

Argosy: Verifying layered storage systems with recovery refinement

Tej Chajed, Joseph Tassarotti, Frans Kaashoek, Nickolai Zeldovich

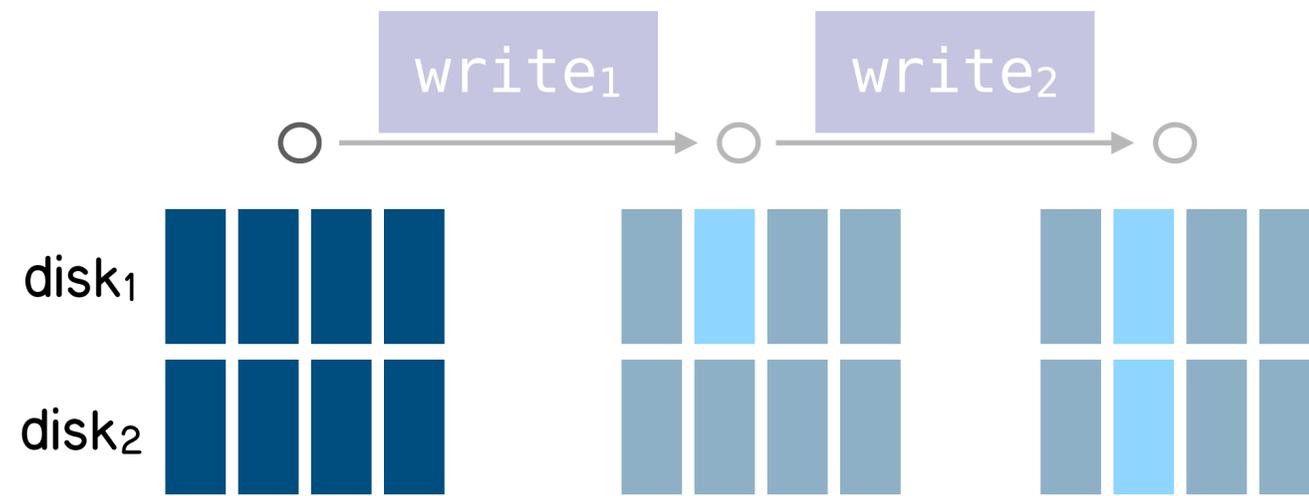
MIT

logical disk 

disk₁ 
disk₂ 

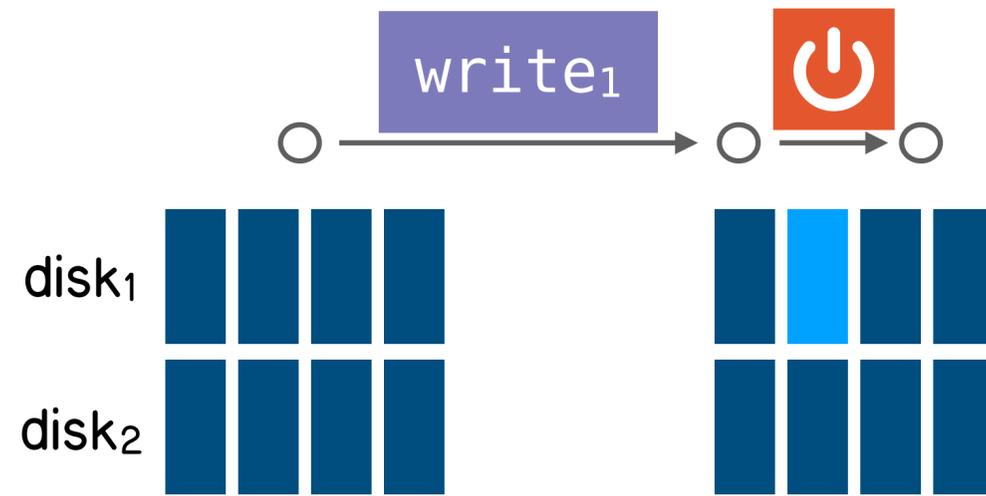
Bob writes a replication system

logical disk

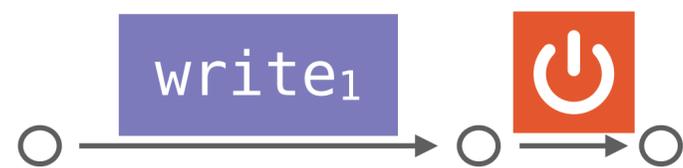
A horizontal row of four dark blue vertical bars representing a logical disk.

