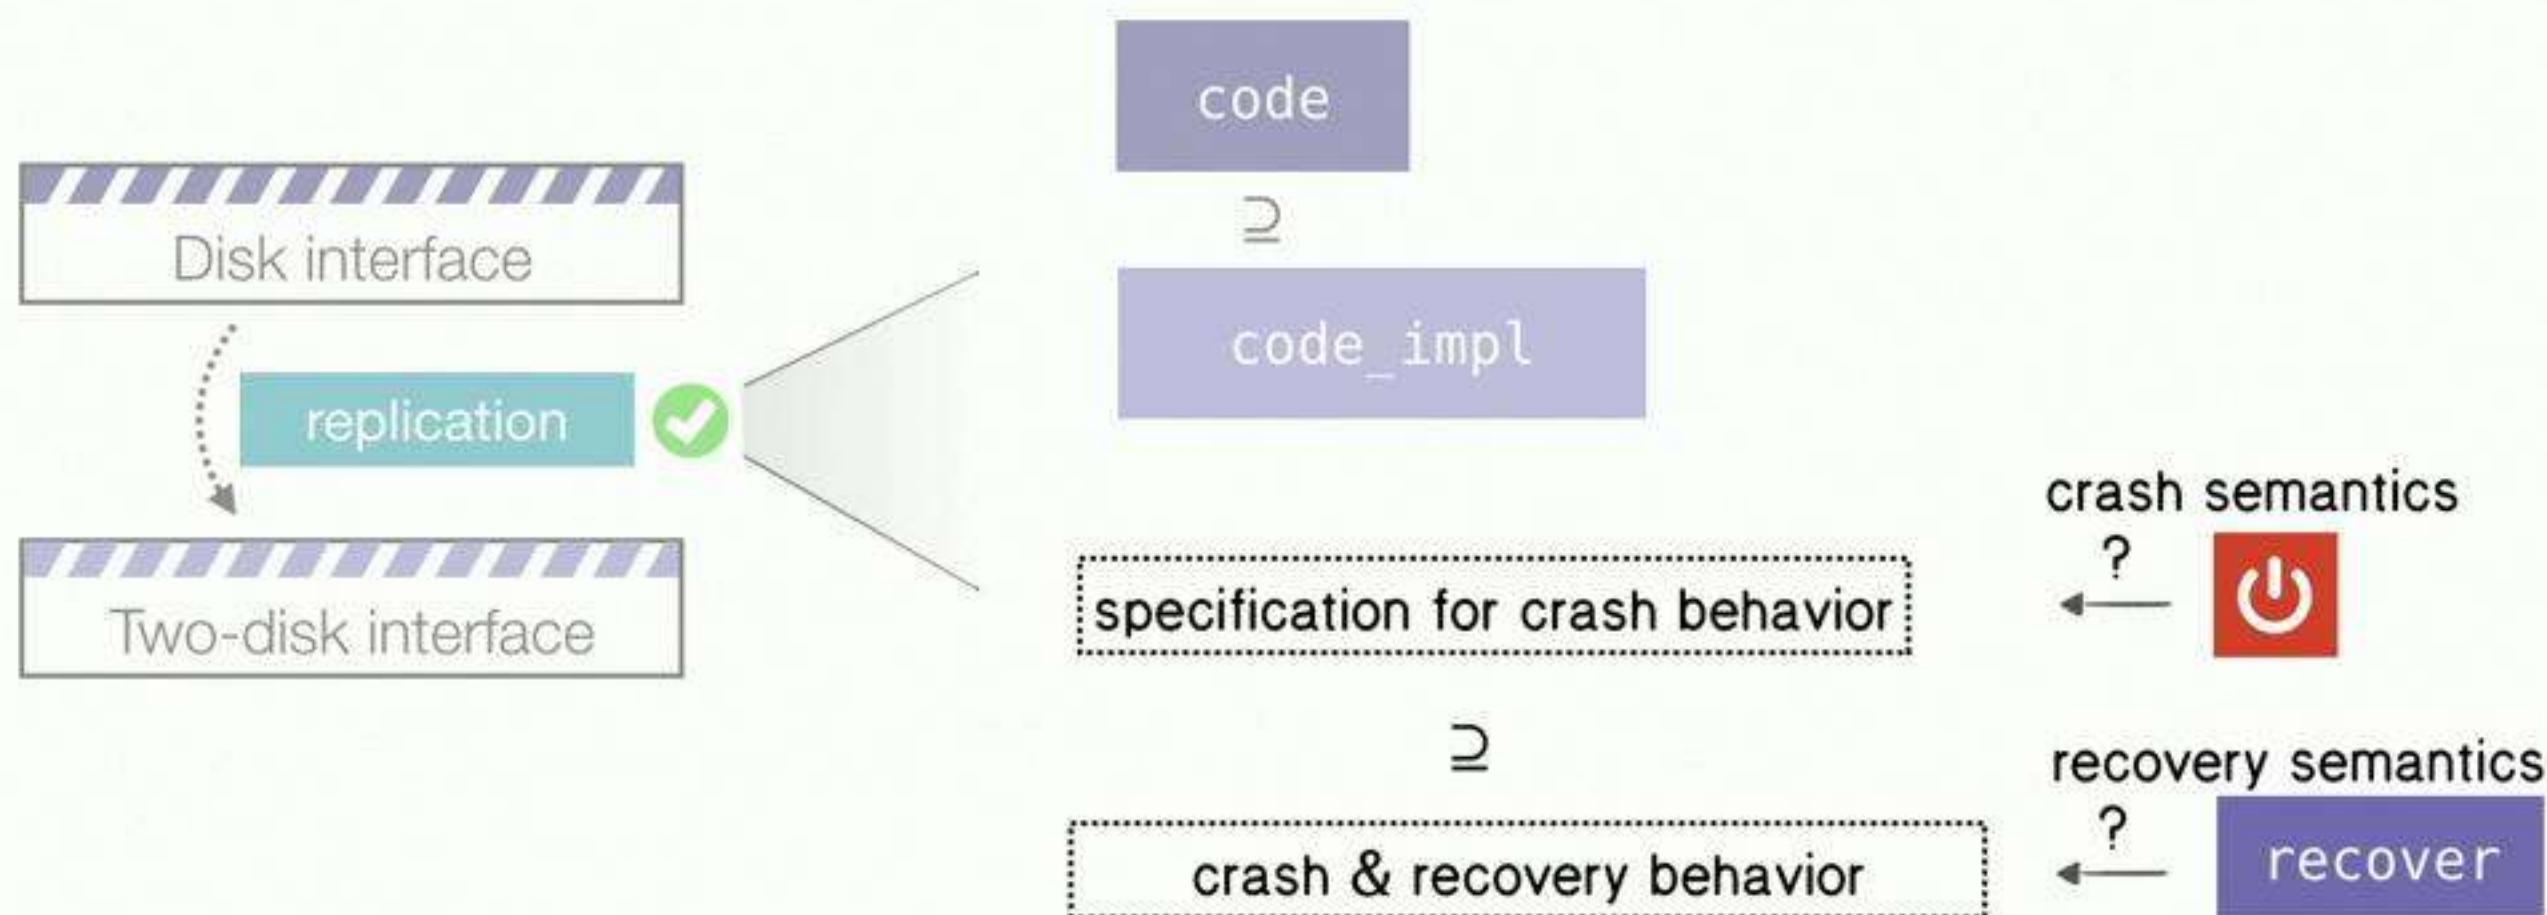


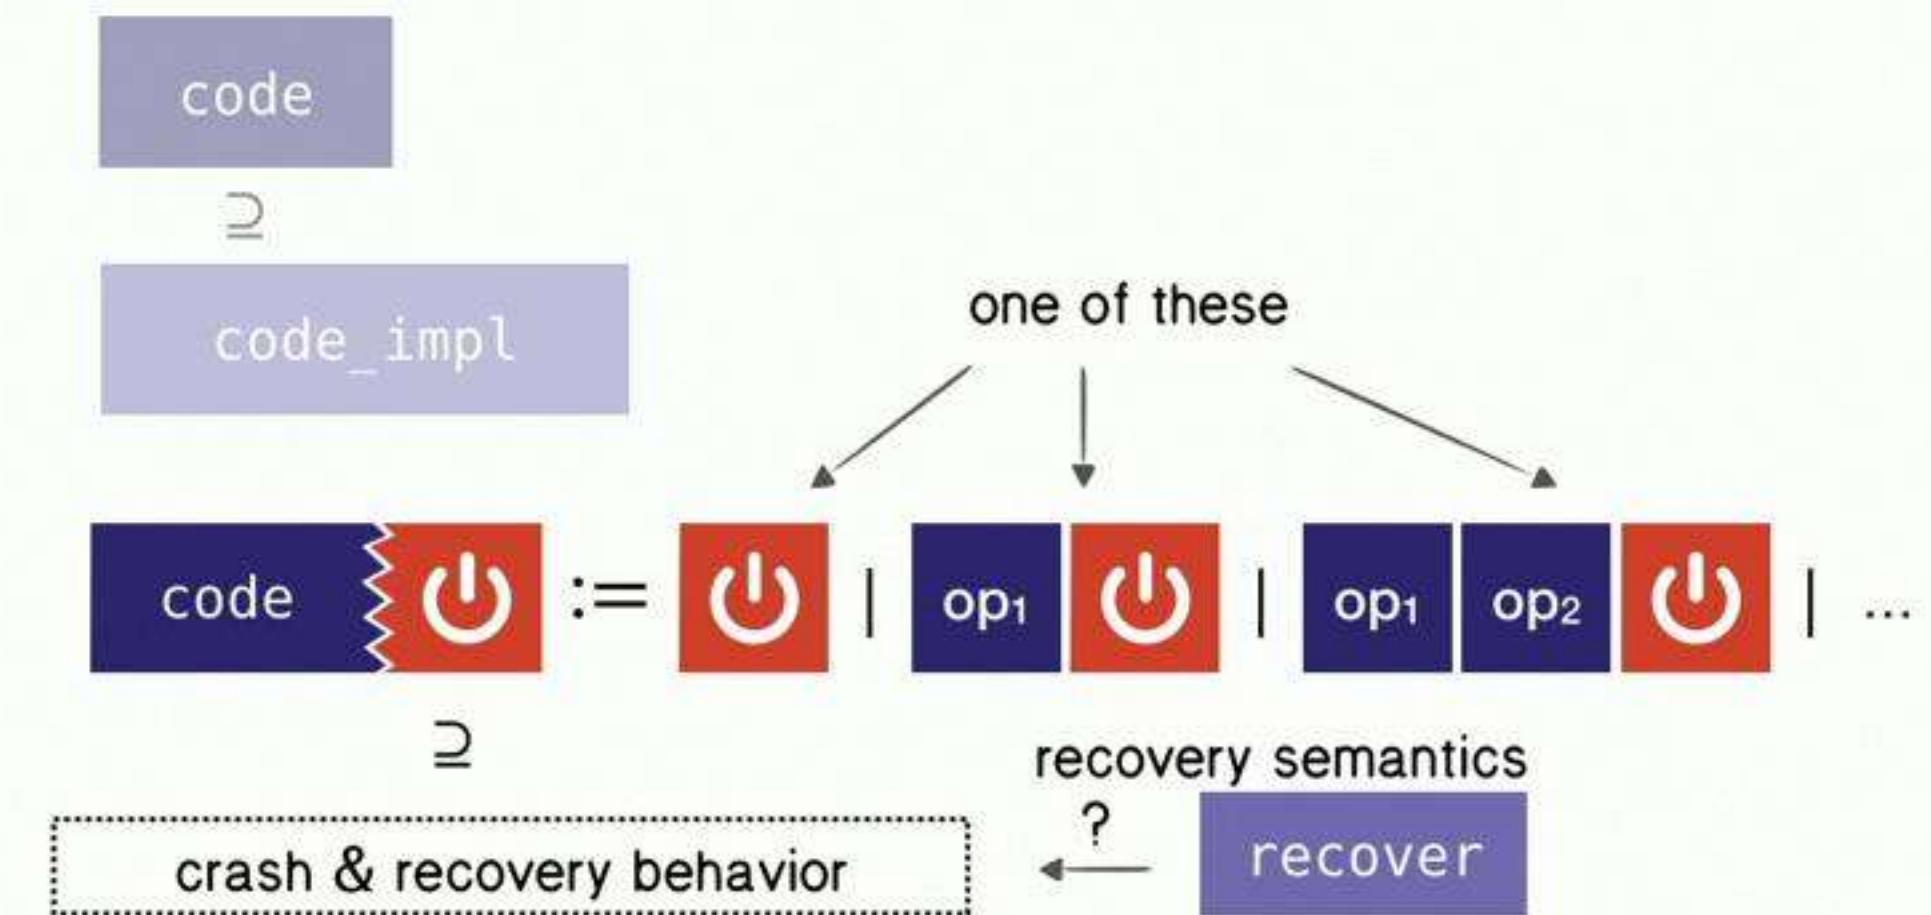
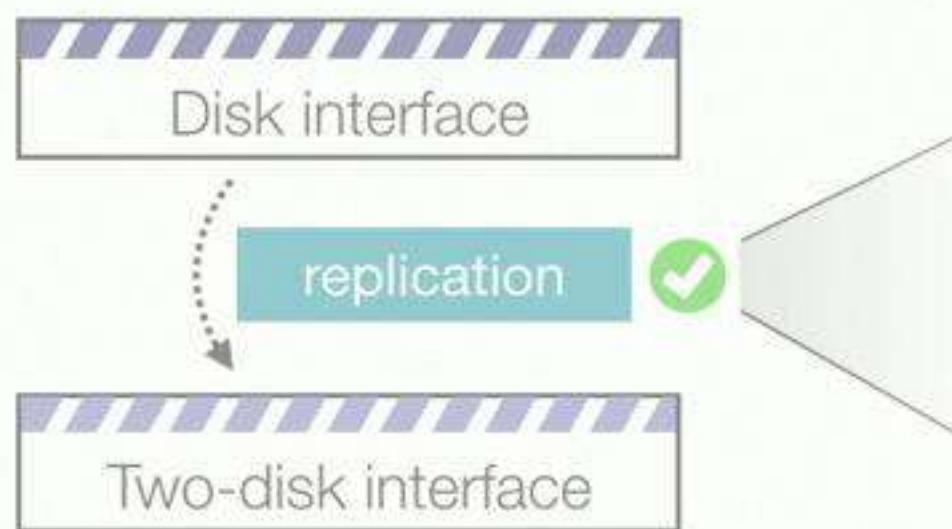


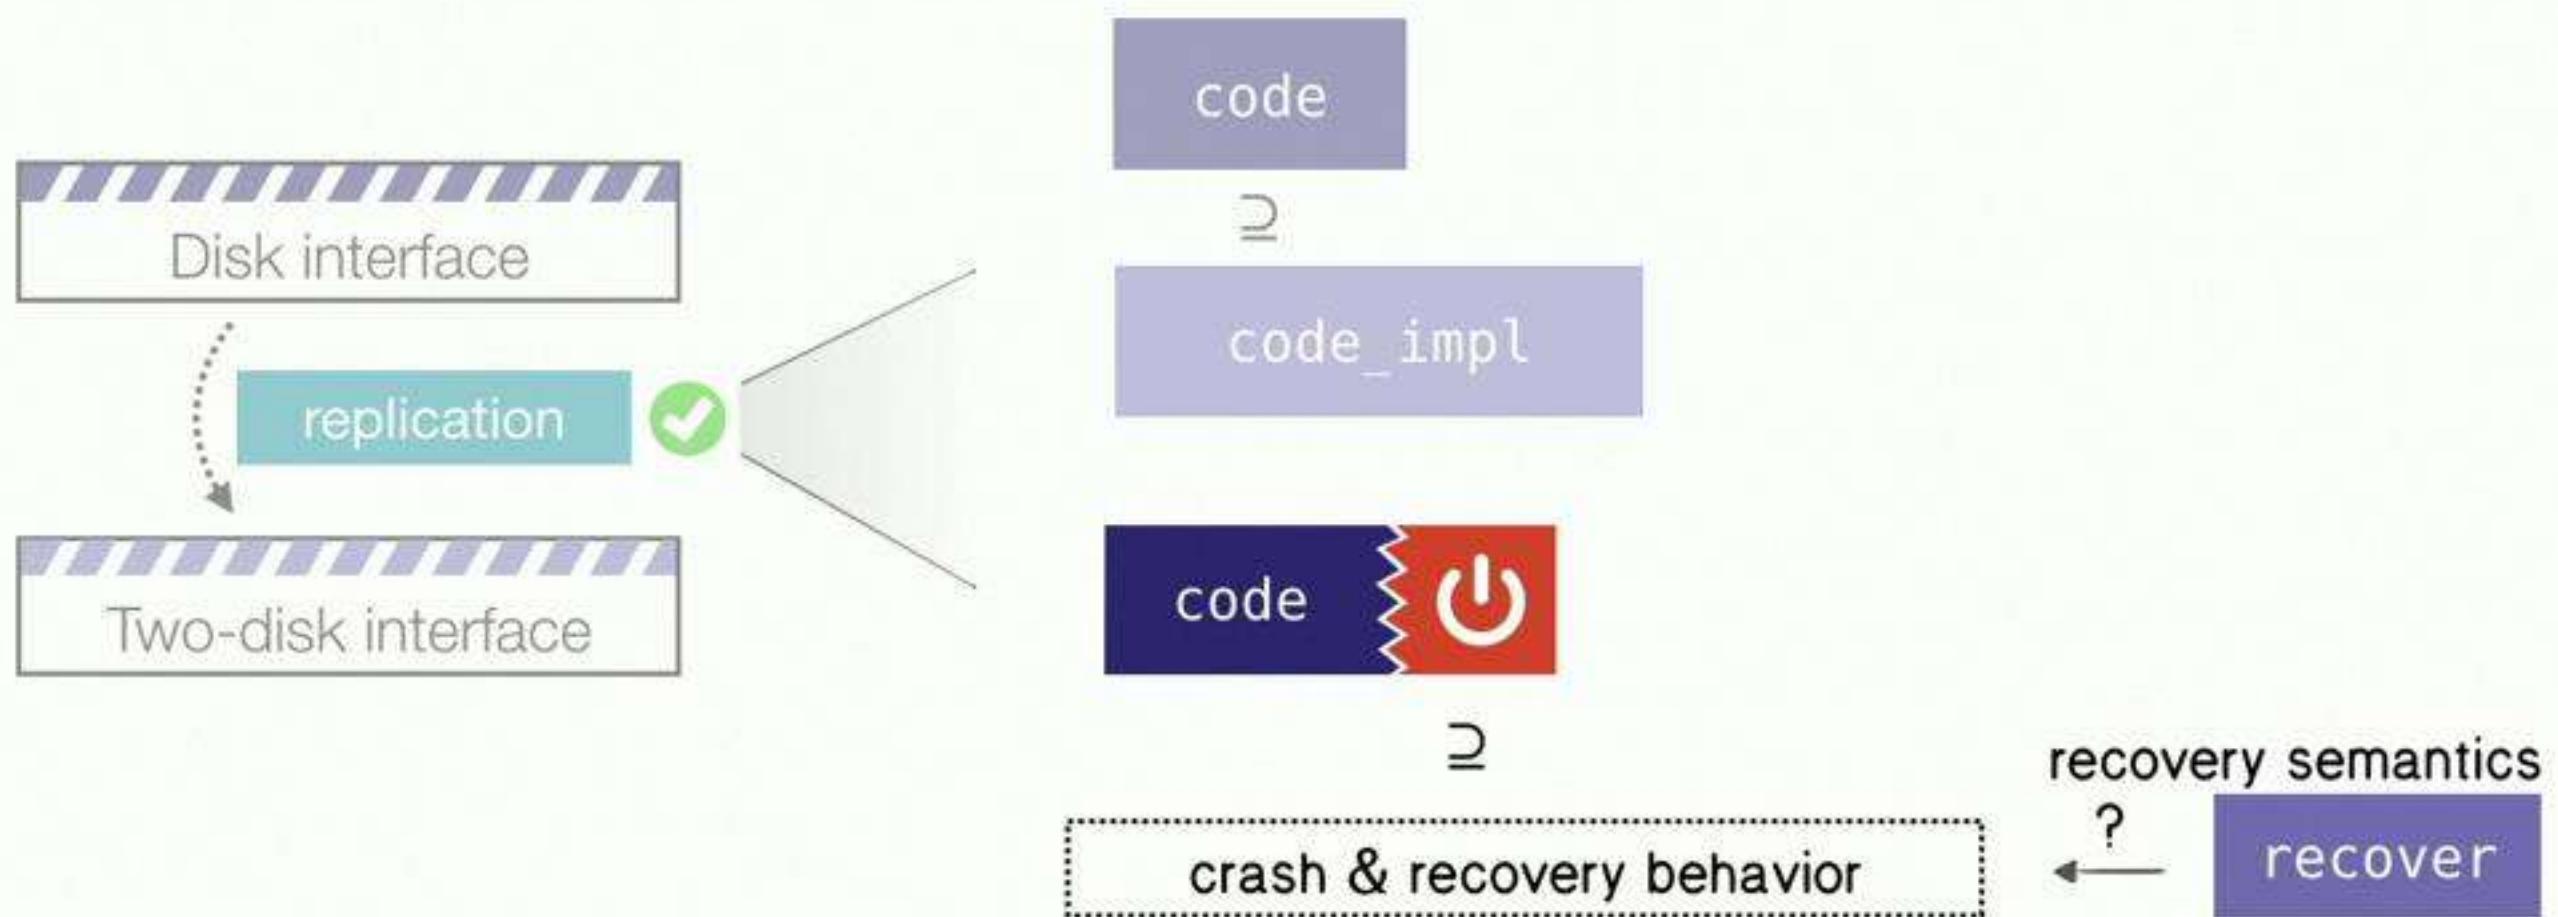
# Extending trace inclusion with recovery

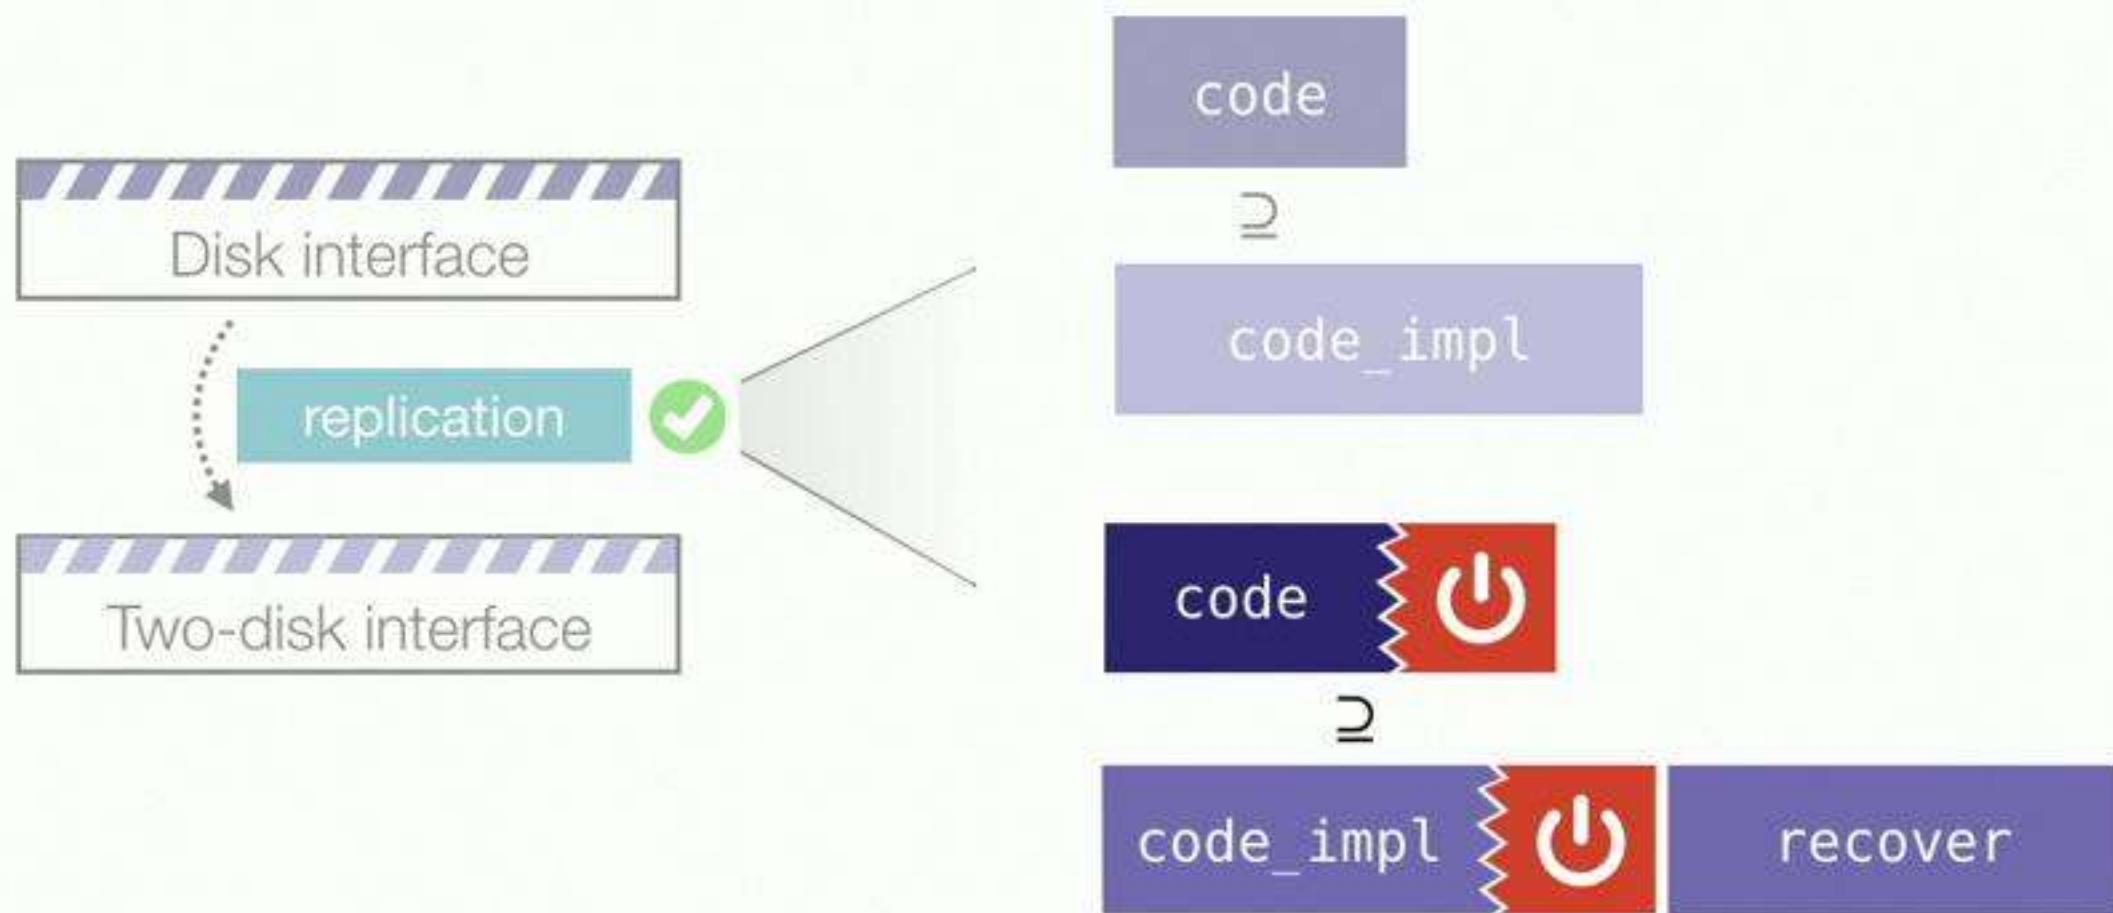


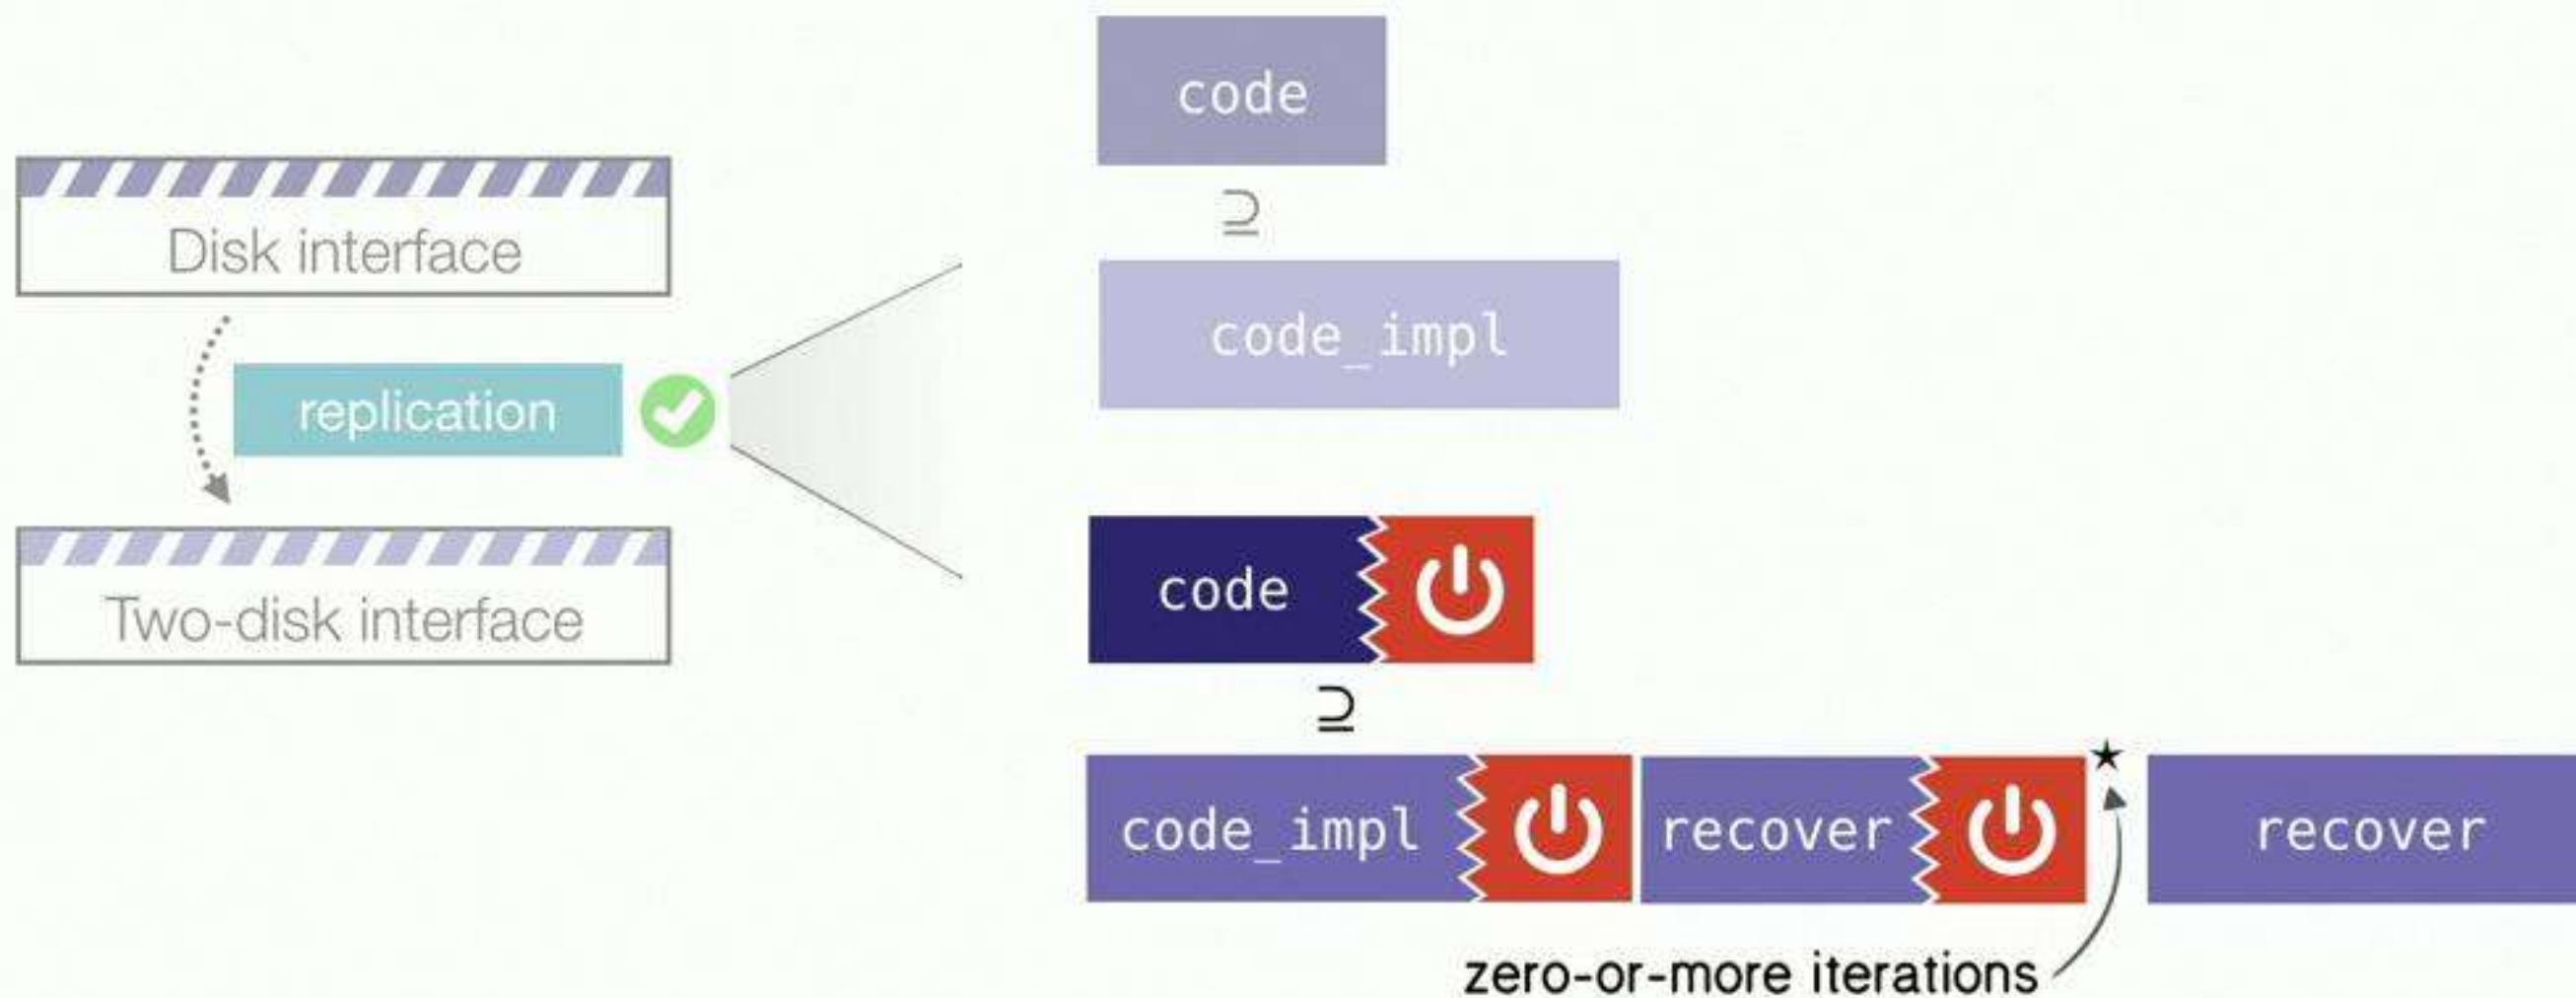
# Extending trace inclusion with recovery



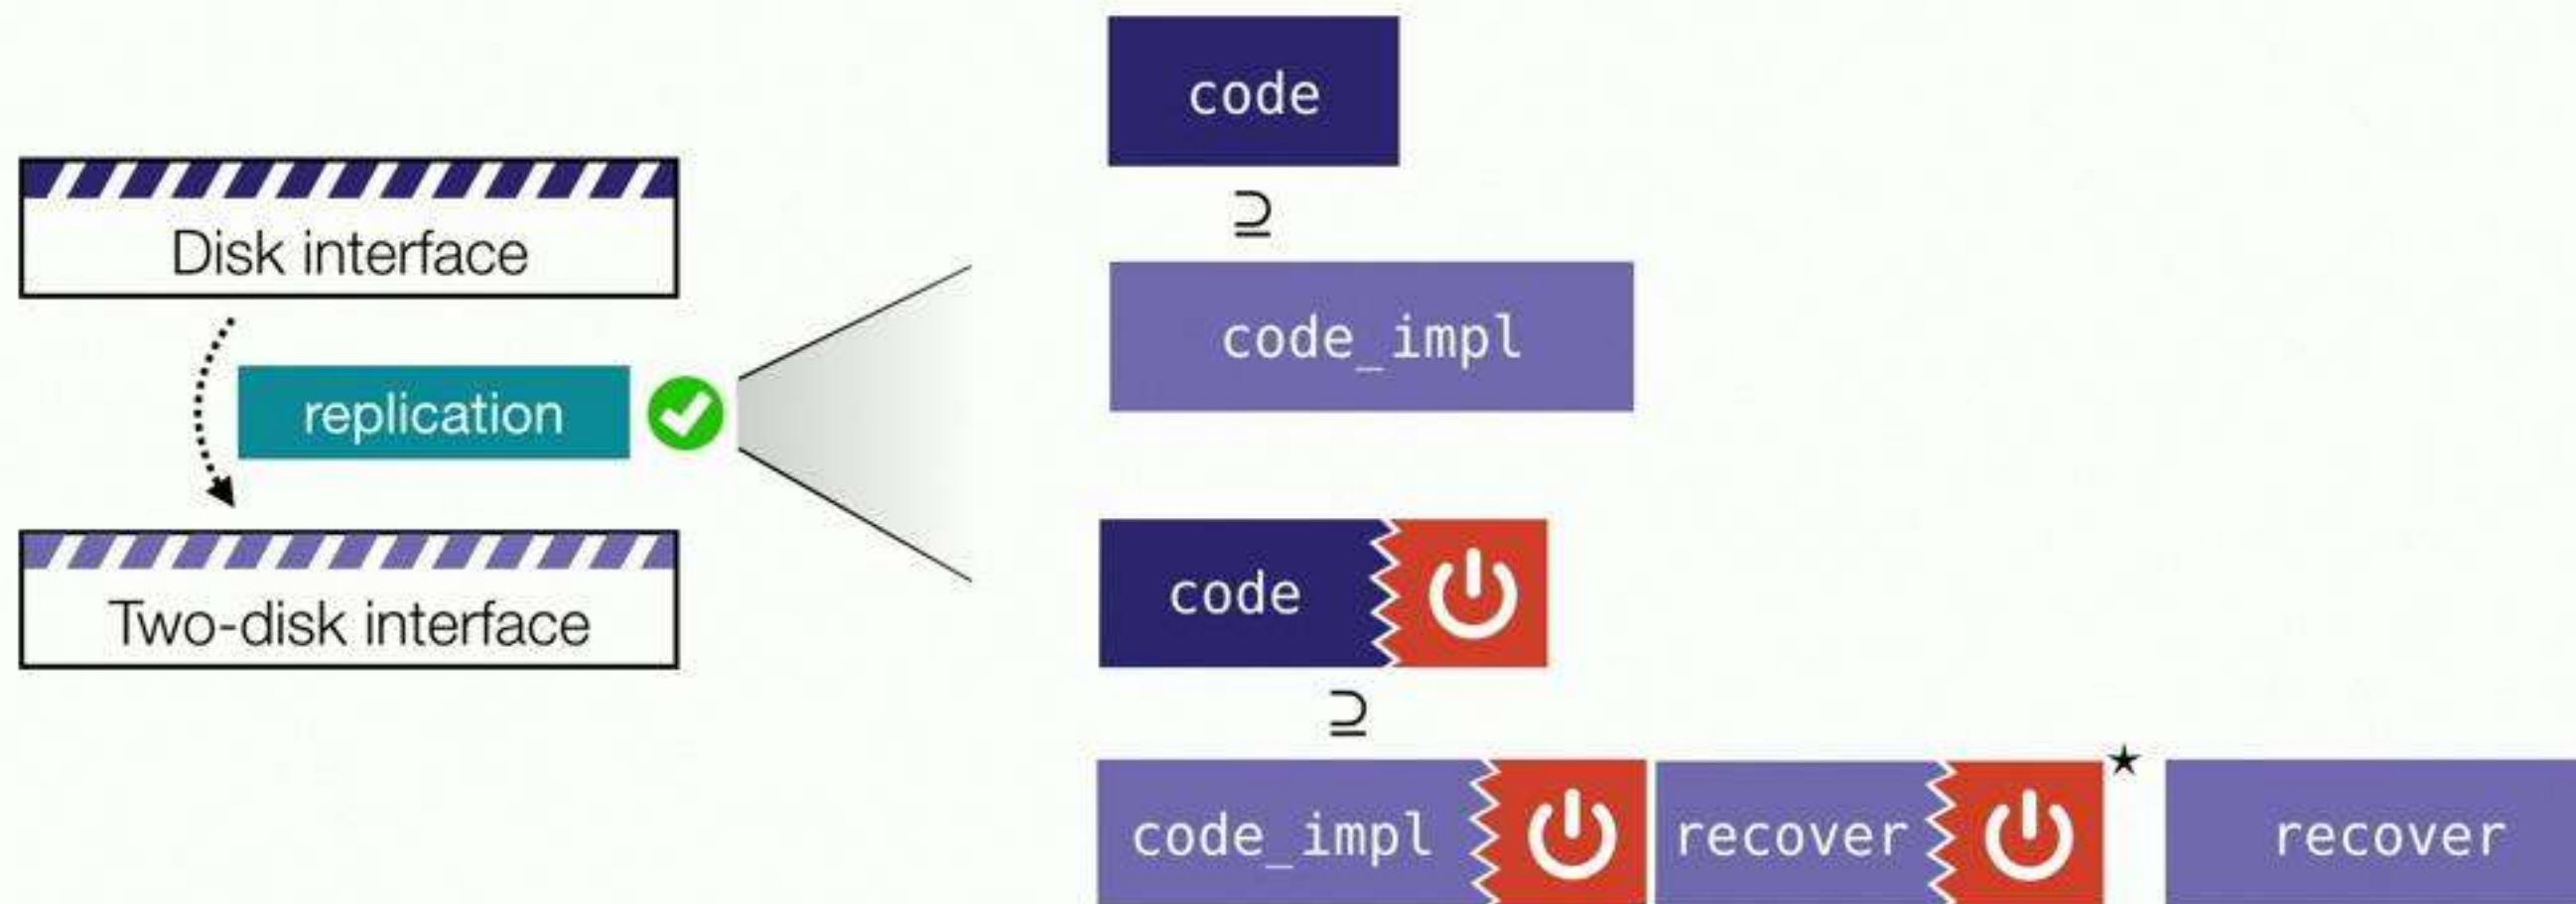








# Trace inclusion, with recovery



# Proving trace inclusion, with recovery



# Proving trace inclusion, with recovery



# Proving trace inclusion, with recovery



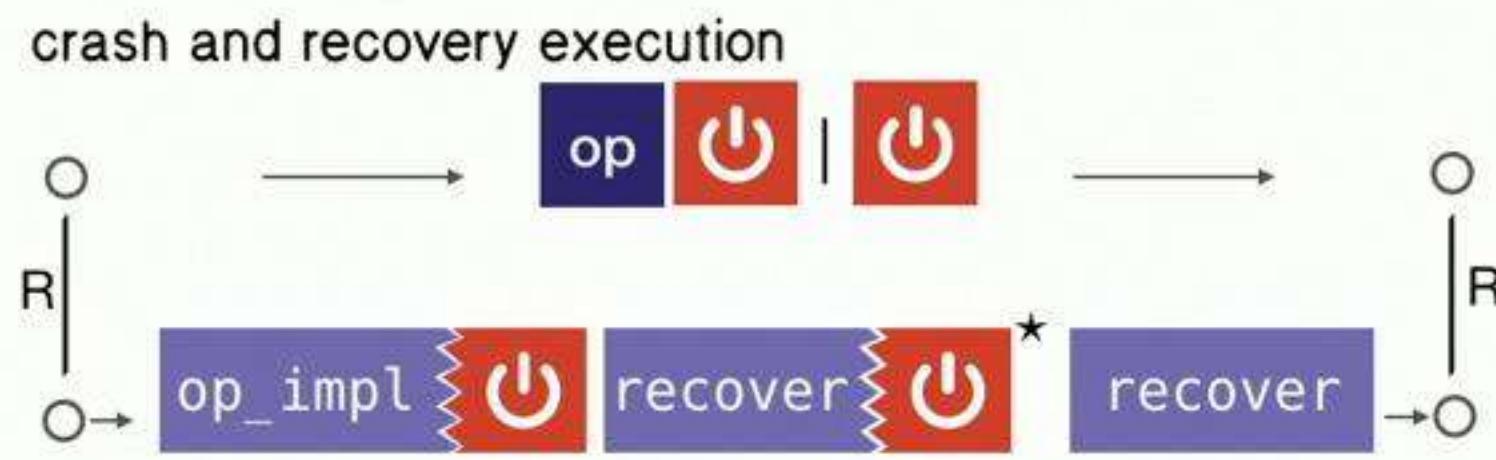
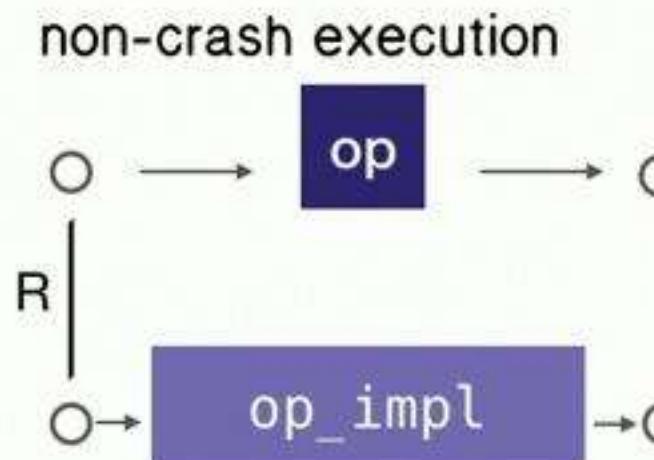
# Proving trace inclusion, with recovery