Bob writes a replication system

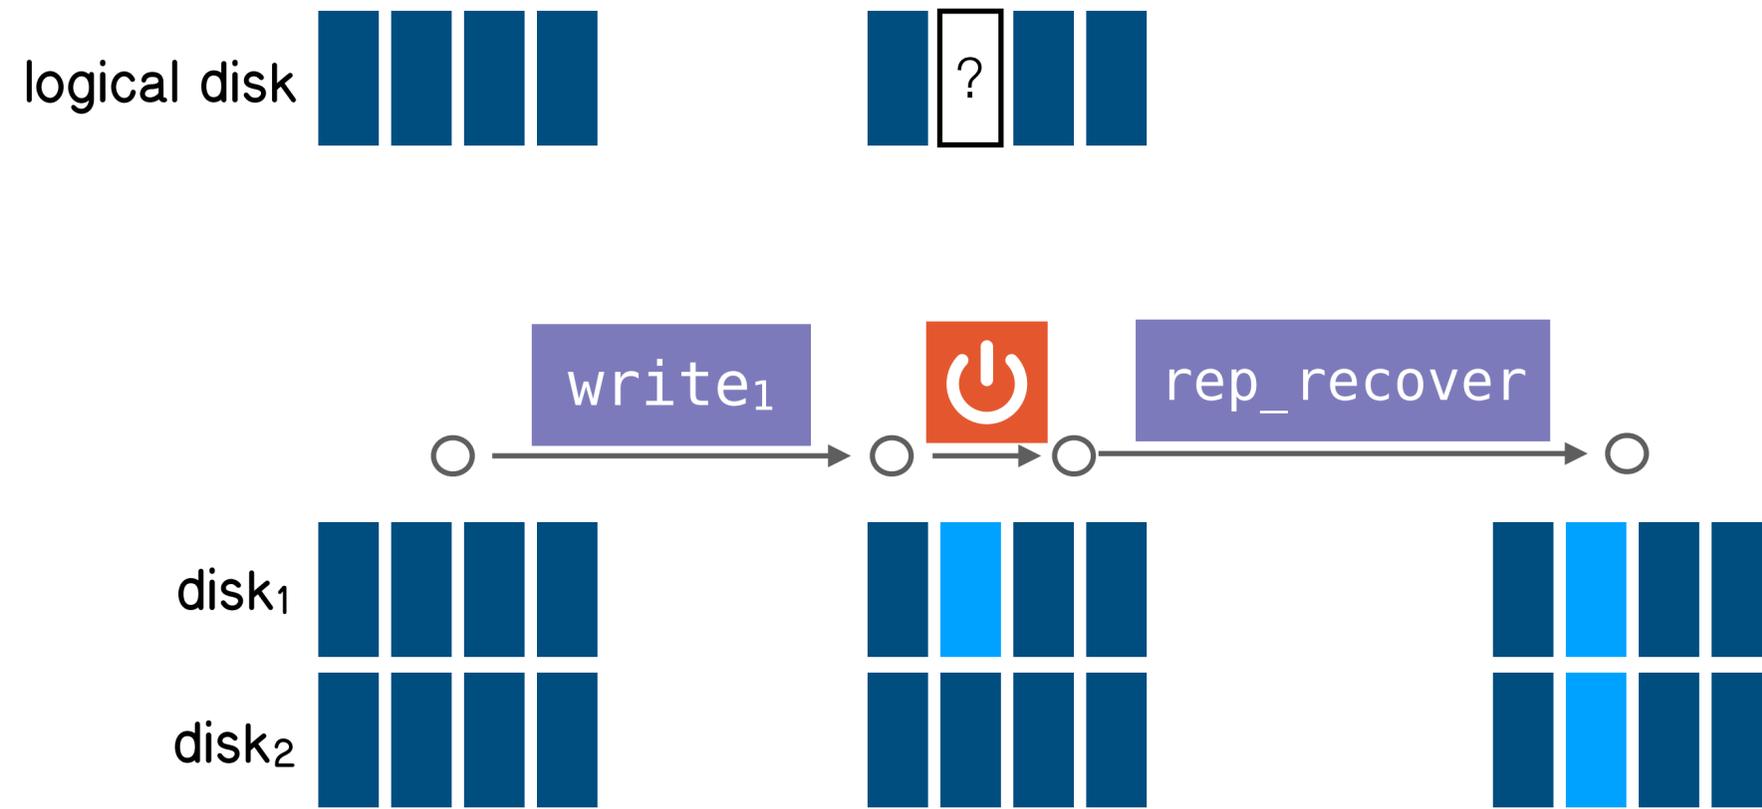
logical disk 



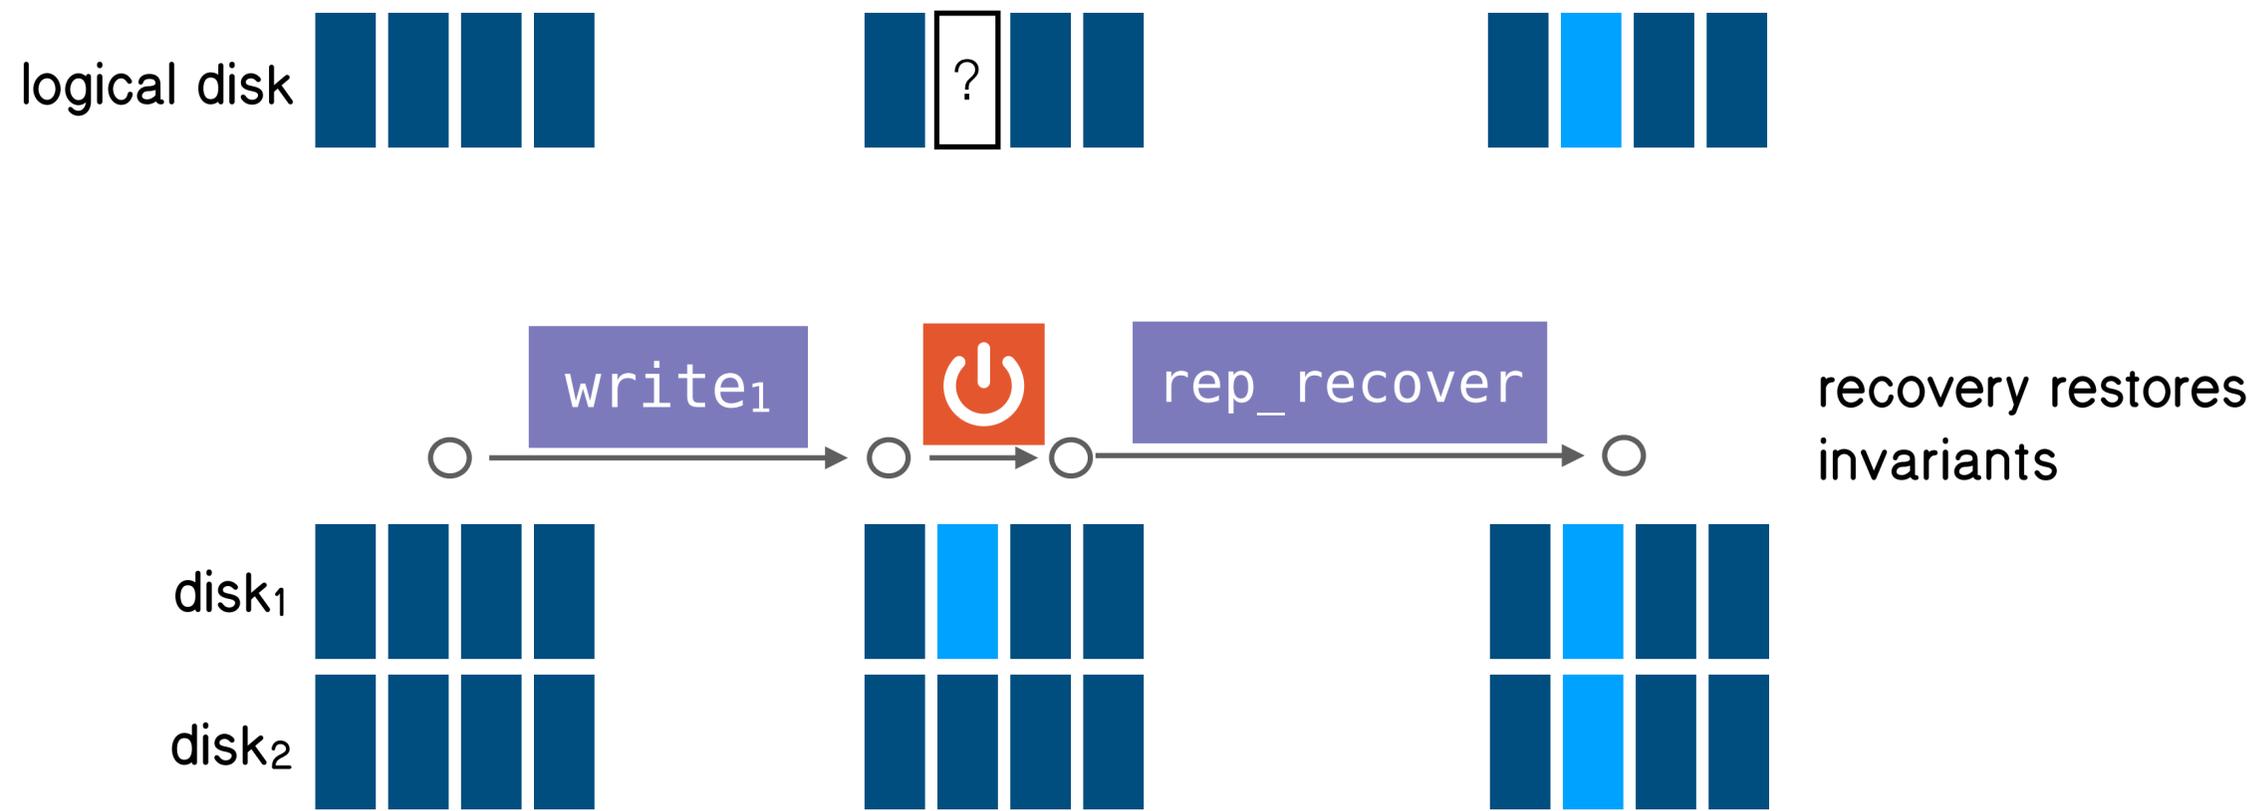