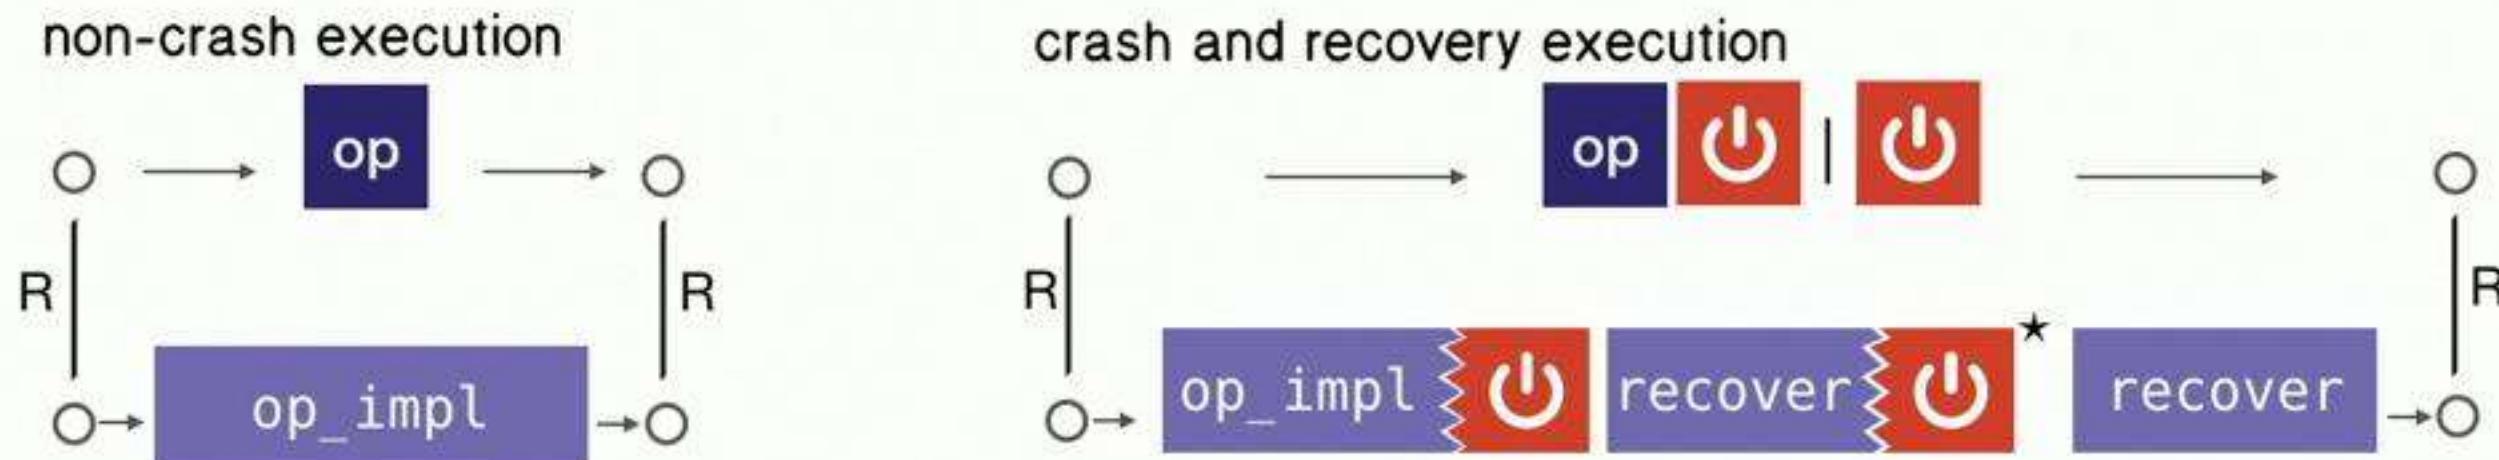
# Proving trace inclusion, with recovery



# Recovery refinement



# Recovery refinement



## Trace inclusion

implies

specification behavior  
 $\supseteq$   
running code behavior

# Composition theorem

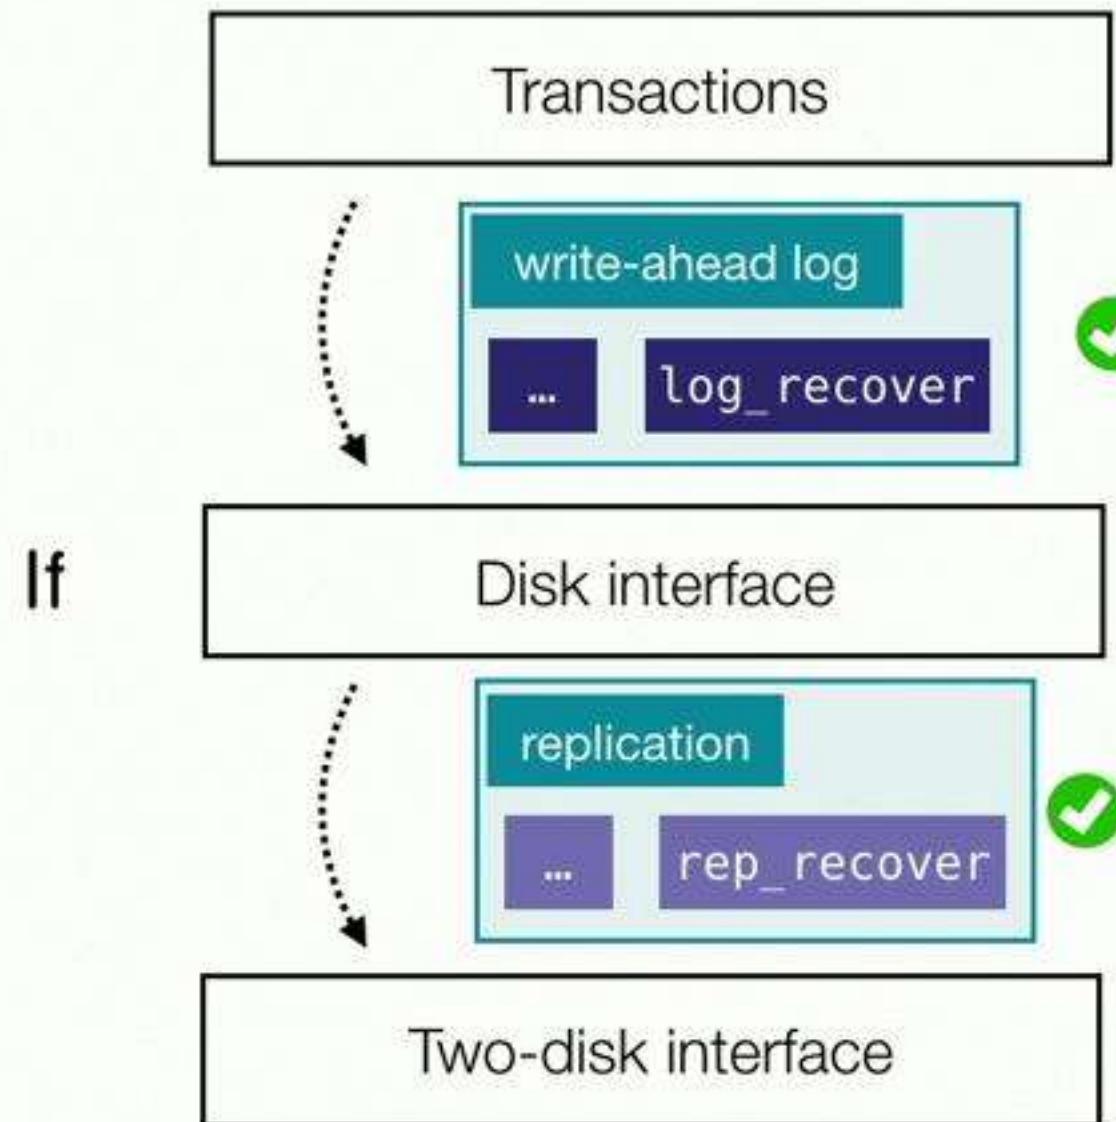
# Kleene algebra for transition relations



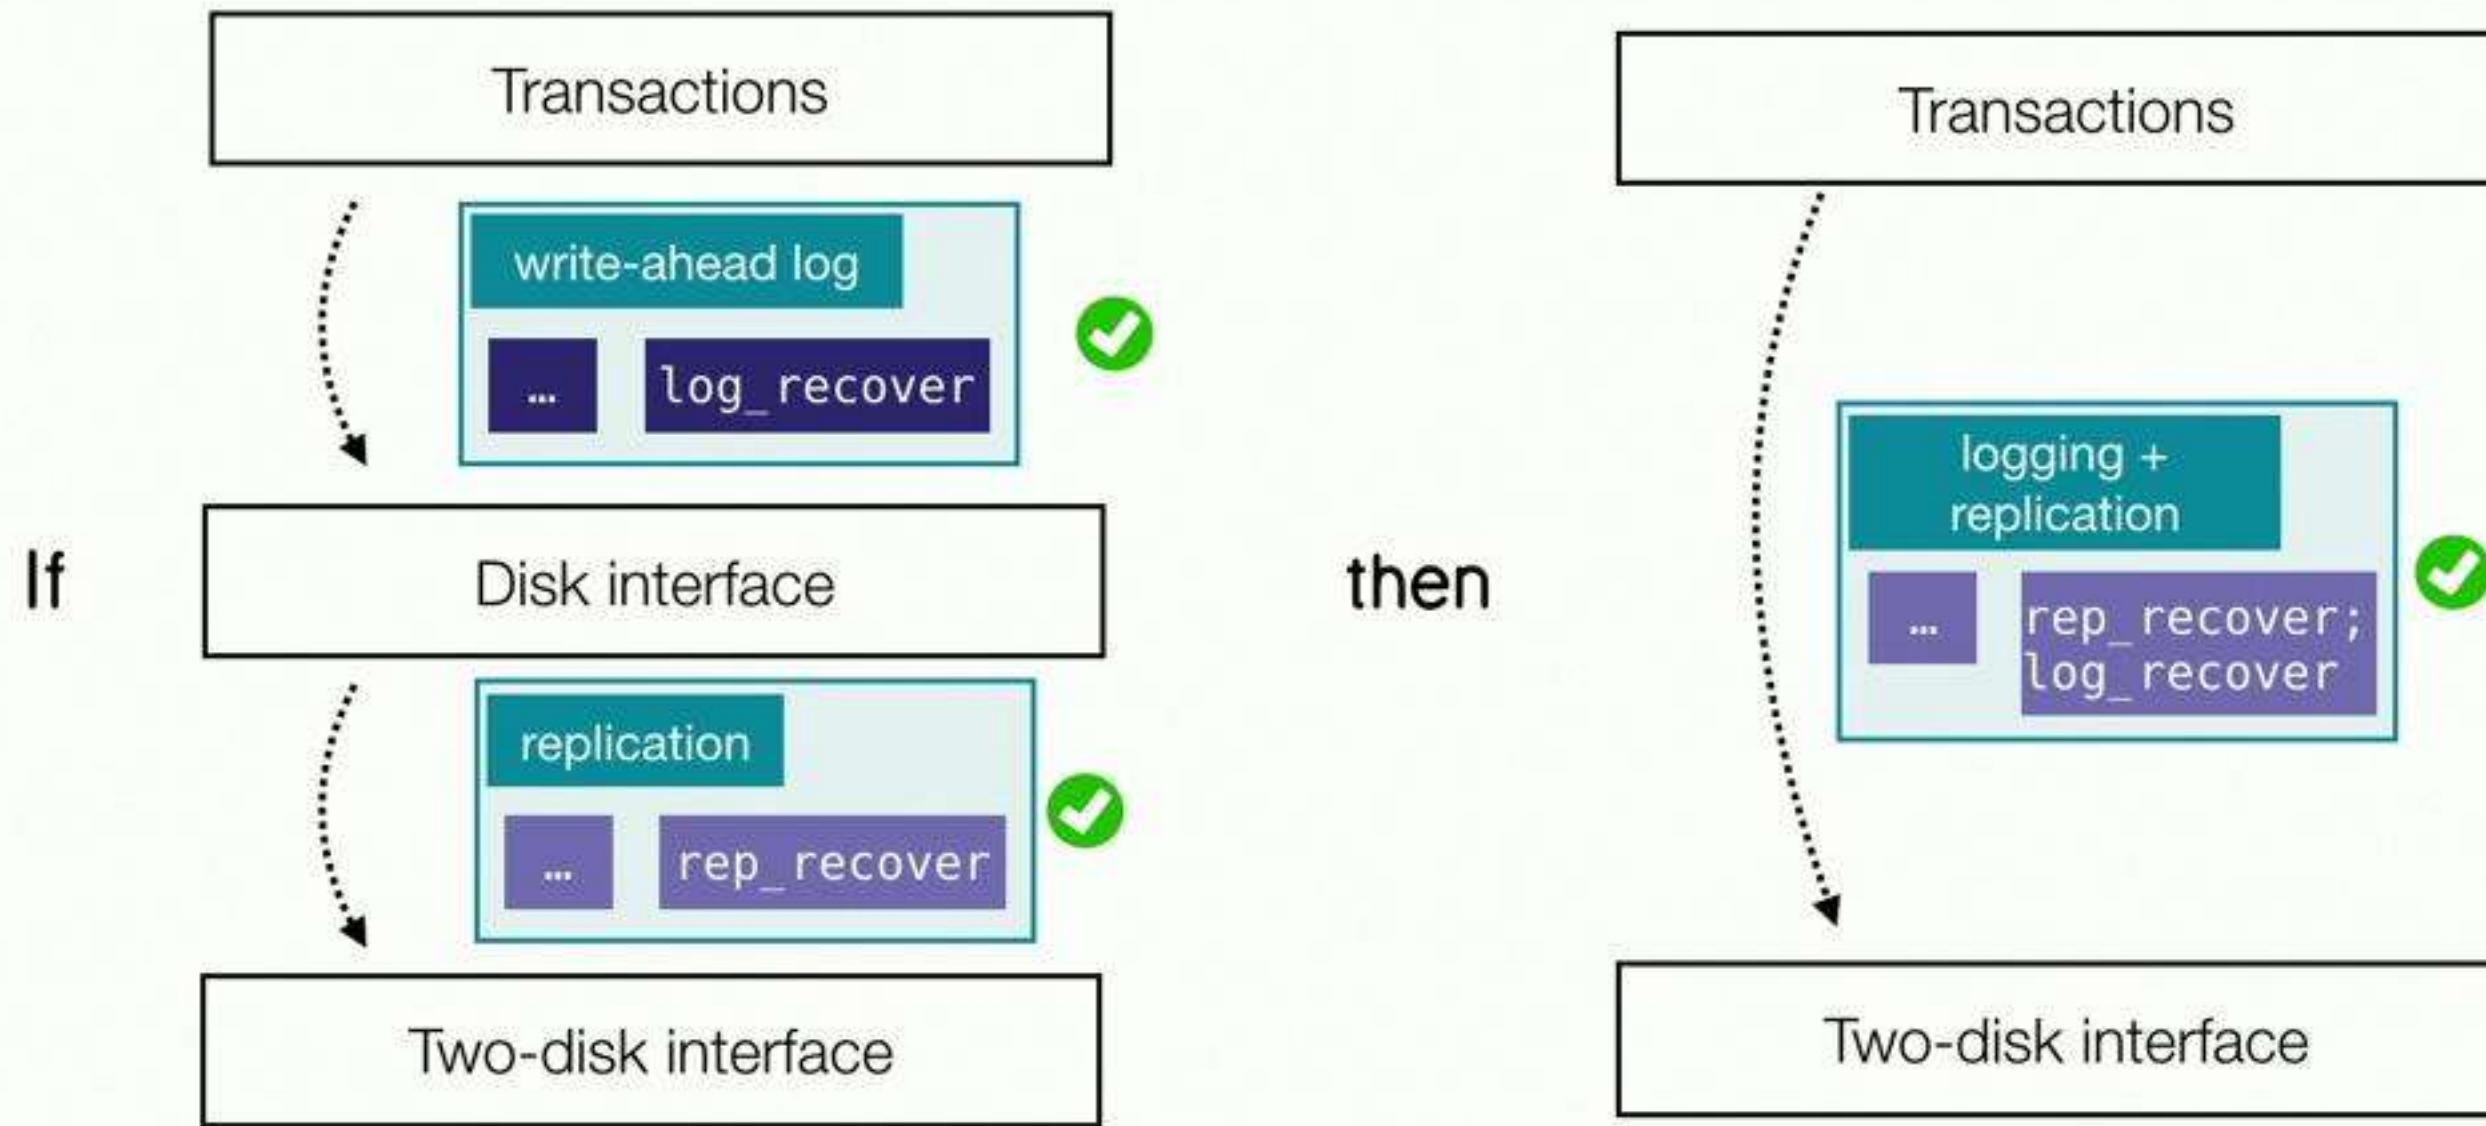
# Kleene algebra for transition relations

expression	matching transitions
$op_1 \mid op_2$	$O \rightarrow op_1 \rightarrow O \rightarrow op_2 \rightarrow O$
$\text{power} \mid op \text{ power}$	$O \rightarrow \text{power} \rightarrow O$ $O \rightarrow op \rightarrow O \rightarrow \text{power} \rightarrow O$
$r \xrightarrow{*} \text{power}^*$	$O \xrightarrow{*} r \xrightarrow{*} \text{power}^* \rightarrow O$ $O \rightarrow r \xrightarrow{*} \text{power}^* \rightarrow O$ $O \rightarrow r \xrightarrow{*} \text{power}^* \rightarrow O \rightarrow r \xrightarrow{*} \text{power}^* \rightarrow O$ $\dots$

# Theorem: recovery refinements compose



# Theorem: recovery refinements compose



# Goal: prove composed recovery correct

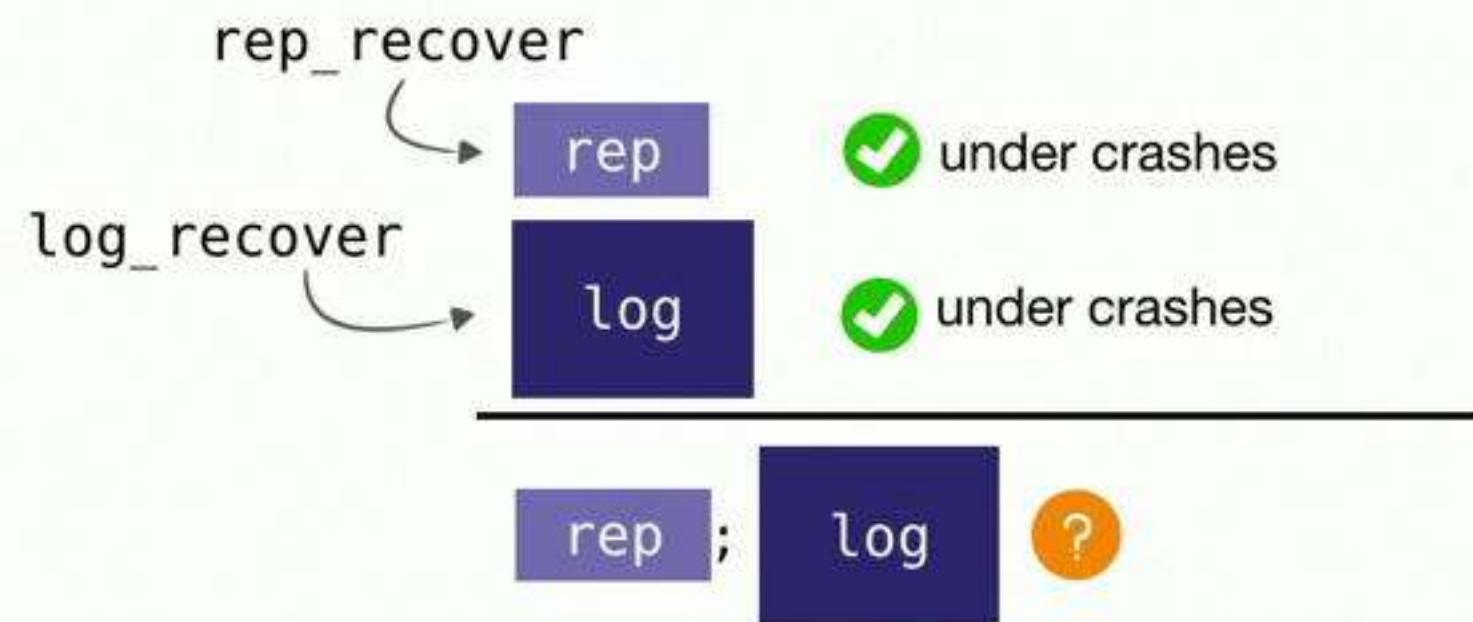
rep\_recover ✓ under crashes

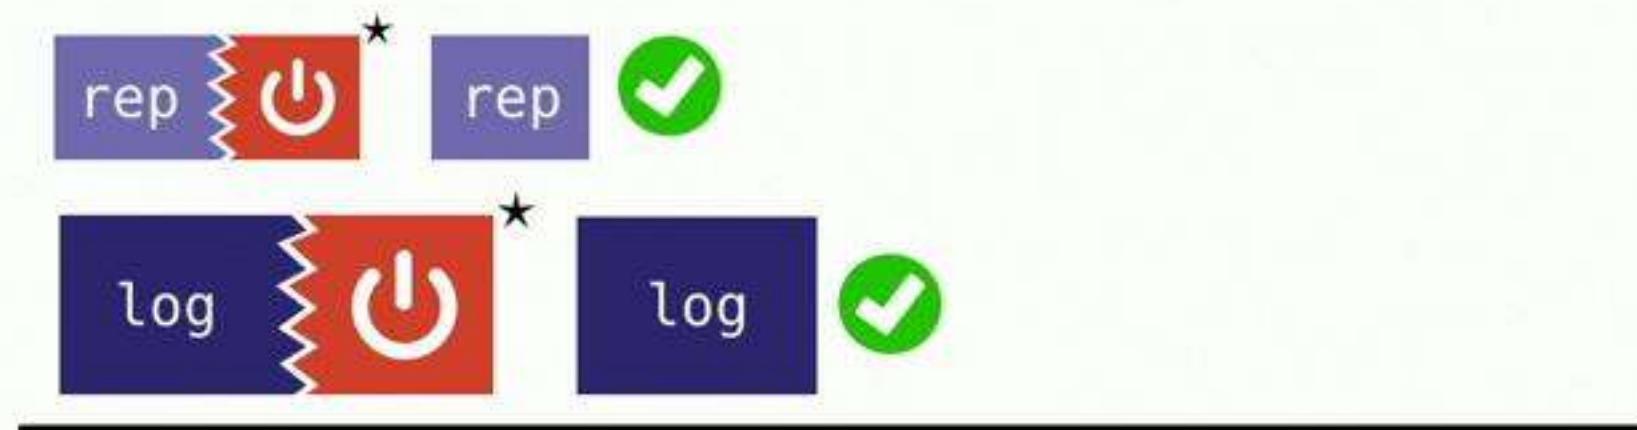
log\_recover ✓ under crashes

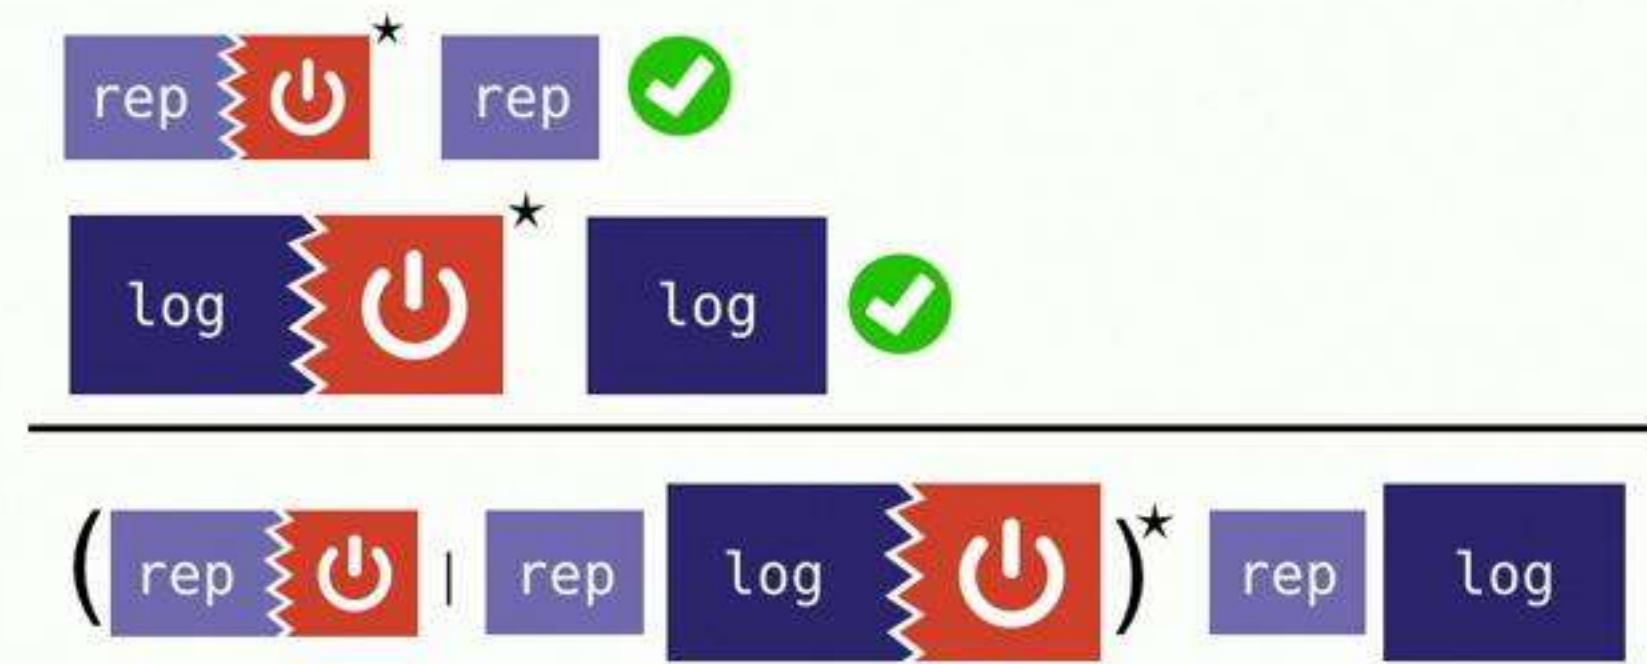
---

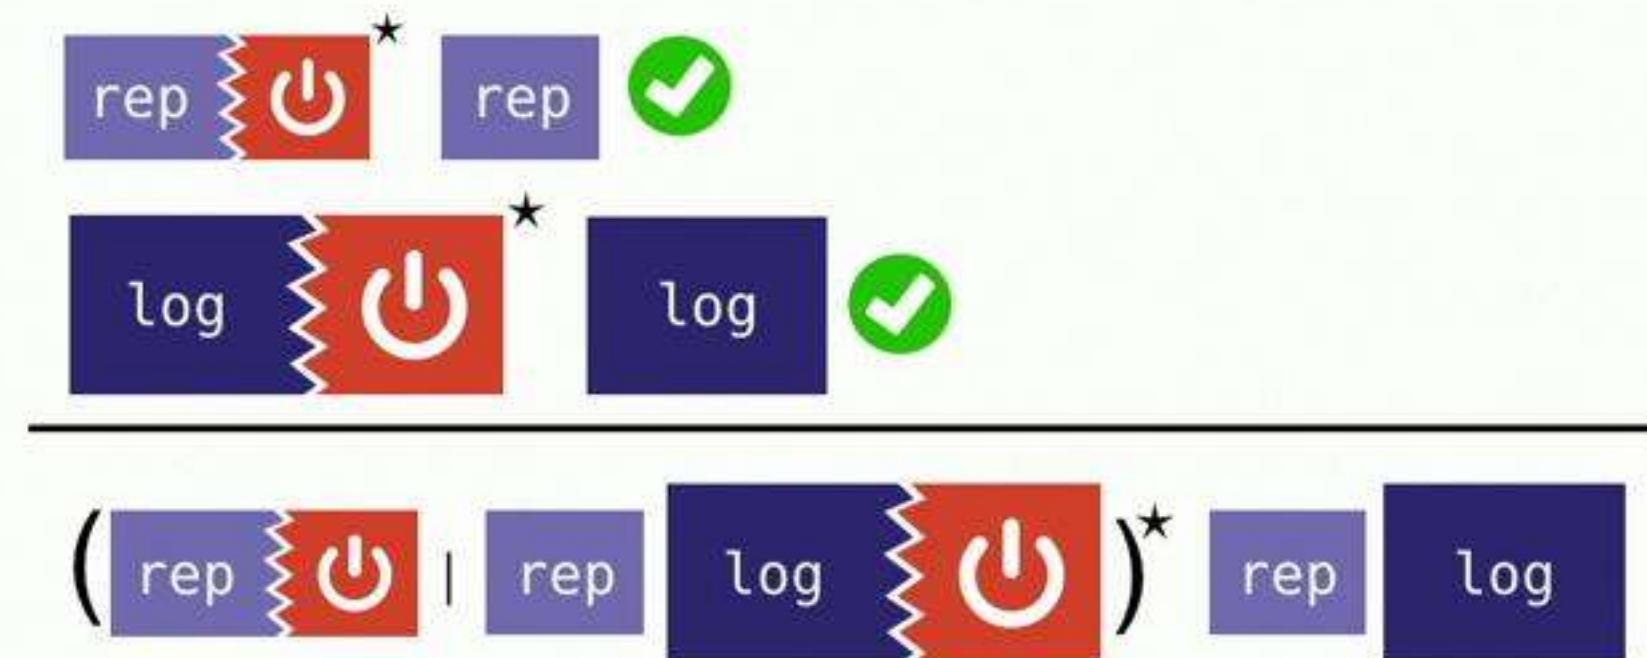
rep\_recover ; log\_recover ?

# Goal: prove composed recovery correct









how to re-use recovery proofs here?

# Using Kleene algebra for reasoning

( rep  | rep log  )<sup>\*</sup> rep log

after *de-nesting*  $(p \mid q)^* = p^* (qp^*)^*$

# Using Kleene algebra for reasoning

$$(\text{rep } \text{power} \mid \text{rep log } \text{power})^* \text{ rep log}$$

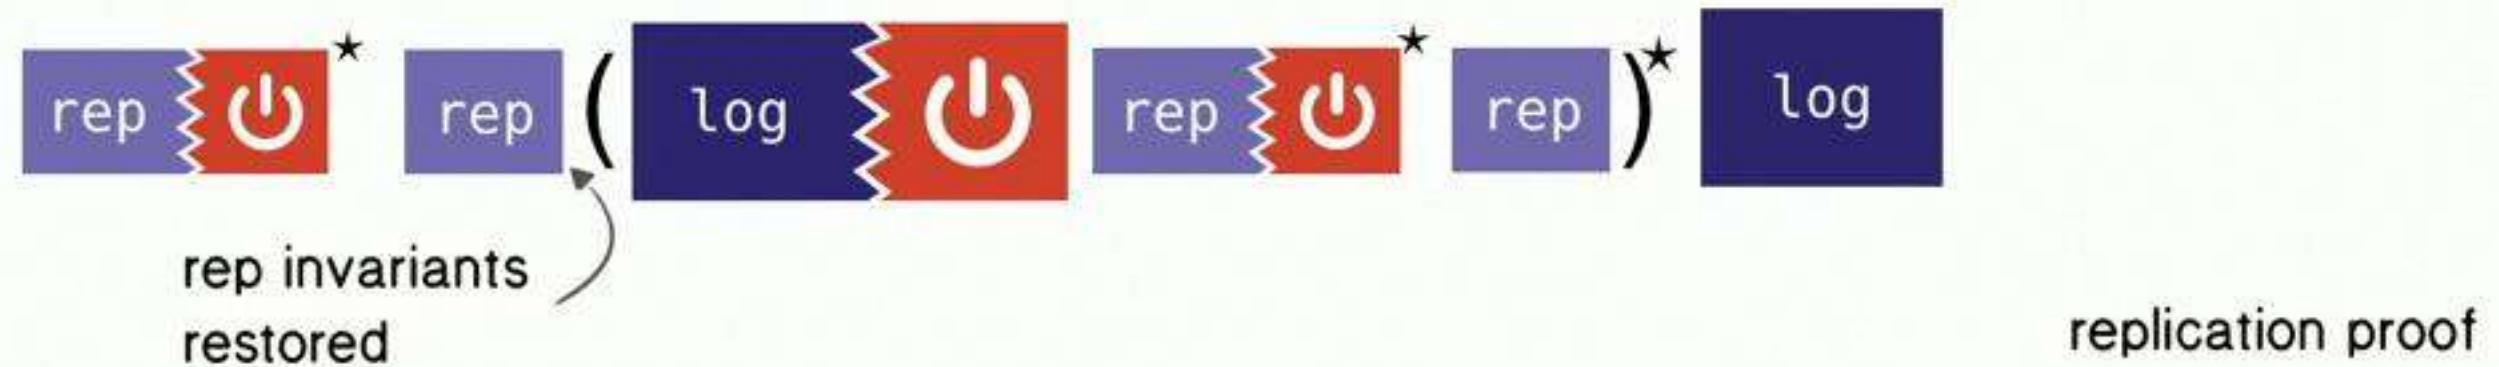
after *de-nesting*  $(p \mid q)^* = p^* (qp^*)^*$

$$= \text{rep power}^* (\text{rep log power} \text{ rep power}^*)^* \text{ rep log}$$

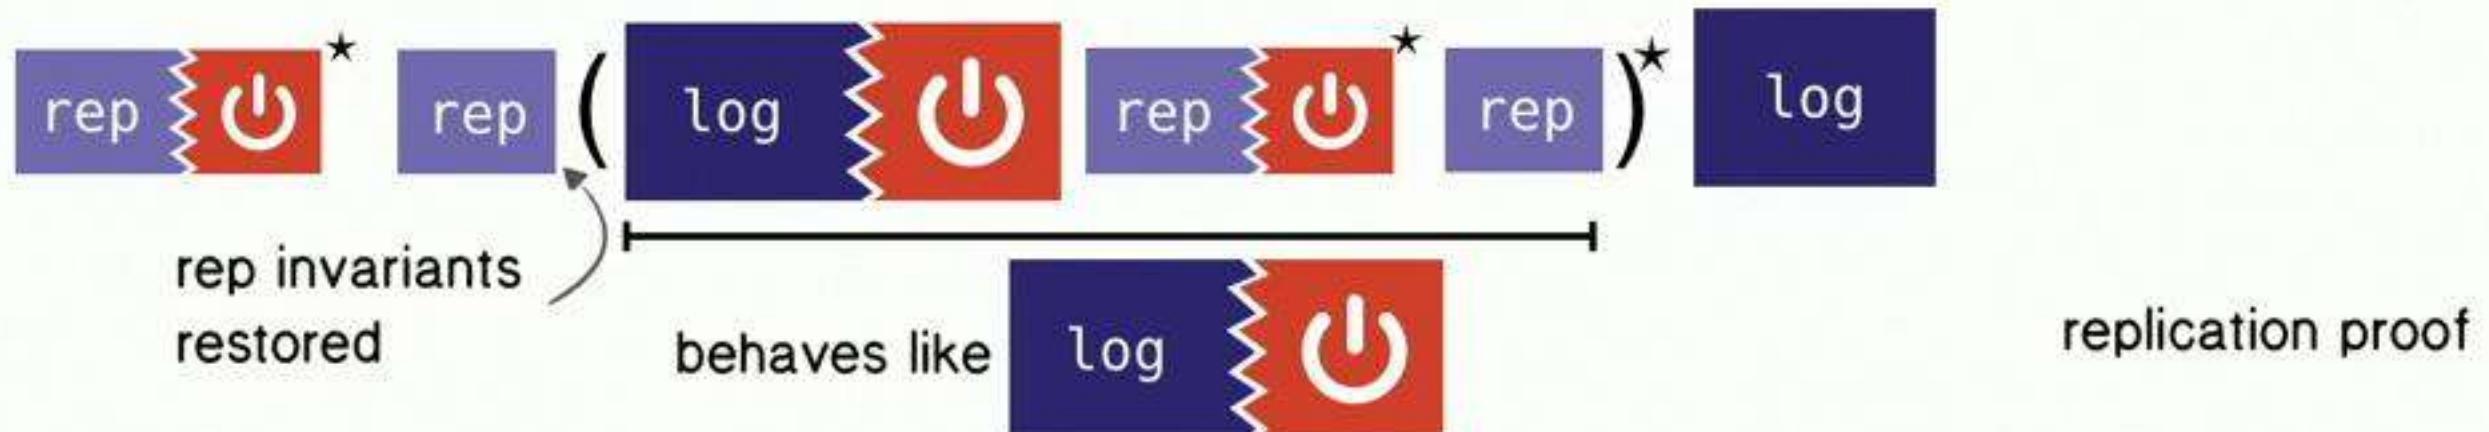
after *sliding*  $(pq)^* p = p(qp)^*$

$$= \text{rep power}^* \text{ rep } (\text{log power} \text{ rep power}^* \text{ rep})^* \text{ log}$$

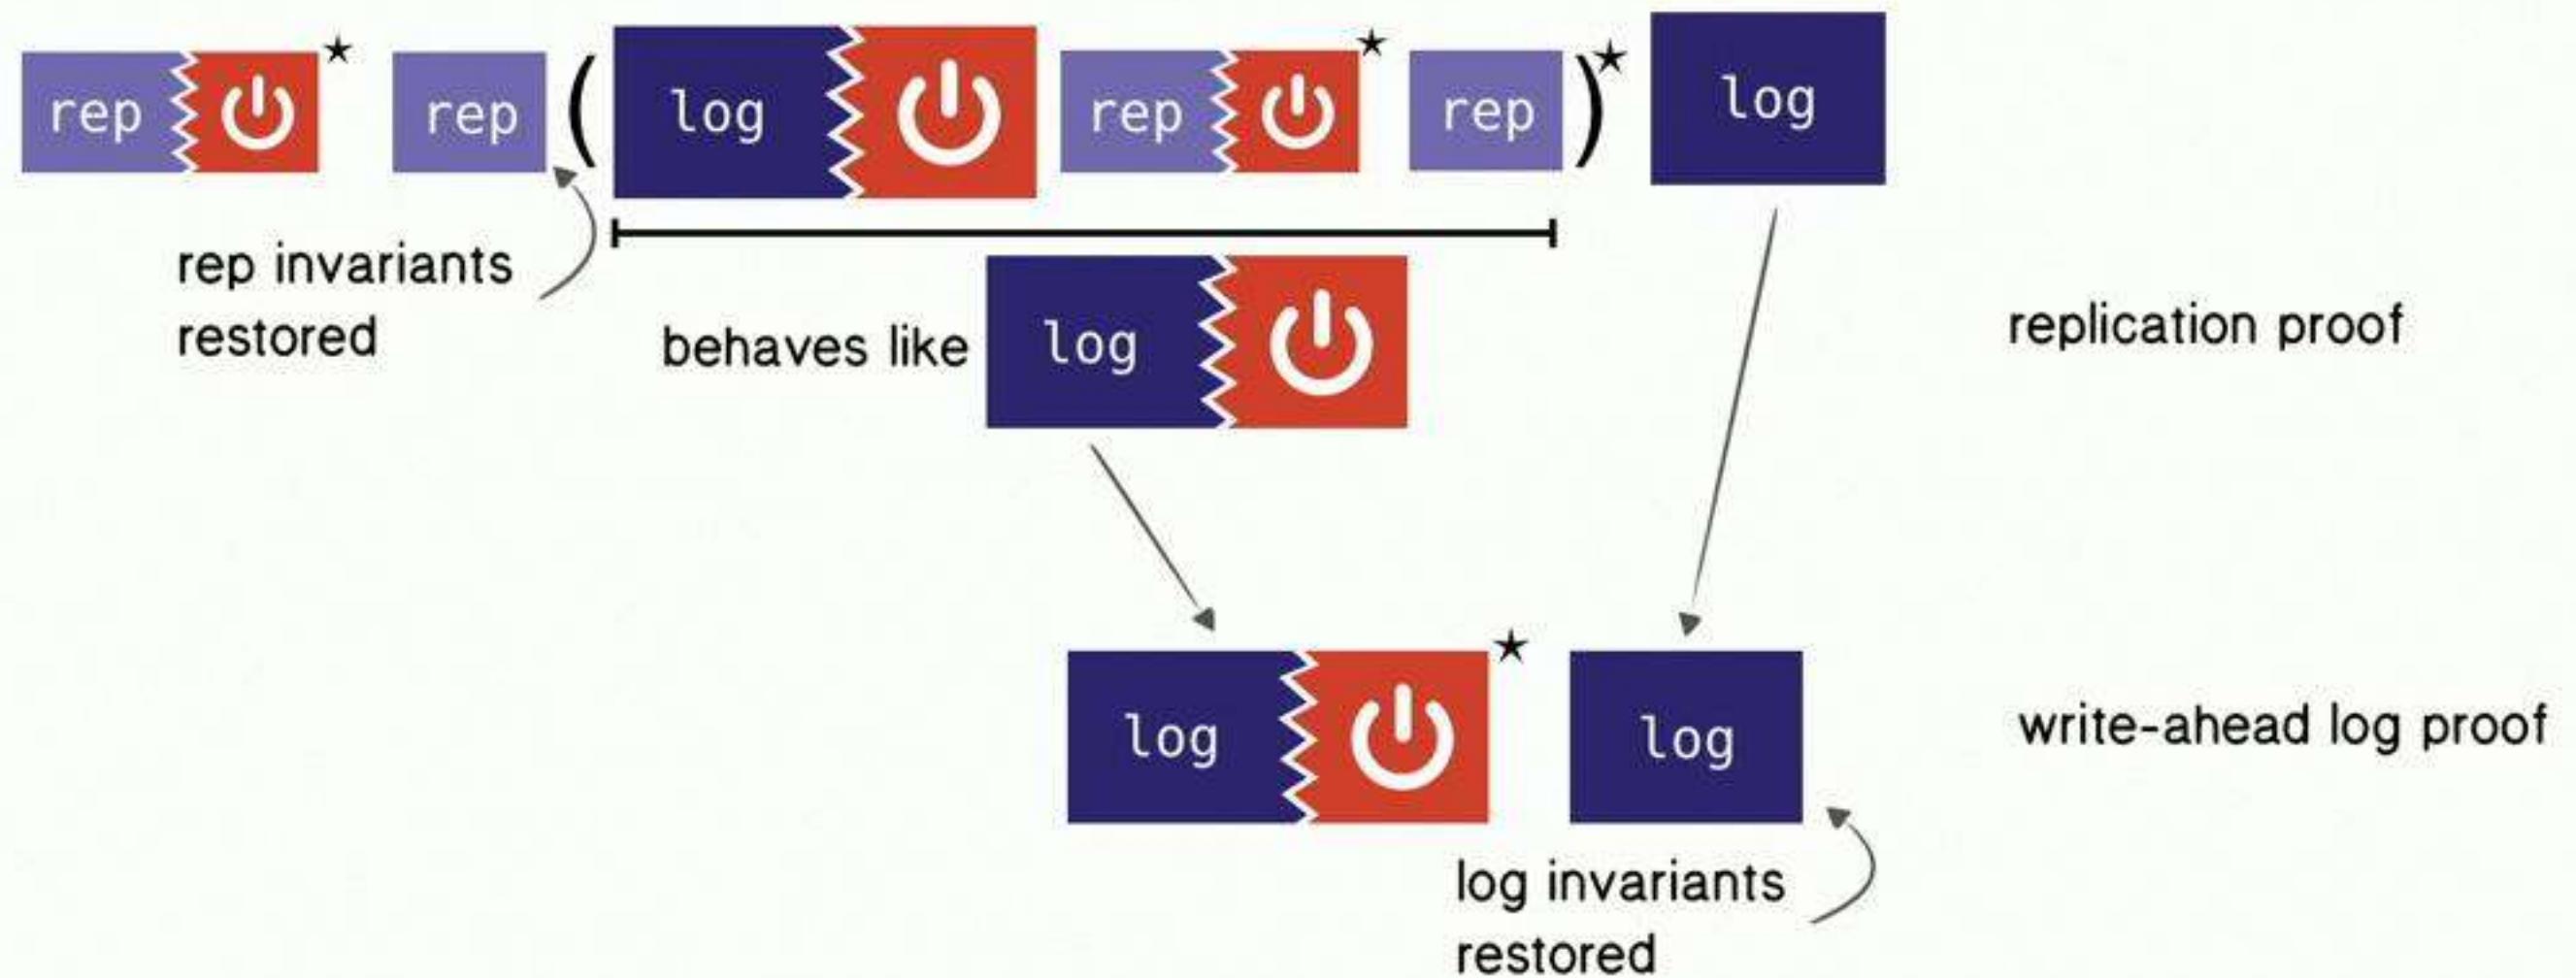
# After rewrite both proofs apply



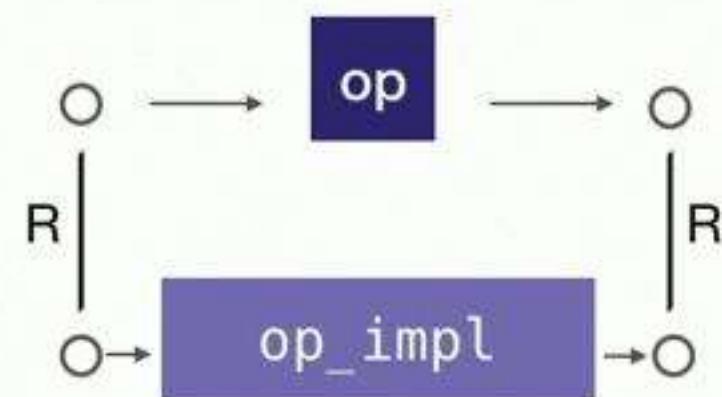
# After rewrite both proofs apply



# After rewrite both proofs apply



# Kleene algebra also helps with refinement



# Kleene algebra also helps with refinement



# An anecdote about modularity

# An anecdote about modularity

## Using Crash Hoare Logic for Certifying the FSCQ File System

Haogang Chen, Daniel Ziegler, Tej Chajed, Adam Chlipala, M. Frans Kaashoek, and Nickolai Zeldovich  
MIT CSAIL

### Abstract

FSCQ is the first file system verified using the Coq proof assistant. Its specification and whose implementation avoids bugs found in other systems, such as performing memory barriers or forgetting to zero out memory that happens at an inopportune time.