Bob writes a replication system



Bob writes a replication system

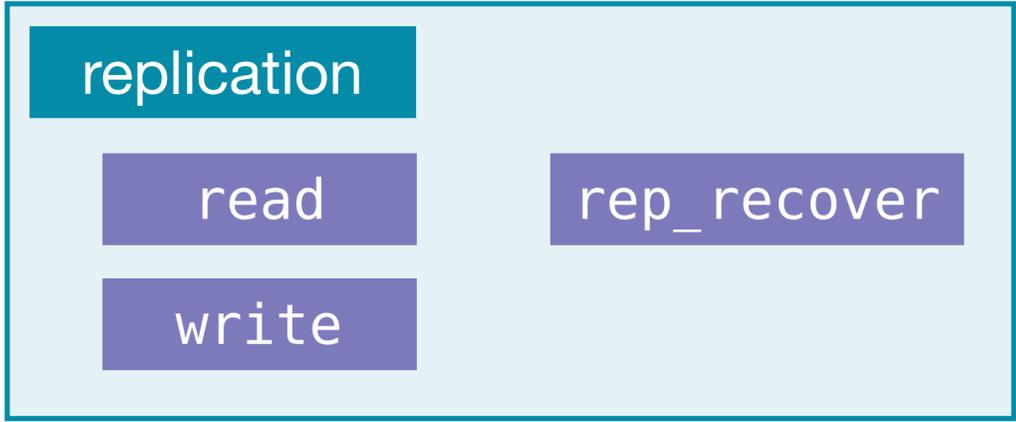


Bob writes a replication system and implements