## Verifying a high-performance crash-safe file system using a tree specification

Haogang Chen,<sup>†</sup> Tej Chajed, Alex Konradi,<sup>‡</sup> Stephanie Wang,<sup>§</sup> Atalay İleri,  
Adam Chlipala, M. Frans Kaashoek, Nickolai Zeldovich  
MIT CSAIL

### ABSTRACT

DFSCQ is the first file system that (1) provides a precise specification for `fsync` and `fdatasync`, which allow applications to achieve high performance and crash safety, and (2) provides a machine-checked proof that its implementation meets this specification. DFSCQ's specification captures the behavior of sophisticated optimizations, including log-

### 1 INTRODUCTION

File systems achieve high I/O performance and crash safety by implementing sophisticated optimizations to increase disk throughput. These optimizations include deferring writing buffered data to persistent storage, grouping many transactions into a single I/O operation, checksumming journal entries, and bypassing the write-ahead log when writing to

# An anecdote about modularity

## Using Crash Hoare Logic for Certifying the FSCQ File System

Haogang Chen, Daniel Ziegler, Tej Chajed, Adam Chlipala, M. Frans Kaashoek, and Nickolai Zeldovich  
MIT CSAIL

### Abstract

FSCQ is the first file system verified using the Coq proof assistant. Its specification and whose implementation avoids bugs in systems, such as performing barriers or forgetting to zero out memory that happens at an inopportune time.

## Verifying a high-performance crash-safe file system using a tree specification (my advisor)

Haogang Chen,<sup>†</sup> Tej Chajed, Alex Konradi,<sup>‡</sup> Stephanie Wang,<sup>§</sup> Atalay İleri,  
Adam Chlipala, M. Frans Kaashoek, Nickolai Zeldovich  
MIT CSAIL

### ABSTRACT

DFSCQ is the first file system that (1) provides a precise specification for `fsync` and `fdatasync`, which allow applications to achieve high performance and crash safety, and (2) provides a machine-checked proof that its implementation meets this specification. DFSCQ's specification captures the behavior of sophisticated optimizations, including log-optimization, write-ahead logging, and journaling.

### 1 INTRODUCTION

File systems achieve high I/O performance and crash safety by implementing sophisticated optimizations to increase disk throughput. These optimizations include deferring writing buffered data to persistent storage, grouping many transactions into a single I/O operation, checksumming journal entries, and bypassing the write-ahead log when writing to

```
def atomic_save(data, path):
    write_all(data, tmp)
    rename(tmp, path)

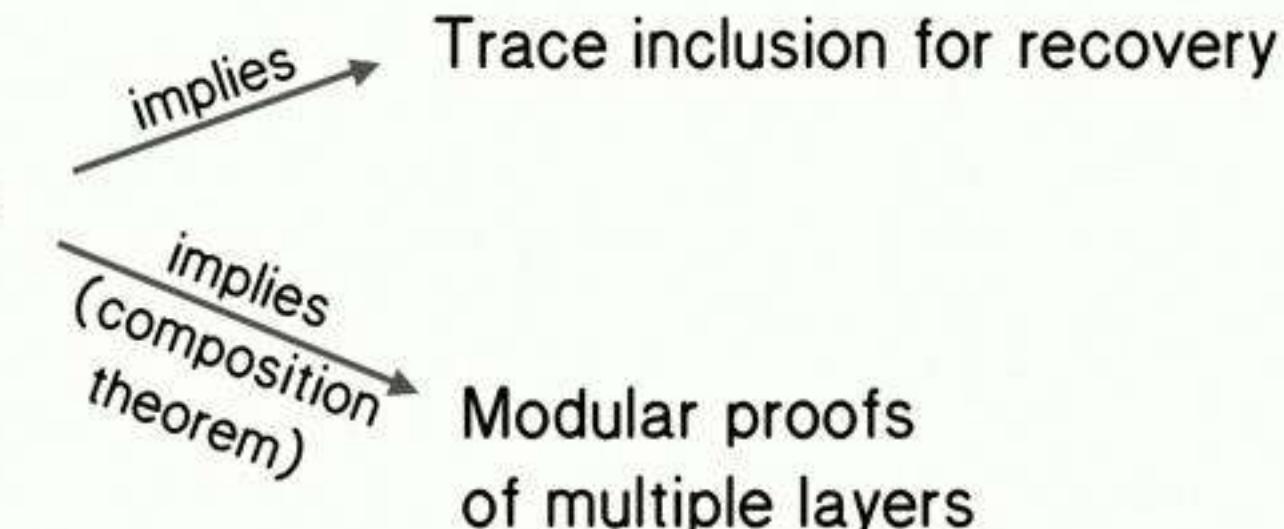
# runs on crash
def recover():
    fs_recover()
    unlink(tmp)
```

```
def atomic_save(data, path):  
    write_all(data, tmp)  
    rename(tmp, path)  
  
    # runs on crash  
def recover(): ← this is non-modular and makes  
    fs_recover() the proof much harder  
    unlink(tmp)
```

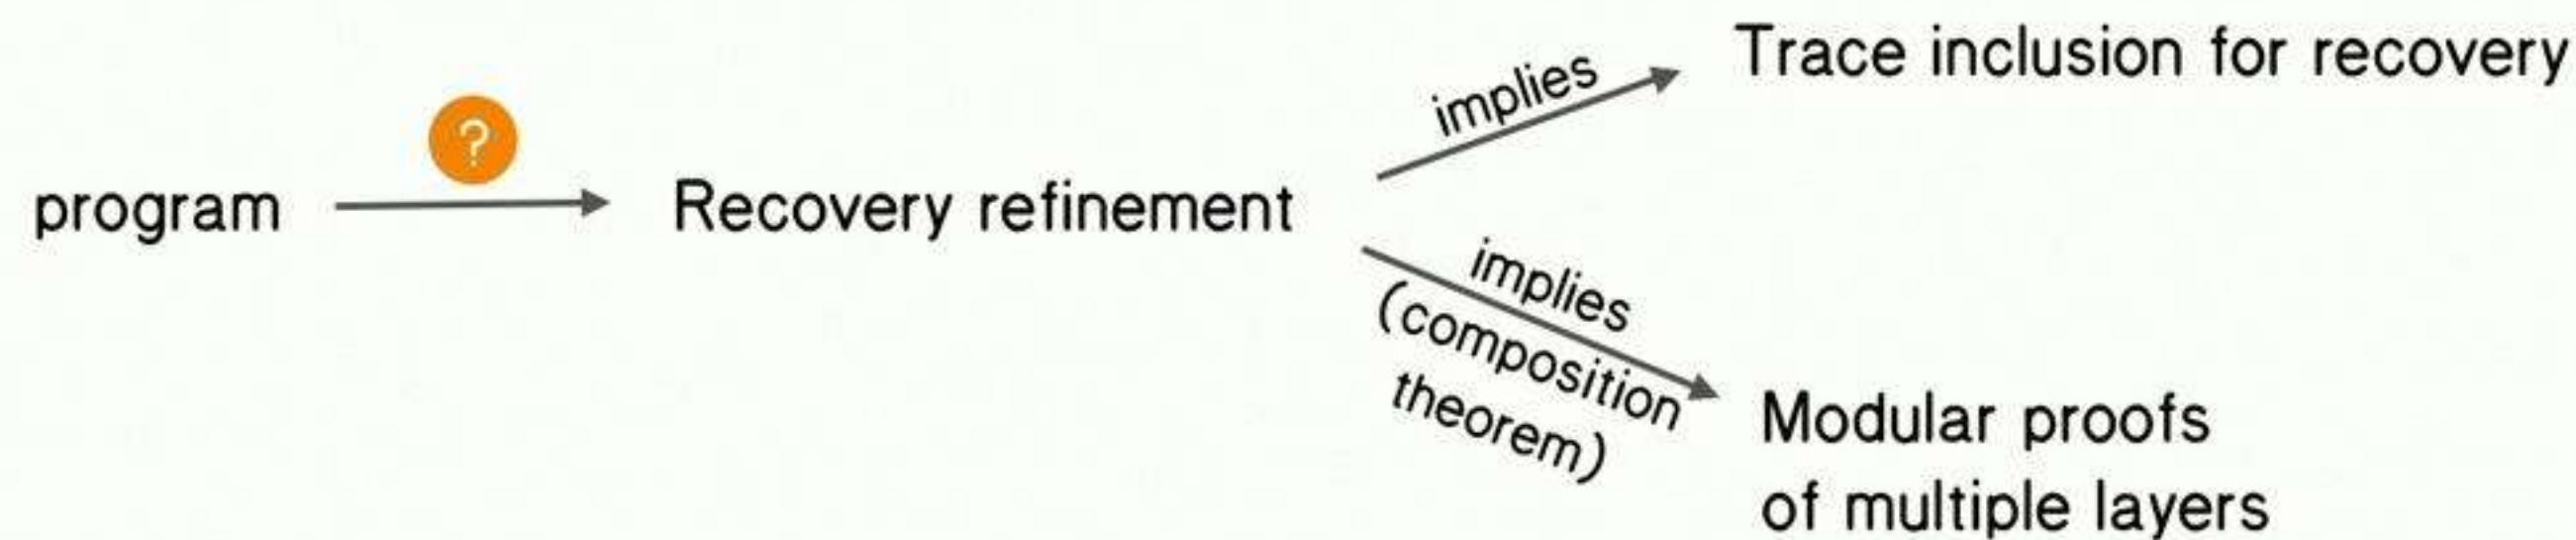
Proving this code correct took 1500 lines of proof code!

# Argosy so far

Recovery refinement



# Argosy so far



# Crash Hoare Logic

# Hoare Logic

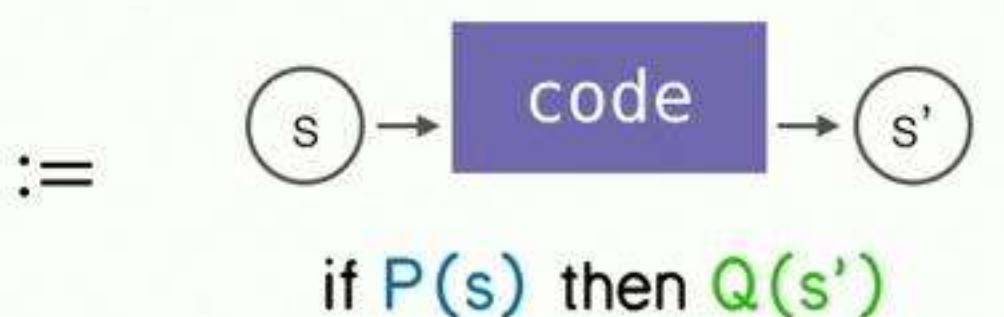
# Hoare Logic

“Hoare triple”

$\{P\}$  code  $\{Q\}$

precondition

postcondition



# Crash Hoare Logic



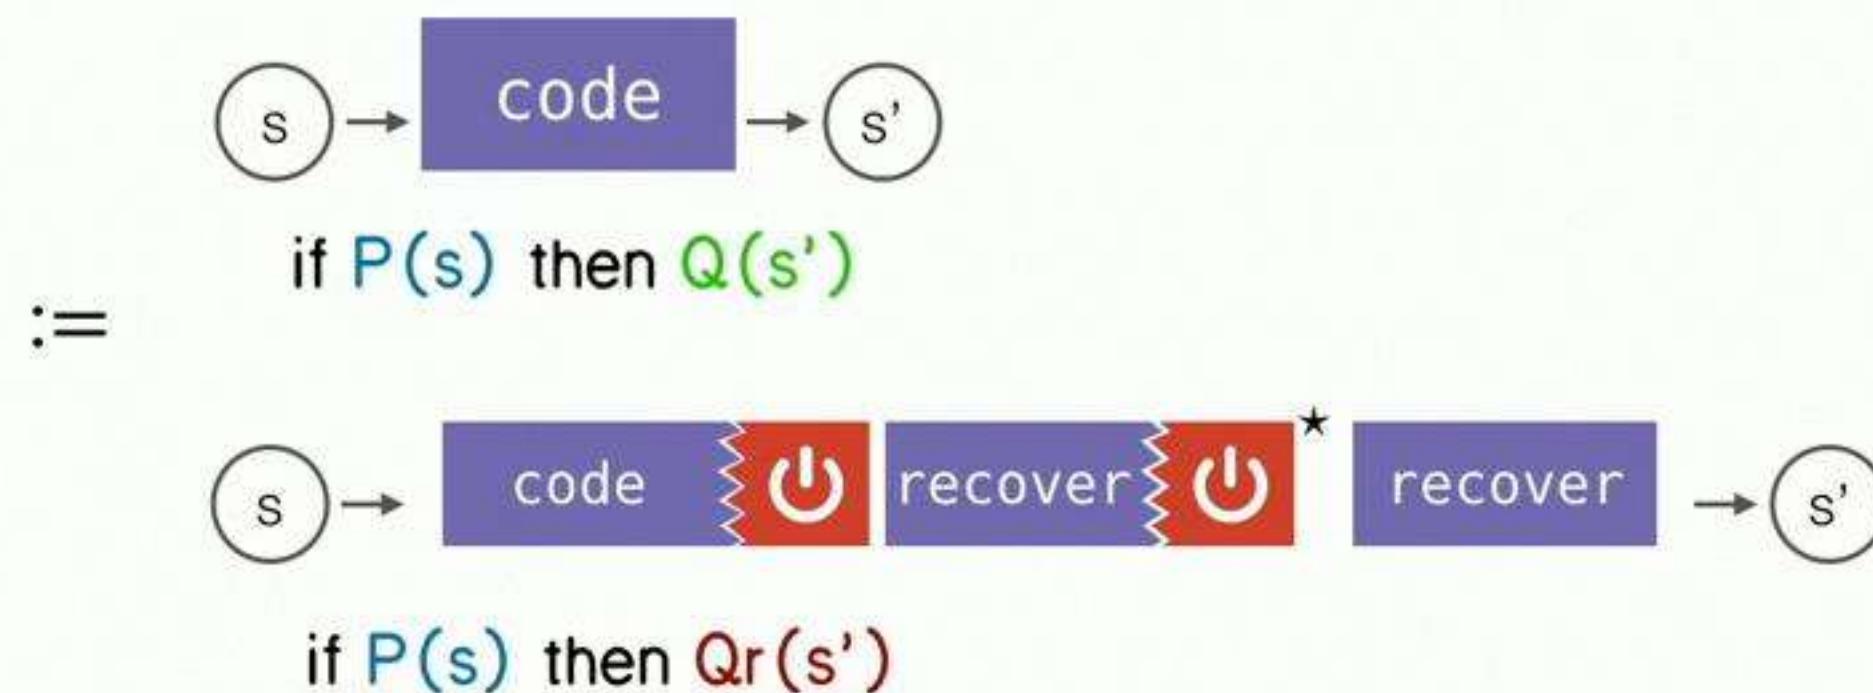
:=      if  $P(s)$  then  $Q(s')$



if  $P(s)$  then  $Q_c(s')$

# Crash Hoare Logic

“recovery specification”  $\{P\}$  code  $\circlearrowleft$  recover  $\{Q\}$   $\{Q_r\}$   
precondition postcondition recovery postcondition