its recovery procedure



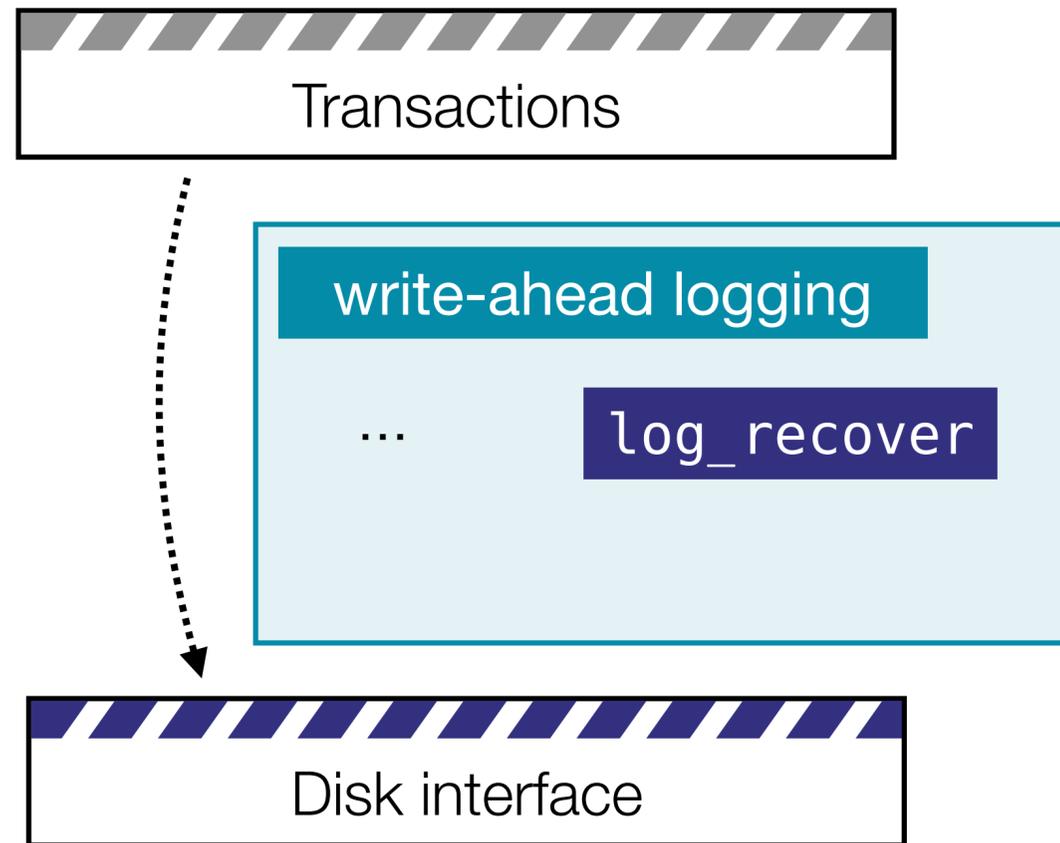
Bob writes a replication system and implements its recovery procedure