# Crash Hoare Logic



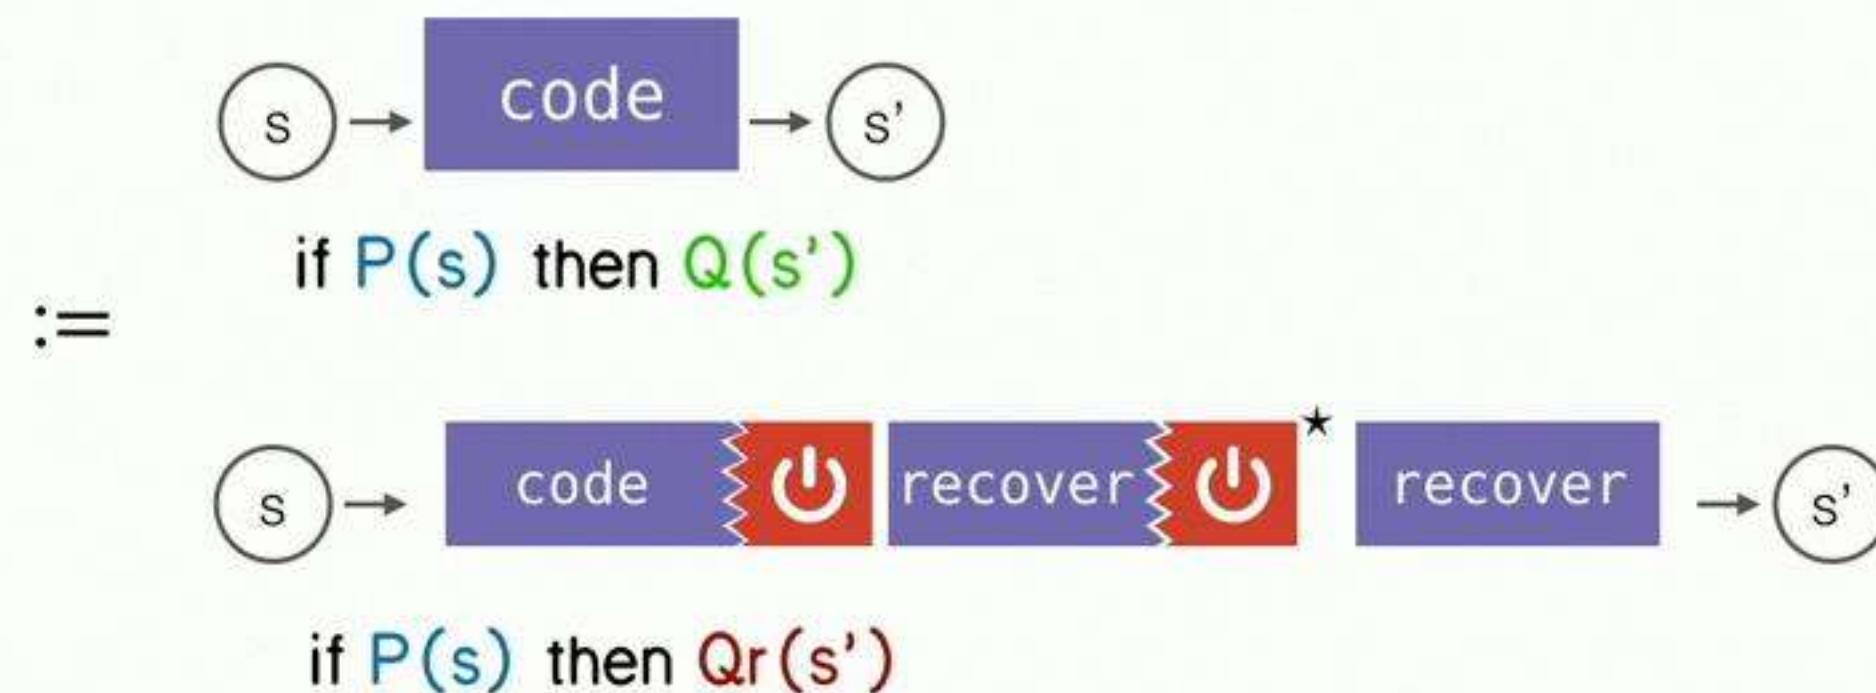
:=      if  $P(s)$  then  $Q(s')$



if  $P(s)$  then  $Qc(s')$

# Crash Hoare Logic

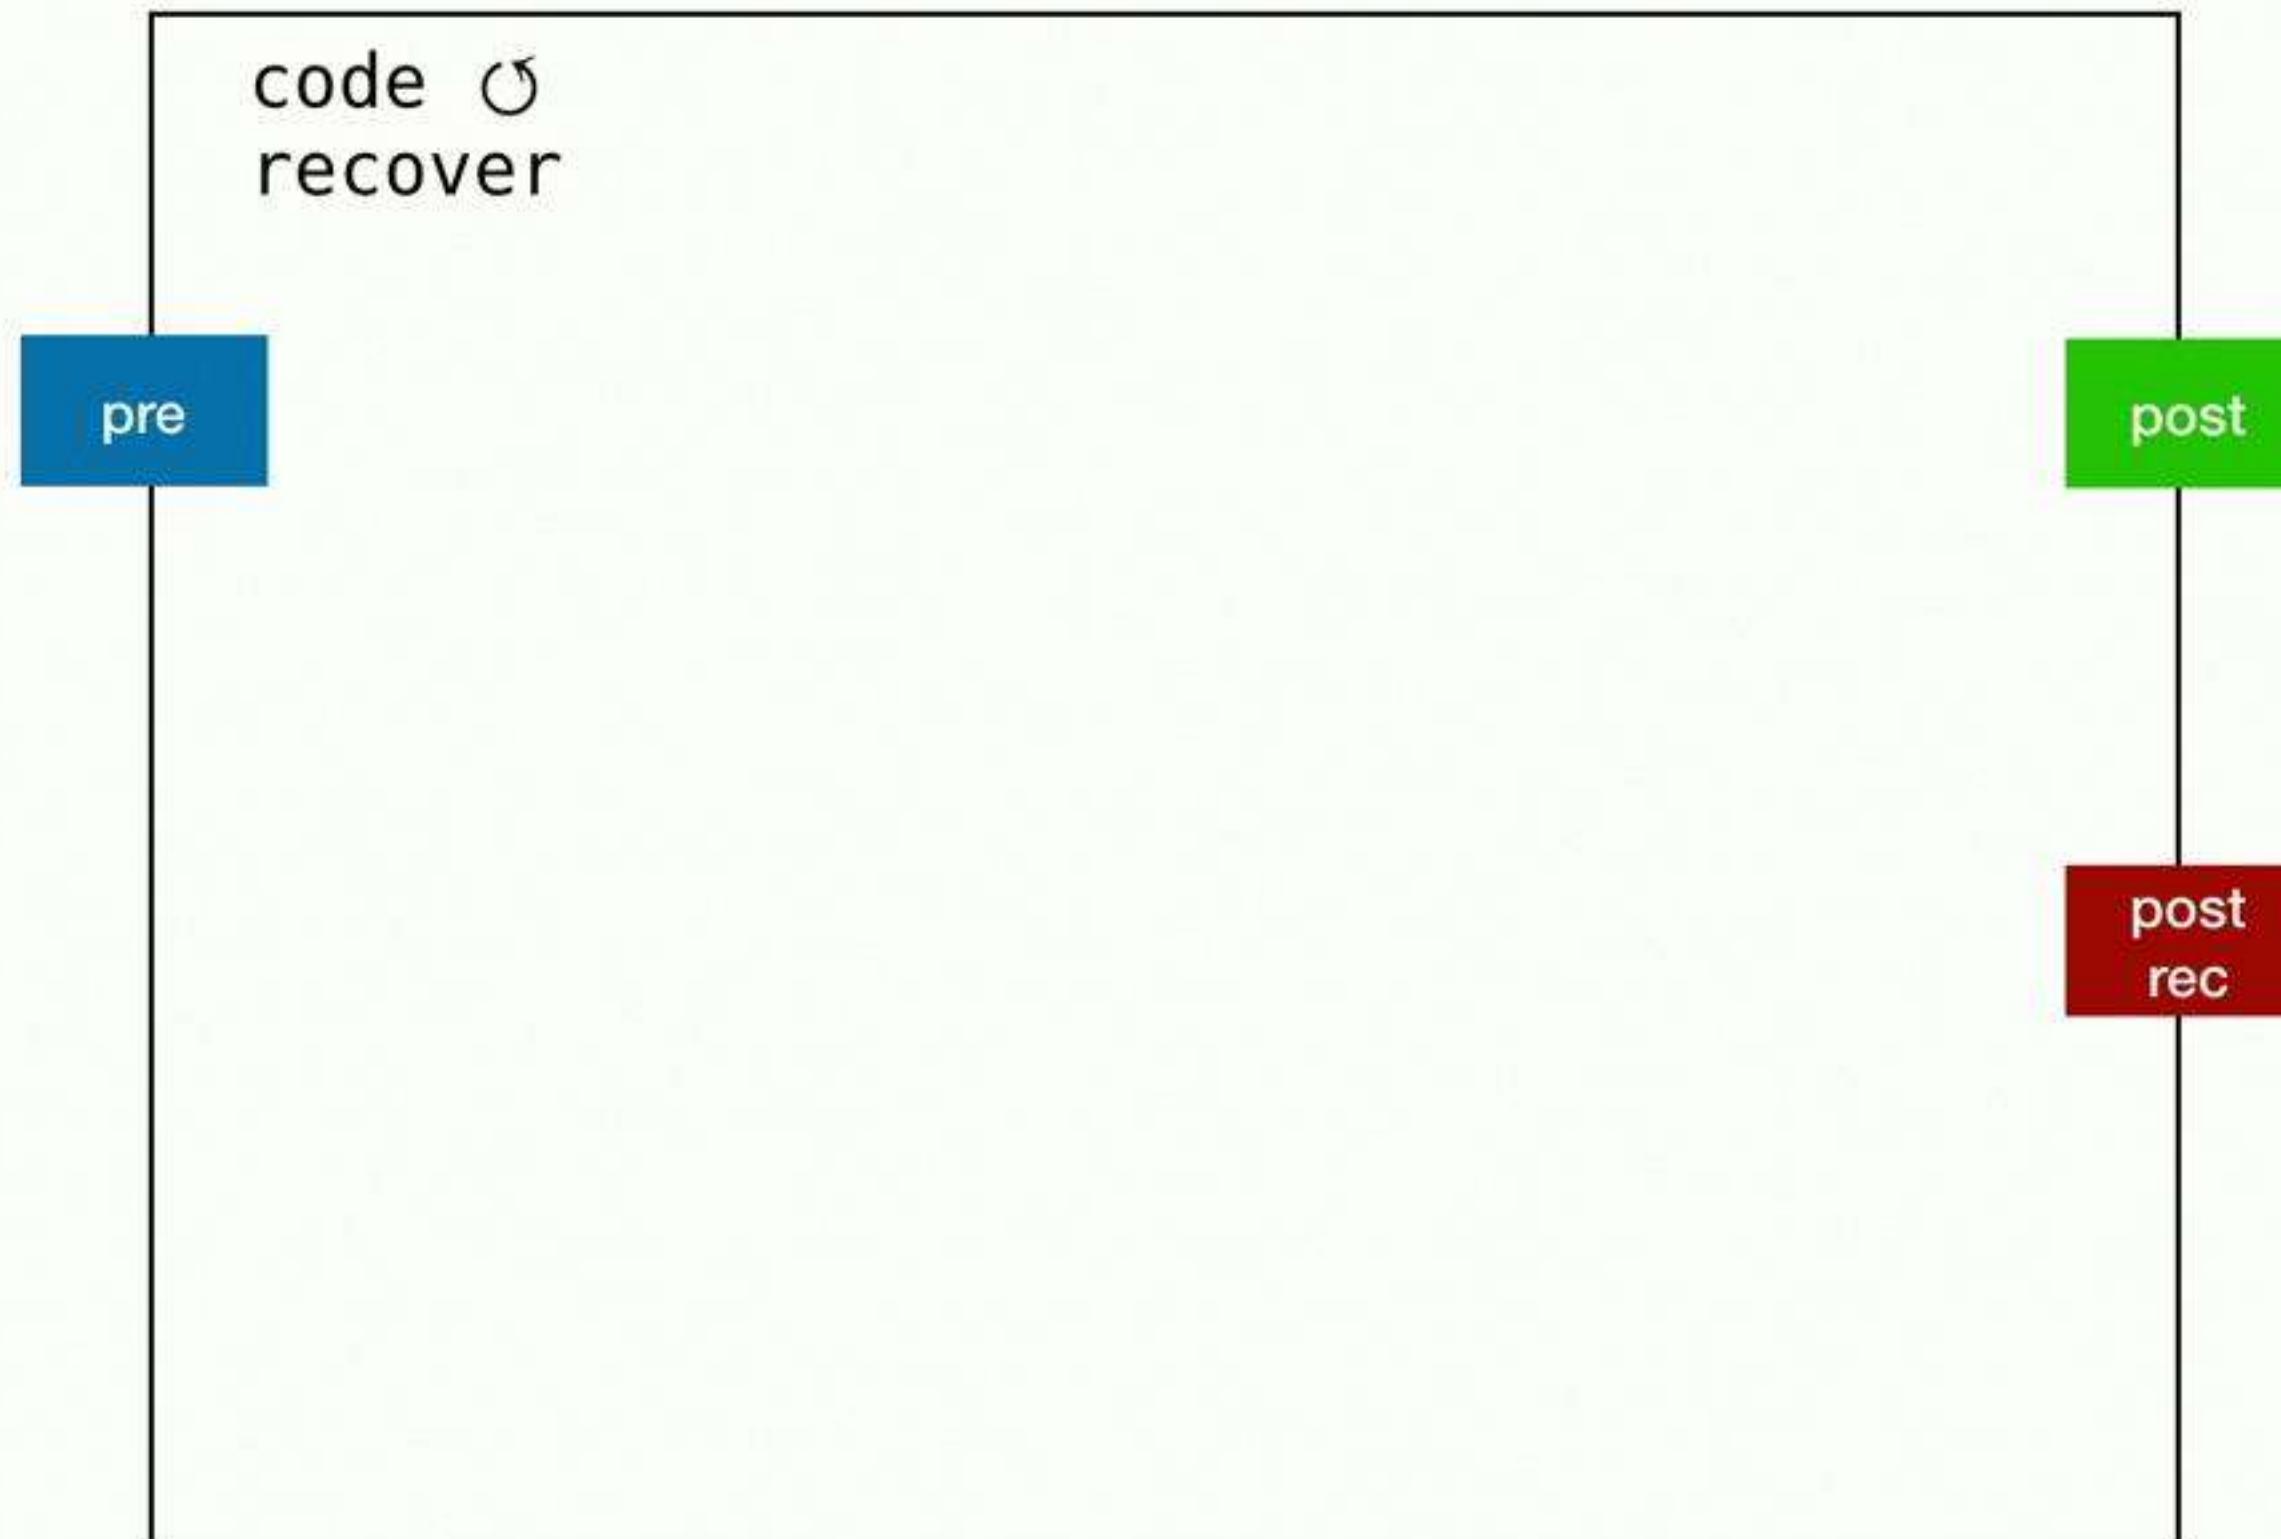
“recovery specification”  $\{P\}$  code  $\circlearrowleft$  recover  $\{Q\}$   $\{Q_r\}$   
precondition postcondition recovery postcondition



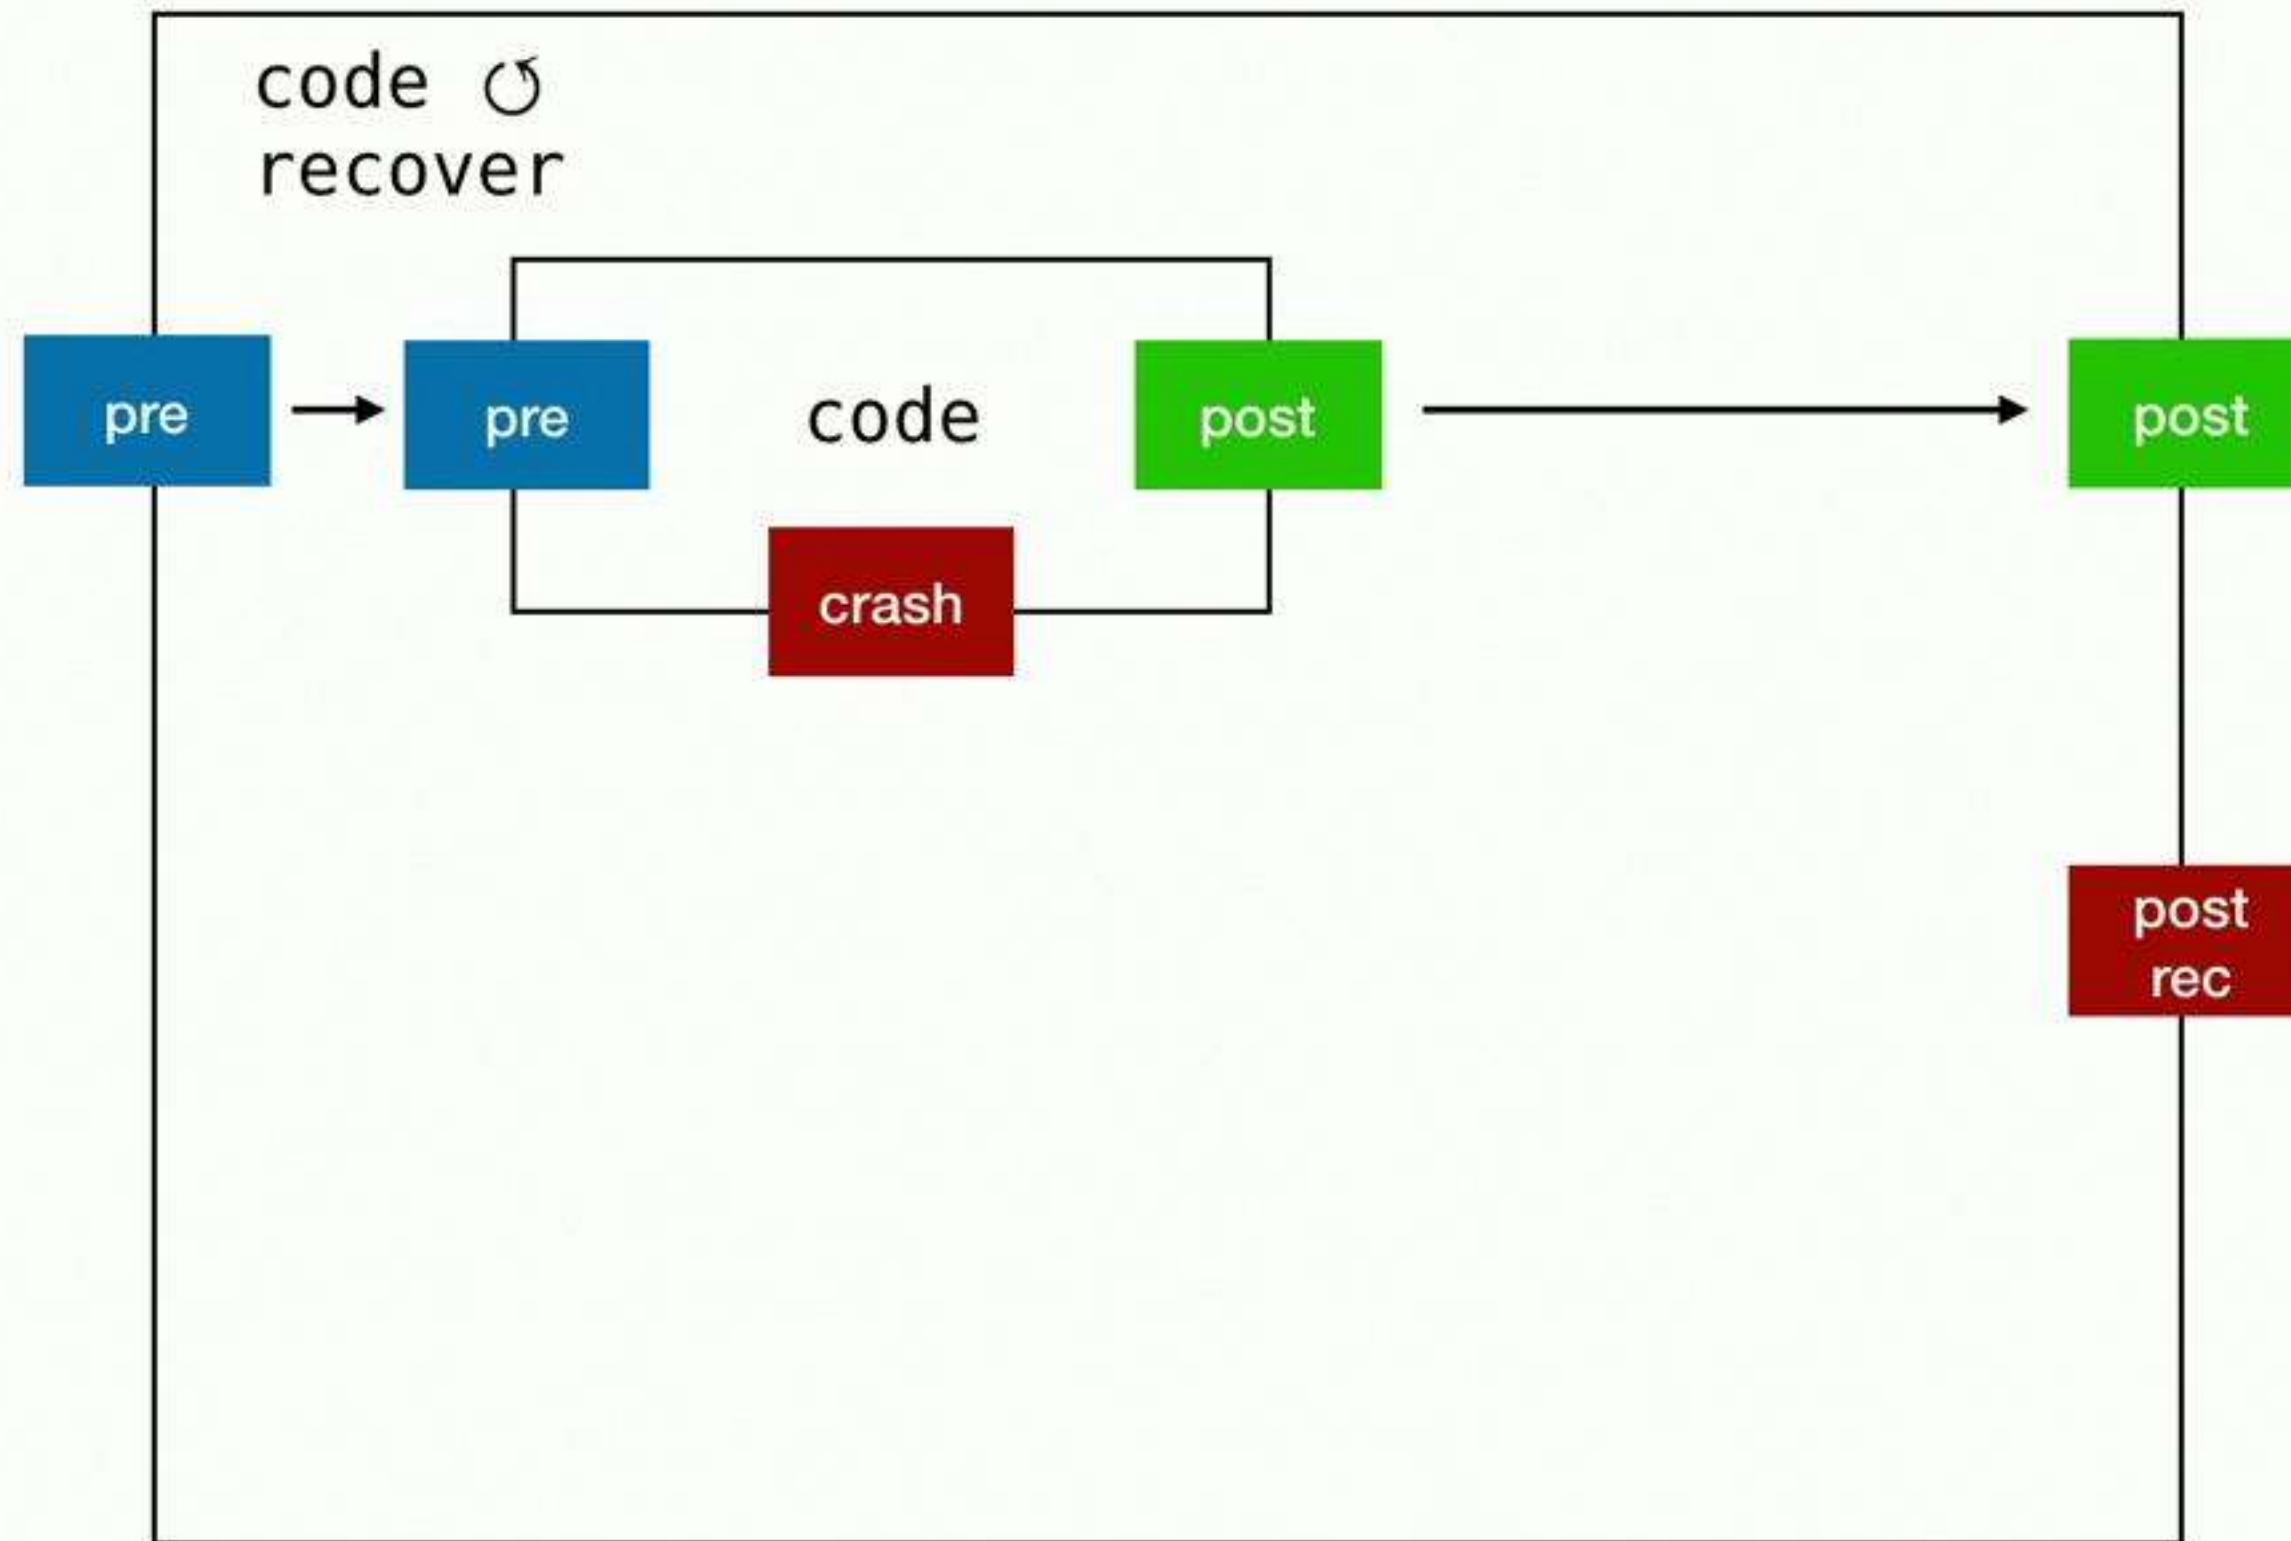
{pre} code ⚡ recover {post} {post rec}



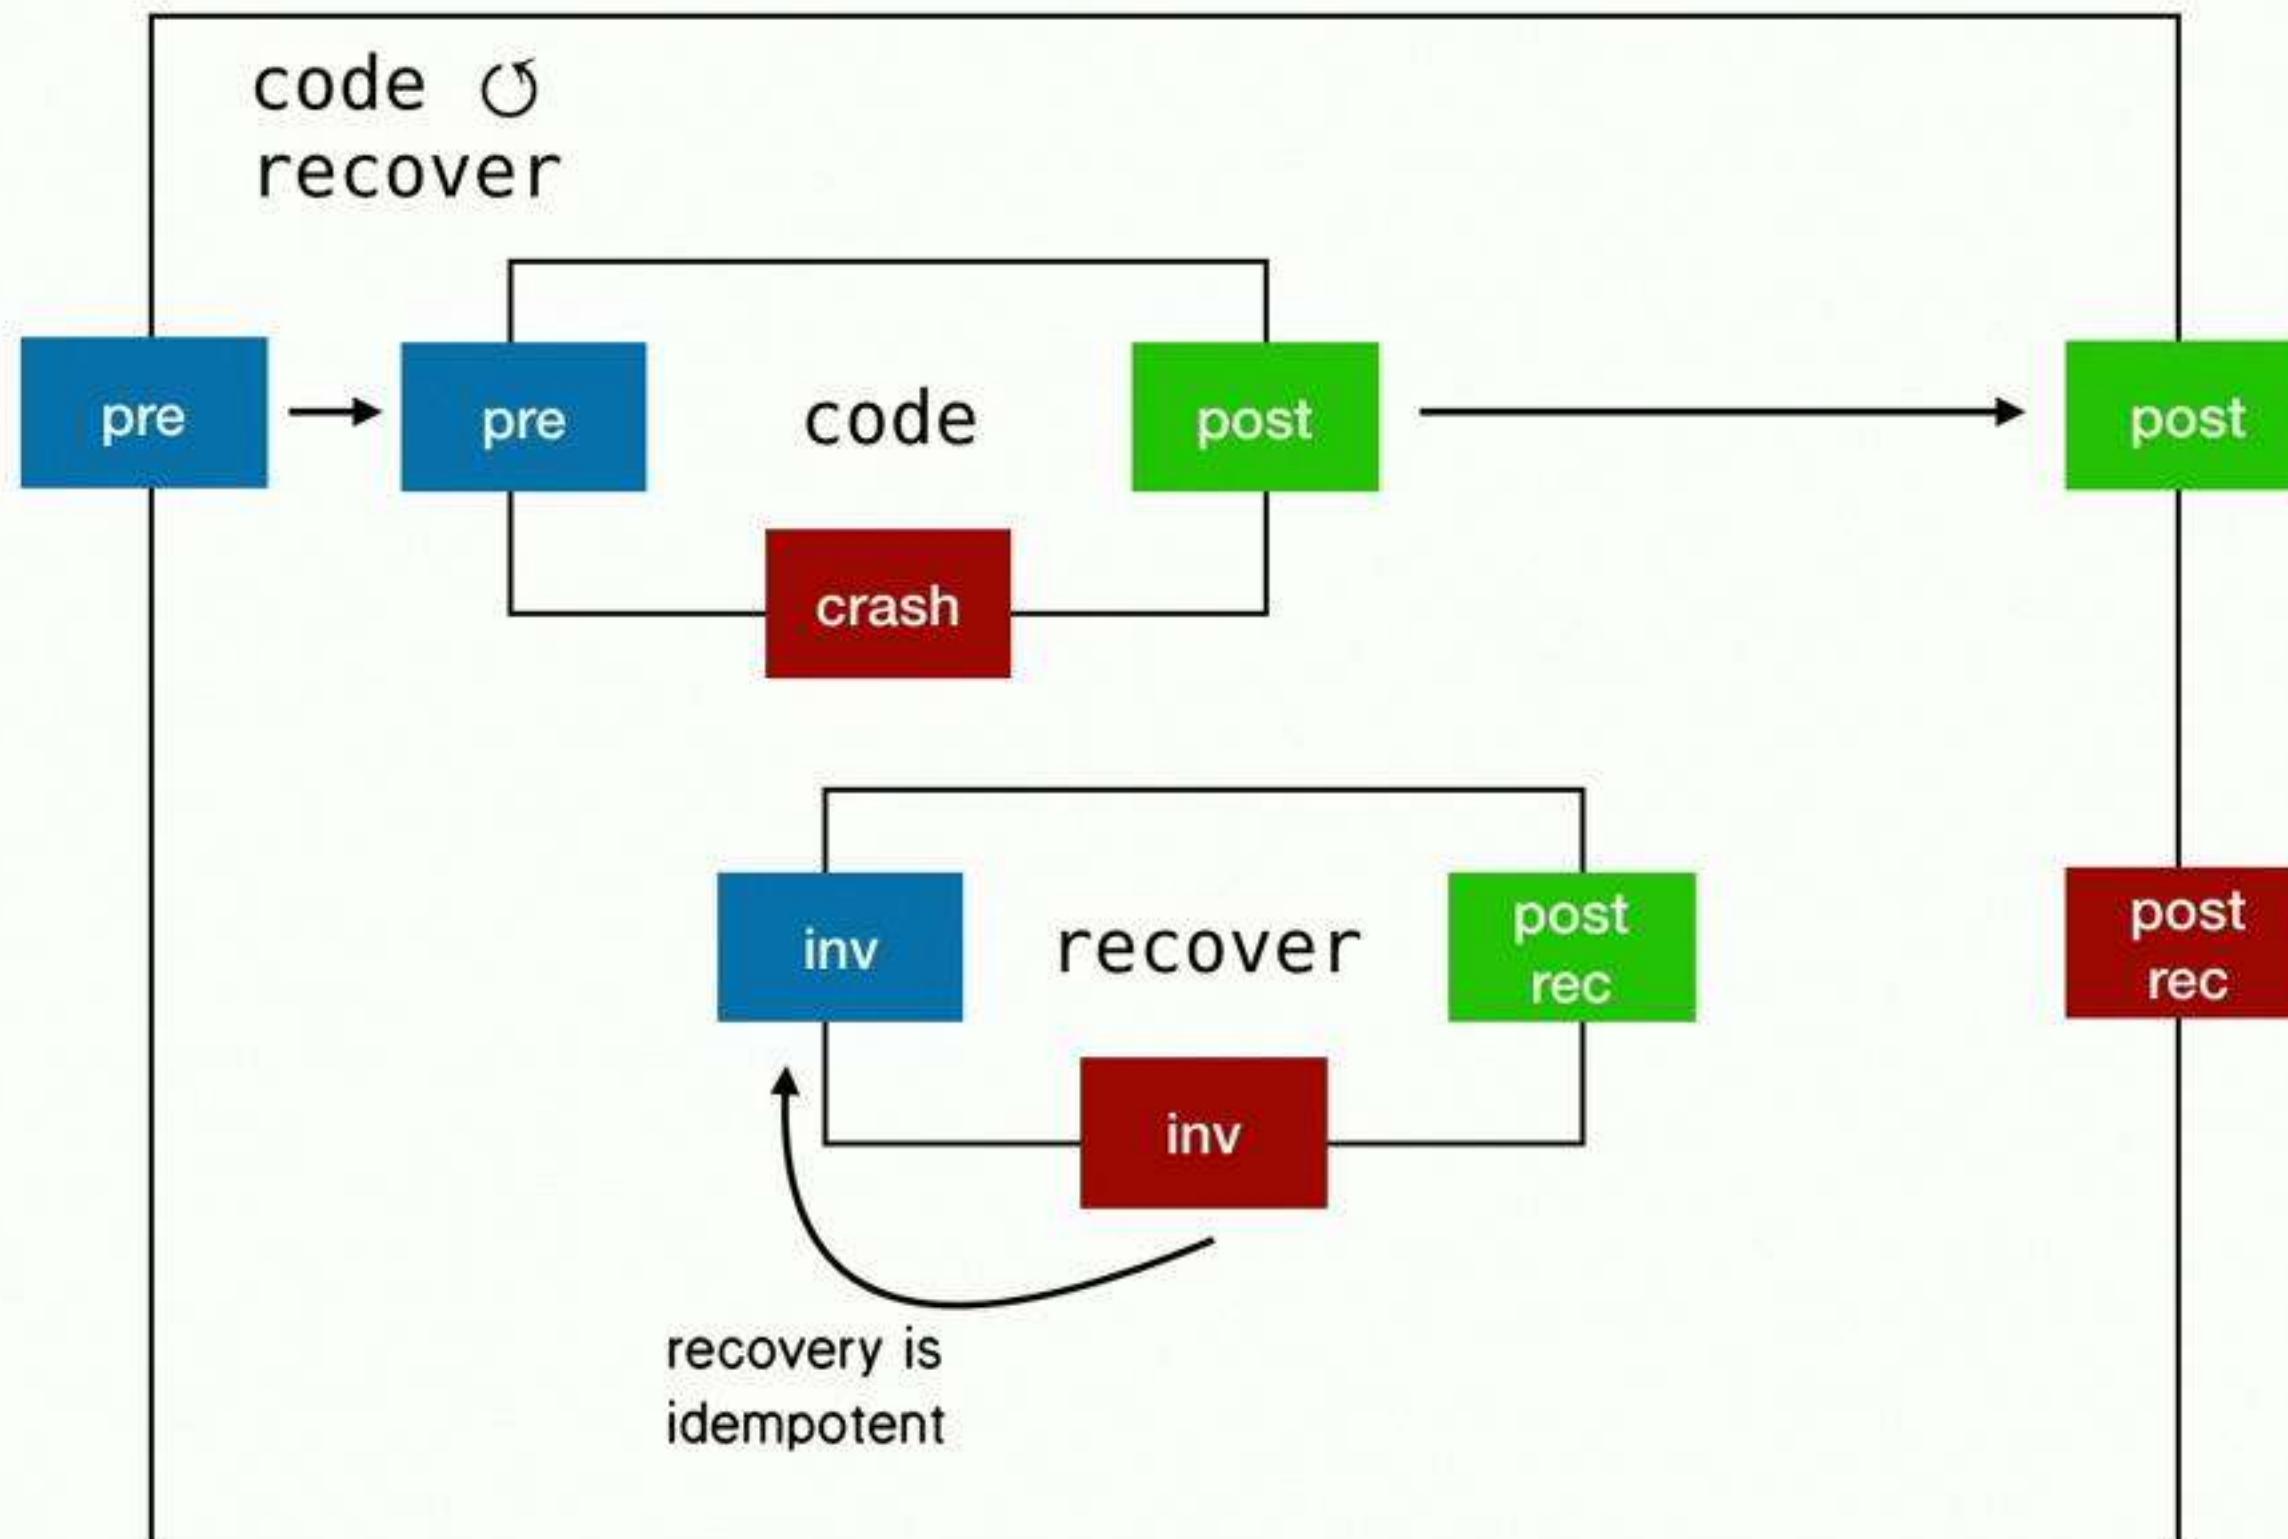
theorem in CHL to prove recovery specs from crash specs



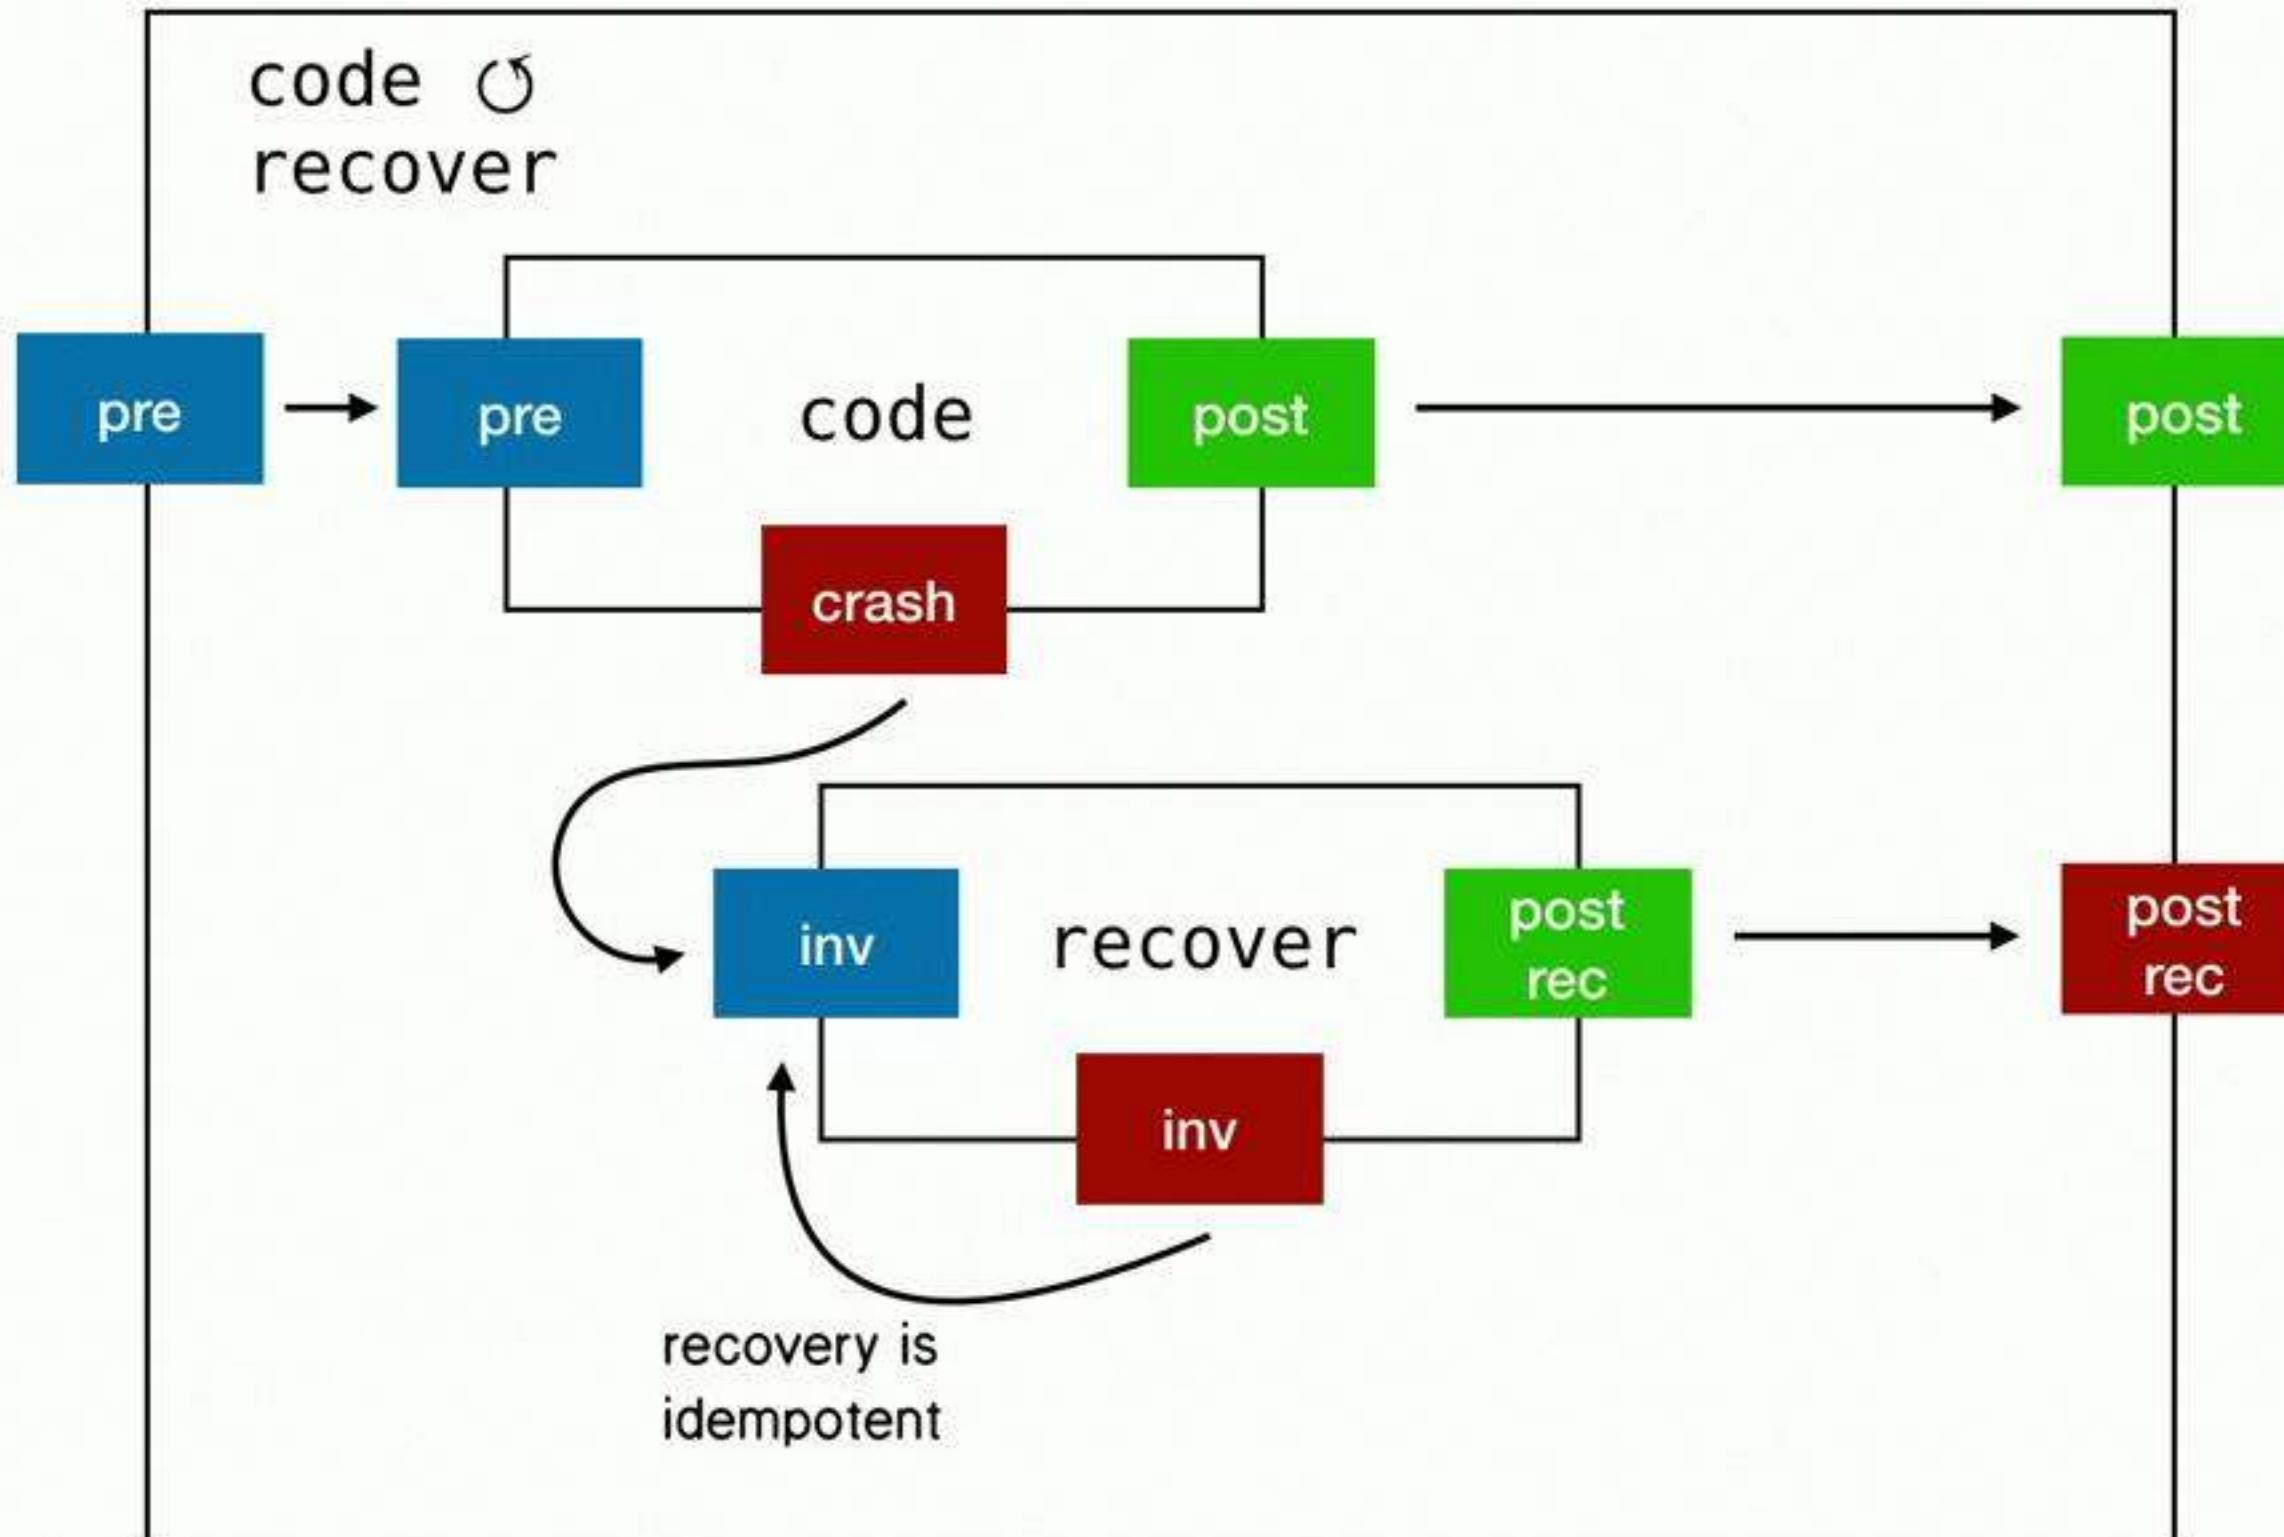
theorem in CHL to prove recovery specs from crash specs



theorem in CHL to prove recovery specs from crash specs



theorem in CHL to prove recovery specs from crash specs



# Argosy connects CHL to recovery refinement

Come up with *abstraction relation*

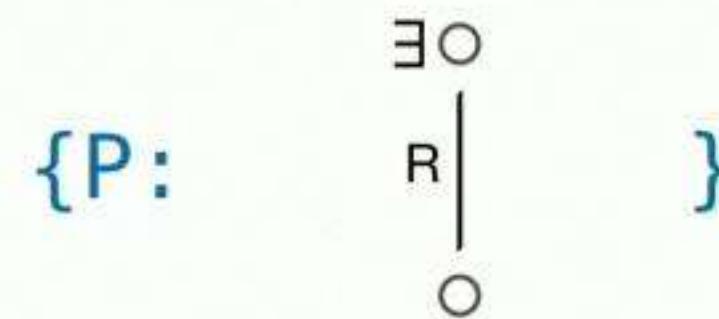
Prove a *refinement specification* for every operation