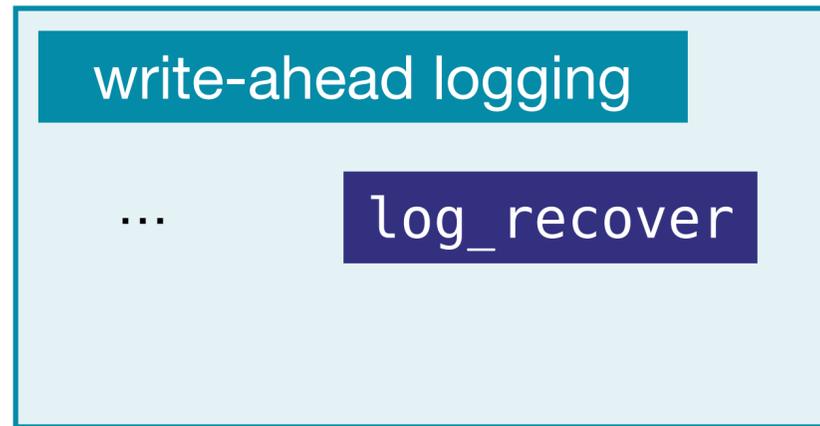
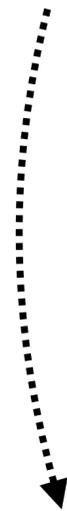
Bob is careful and writes a machine-checked proof of correctness



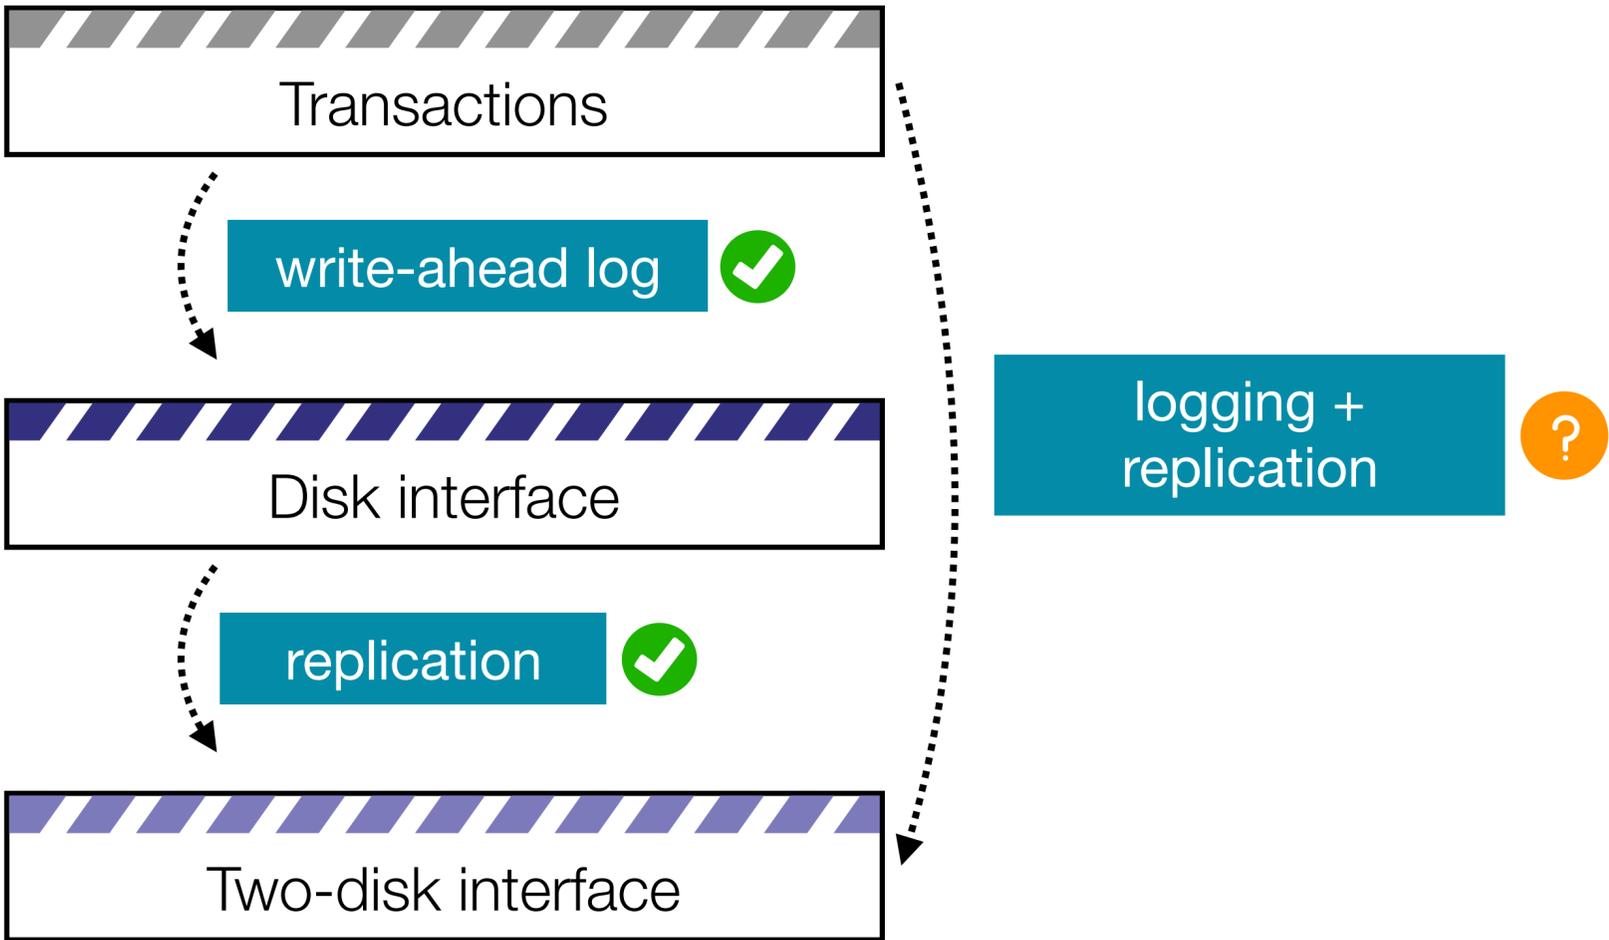
read and write are atomic if you run rep_recover after every crash