Gives recovery refinement for implementation

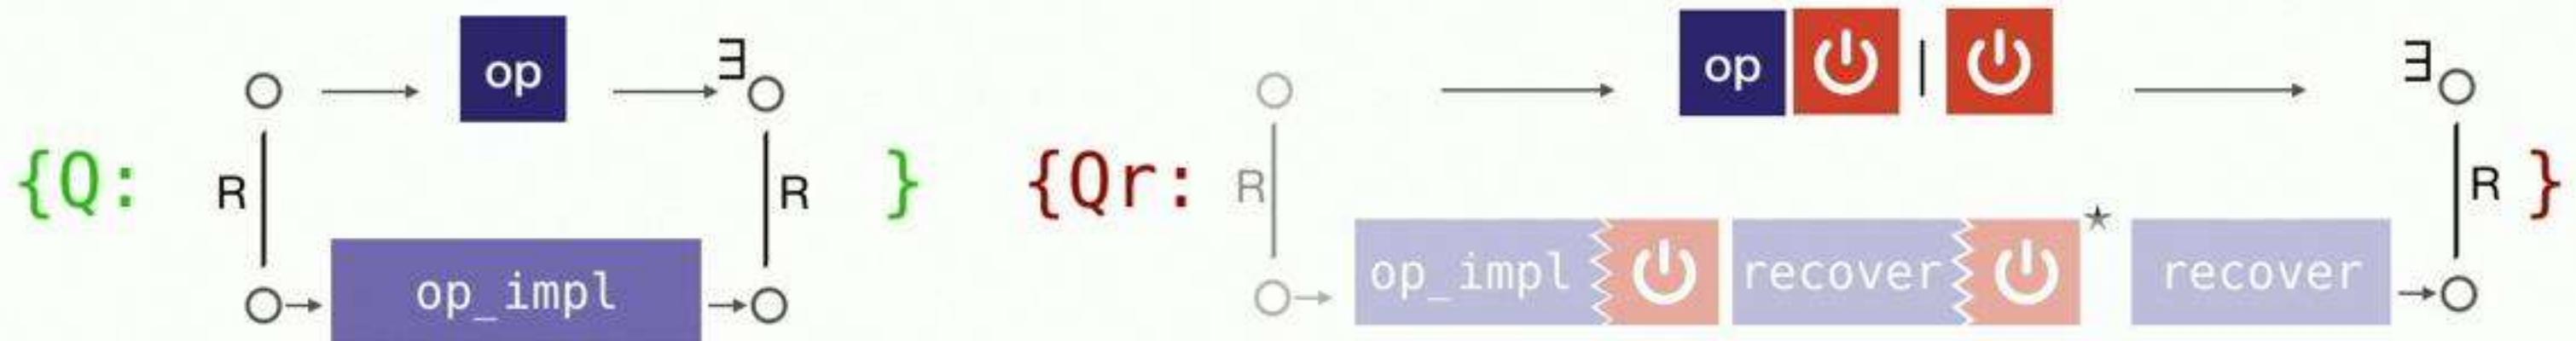
# Recovery refinement as a CHL spec

$\{P\} \text{ op\_impl } \circ \text{ recover } \{Q\} \{Qr\}$

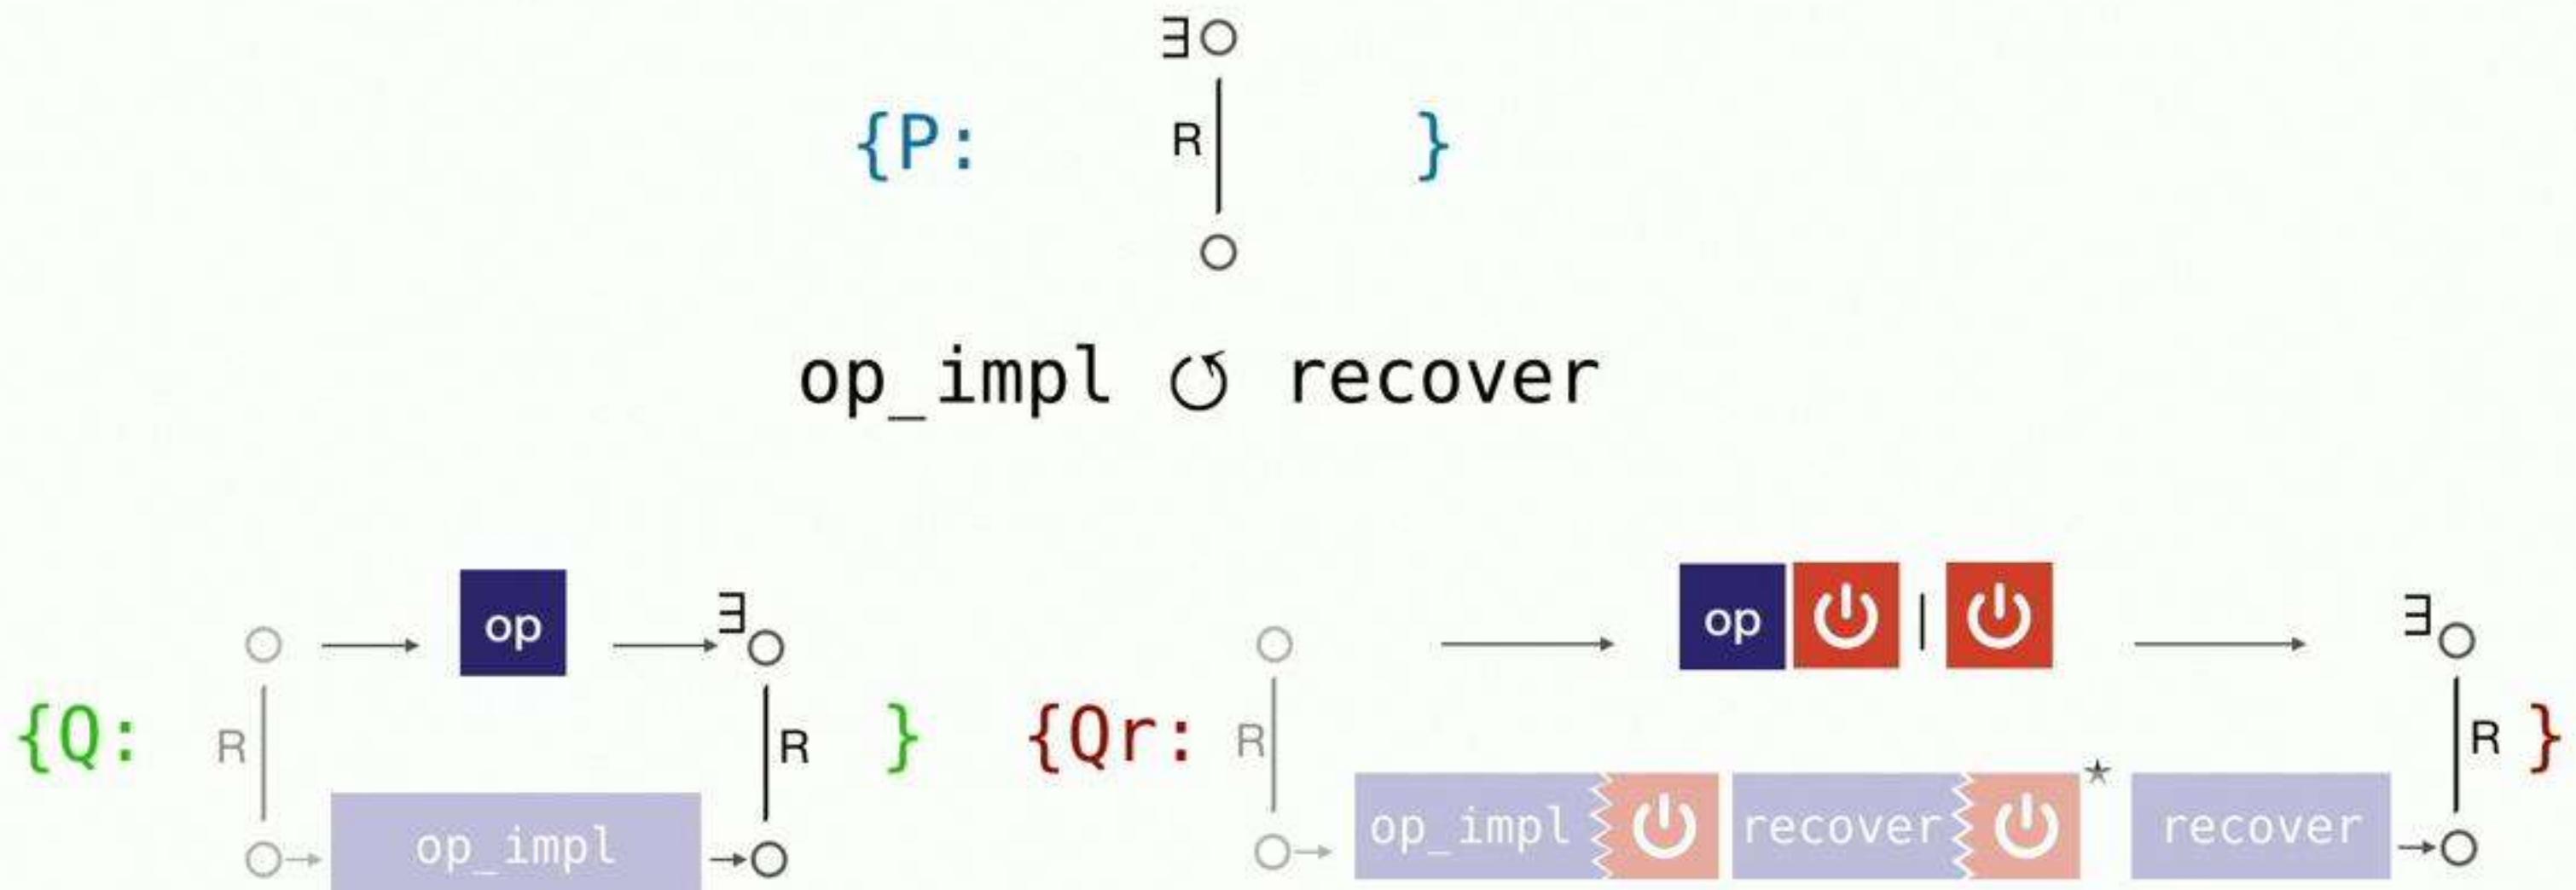
# Recovery refinement as a CHL spec

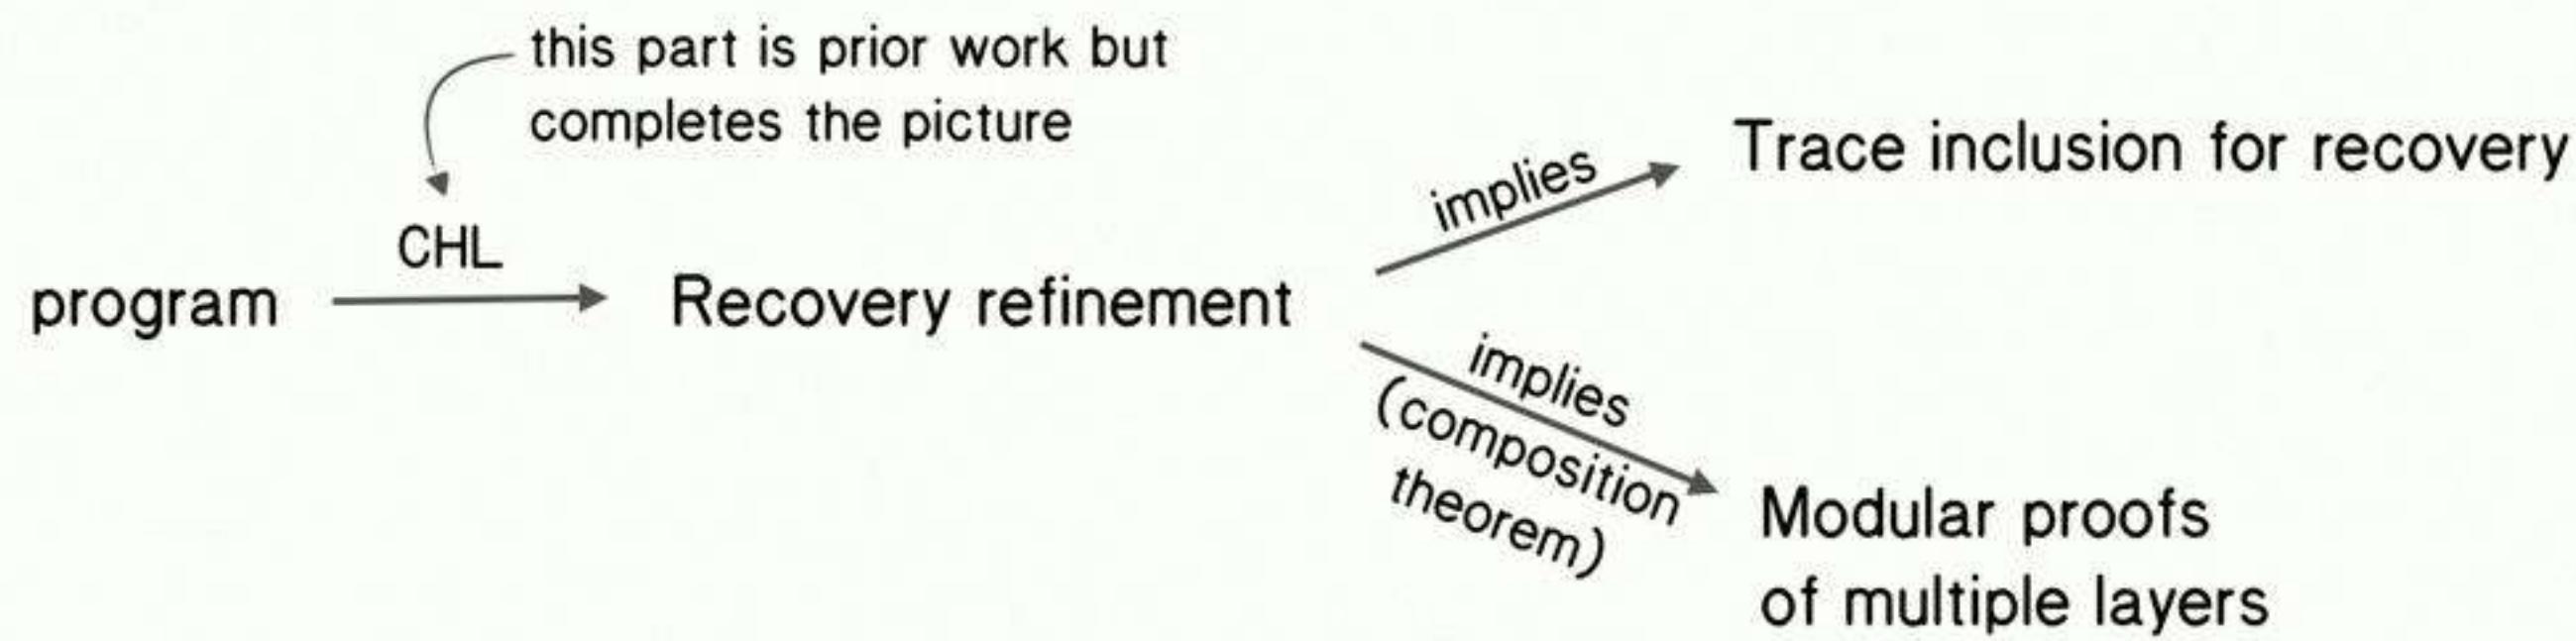


$op\_impl \circ recover$



# Recovery refinement as a CHL spec





# Argosy is implemented and verified in Coq



3,200 lines for framework

4,000 lines for verified example (logging + replication)

Example extracts to Haskell and runs

[github.com/mit-pdos/argosy](https://github.com/mit-pdos/argosy)

# Future work

# Concurrency

Extending Concurrent Separation Logic [originally 2007]

Implemented using Iris [originally POPL 2015]

# Better story for running code

Currently extract to Haskell

Performance problems (esp. for concurrency)

New plan: import Go into Coq

# Usability for students

Argosy spun off from course infrastructure

Now want to backport improvements

# Argosy: modular proofs of layered storage systems



Kleene algebra

( rep  | rep  log  )<sup>\*</sup>

# Argosy: modular proofs of layered storage systems



## Kleene algebra

( rep  | rep log  )\*

## recovery refinemen



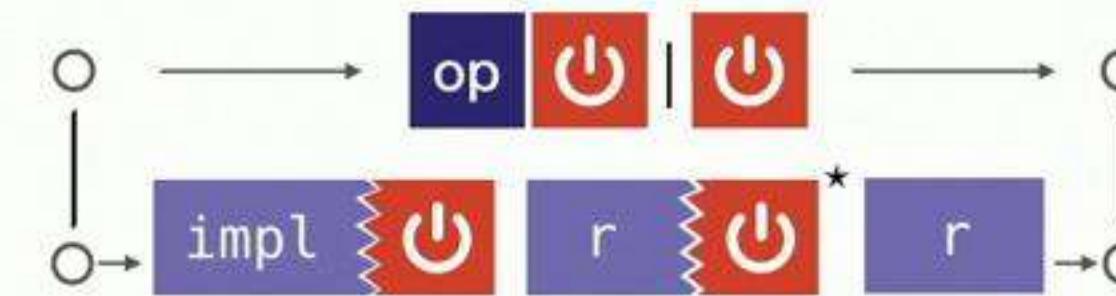
# Argosy: modular proofs of layered storage systems



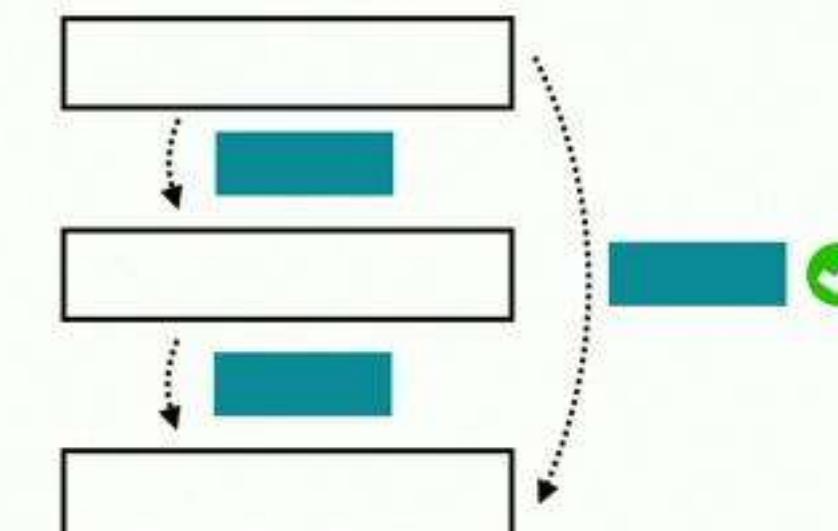
Kleene algebra

$$(\text{rep} \xrightarrow{\quad} \text{op} \mid \text{rep} \xrightarrow{\quad} \text{log} \xrightarrow{\quad} \text{op})^*$$

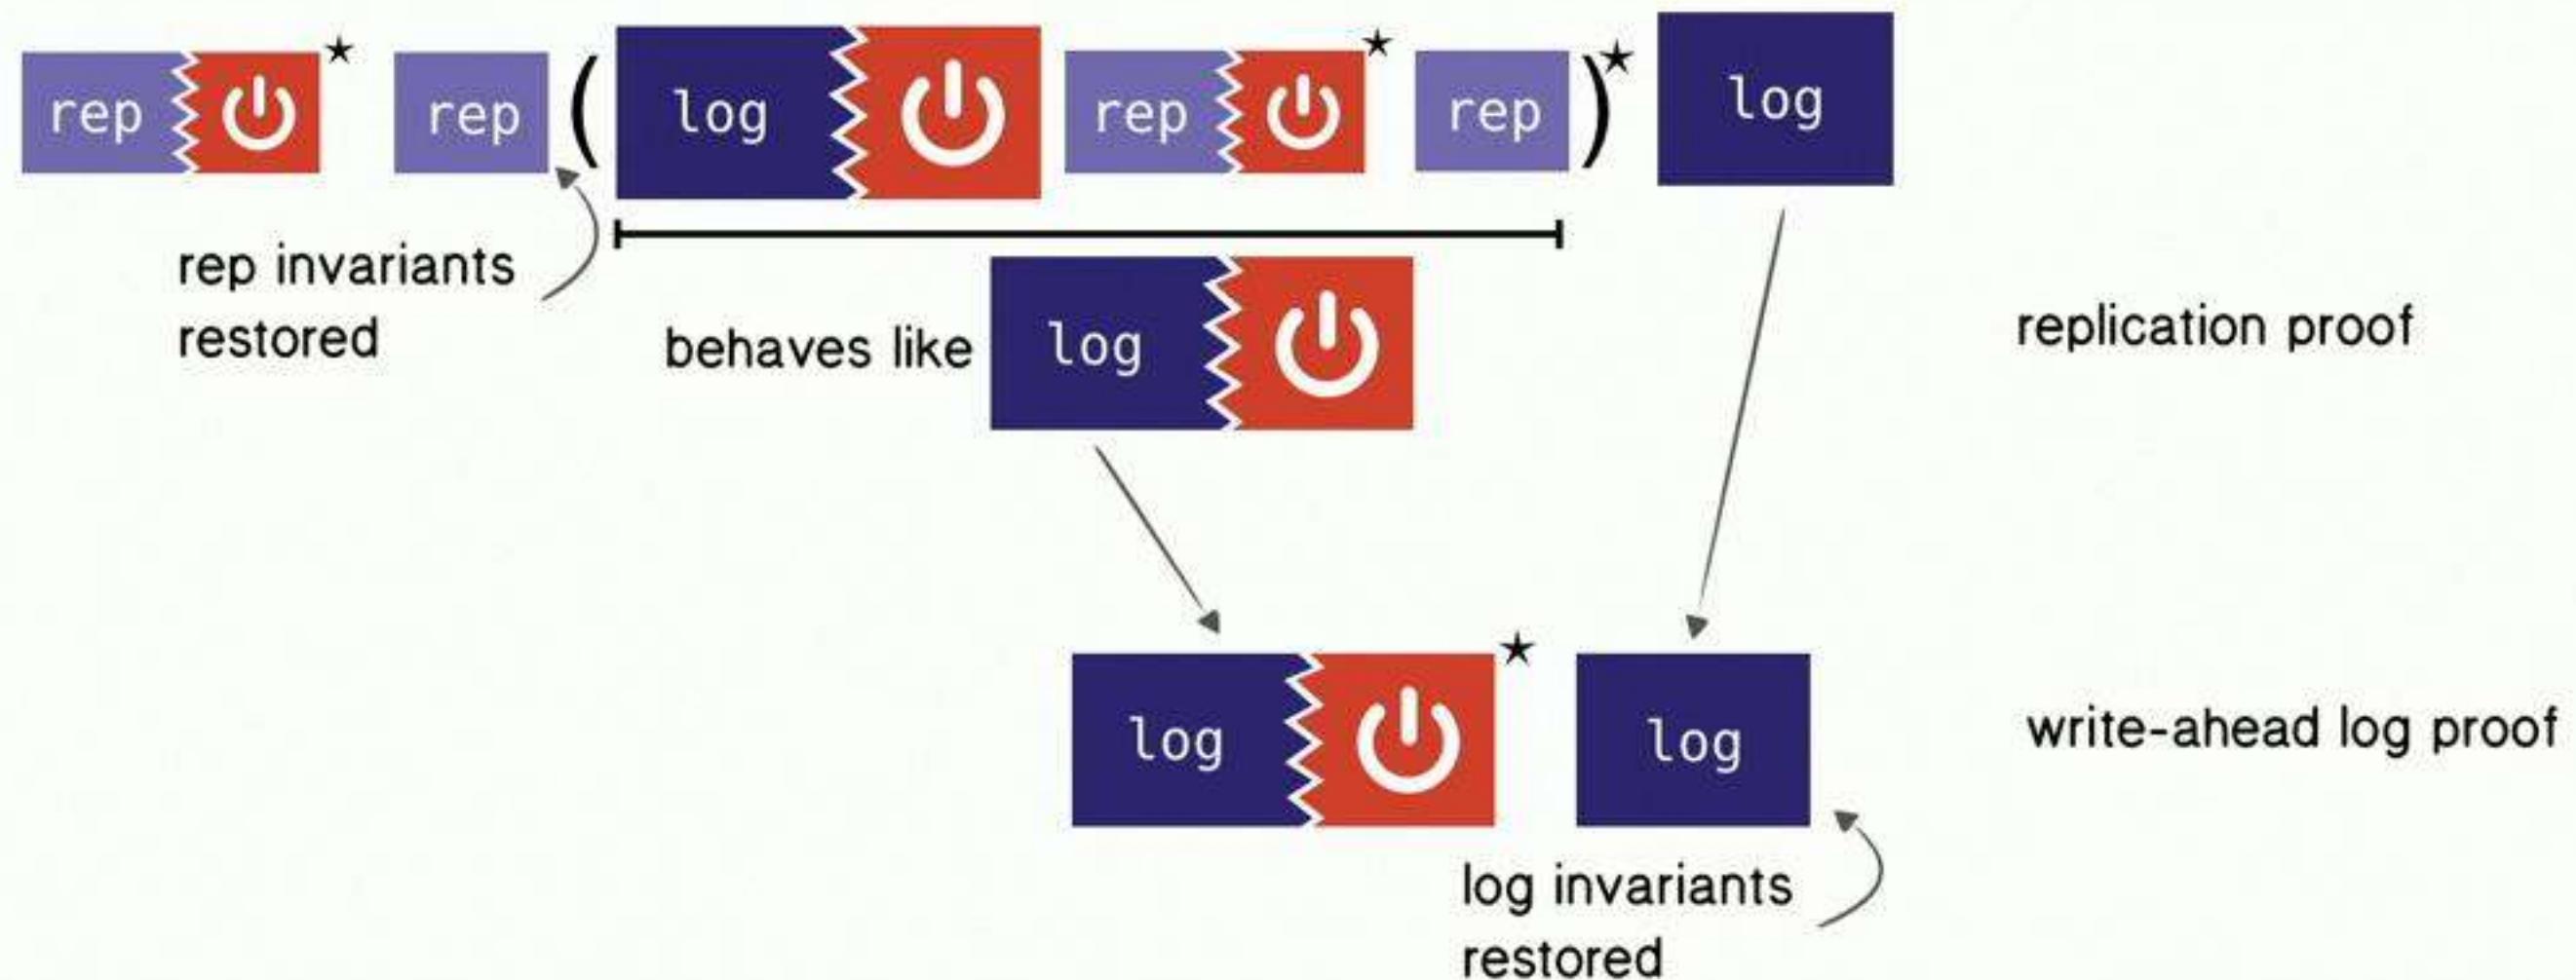
recovery refinement



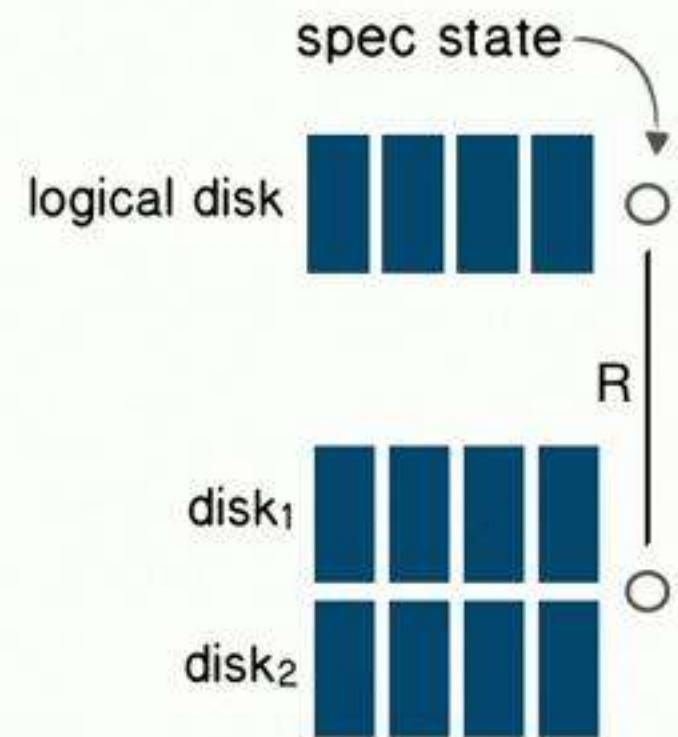
modular proofs



# After rewrite both proofs apply



# Proving correctness with an abstraction relation



1. developer provides abstraction relation  $R$

```
def atomic_save(data, path):
    write_all(data, tmp)
    rename(tmp, path)

# runs on crash
def recover():
    fs_recover()
    unlink(tmp)
```