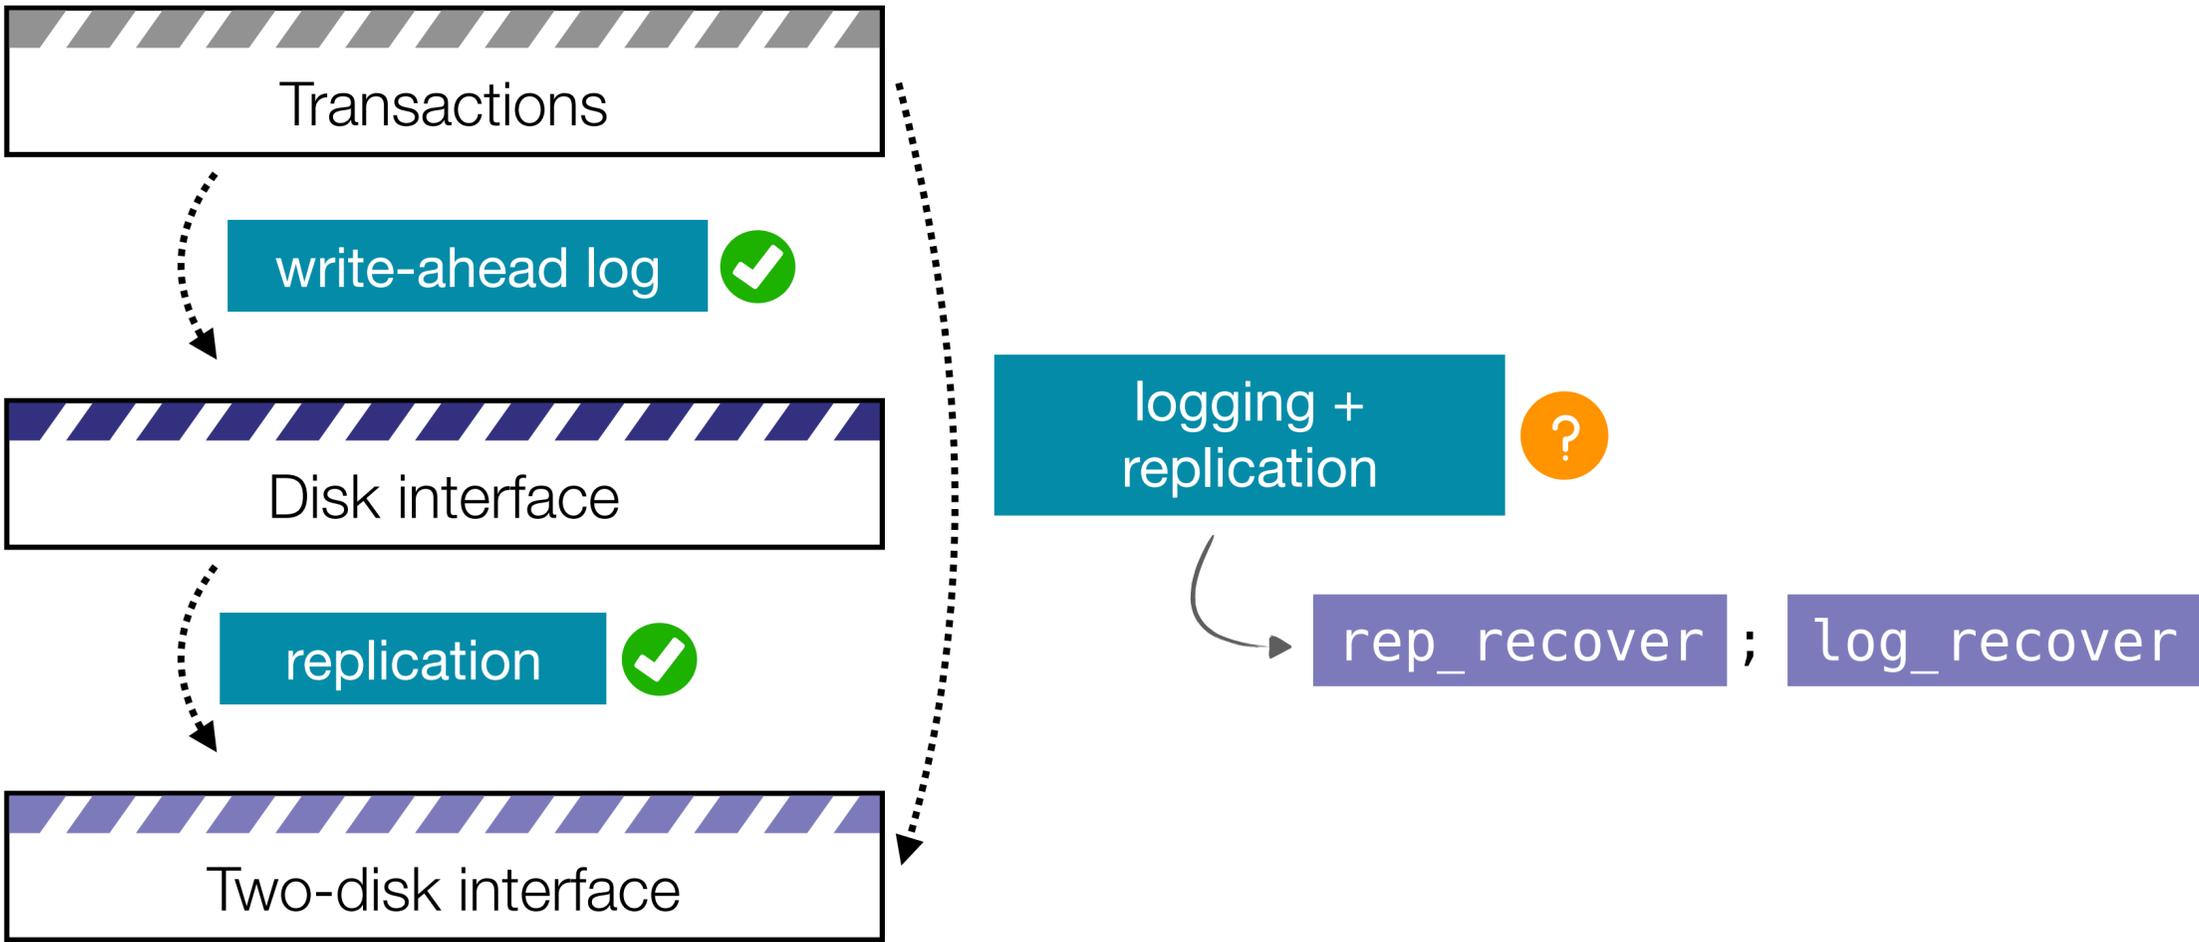






✔ ops are atomic if you run
log_recover after every crash





Challenge: crashes during composed recovery

`rep_recover` ✓ under crashes

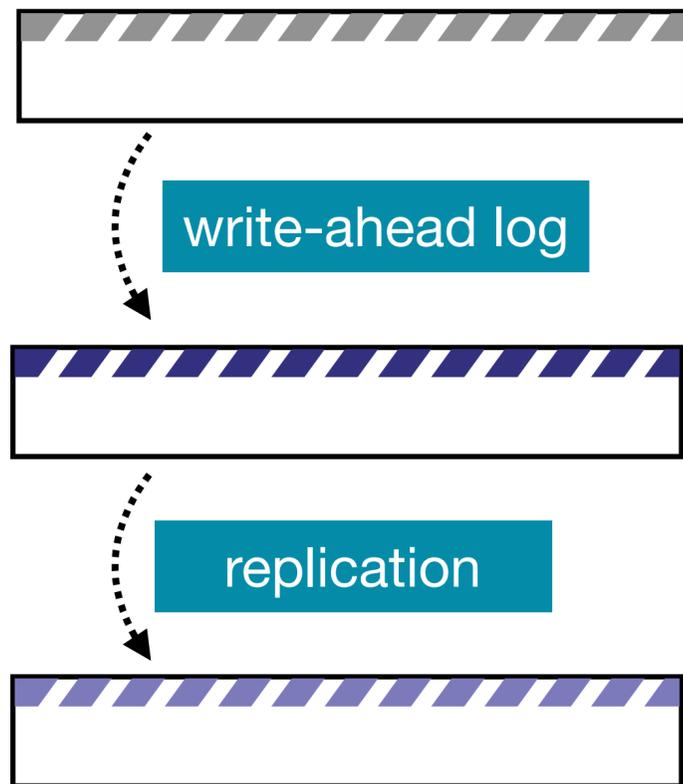
`log_recover` ✓ under crashes

`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

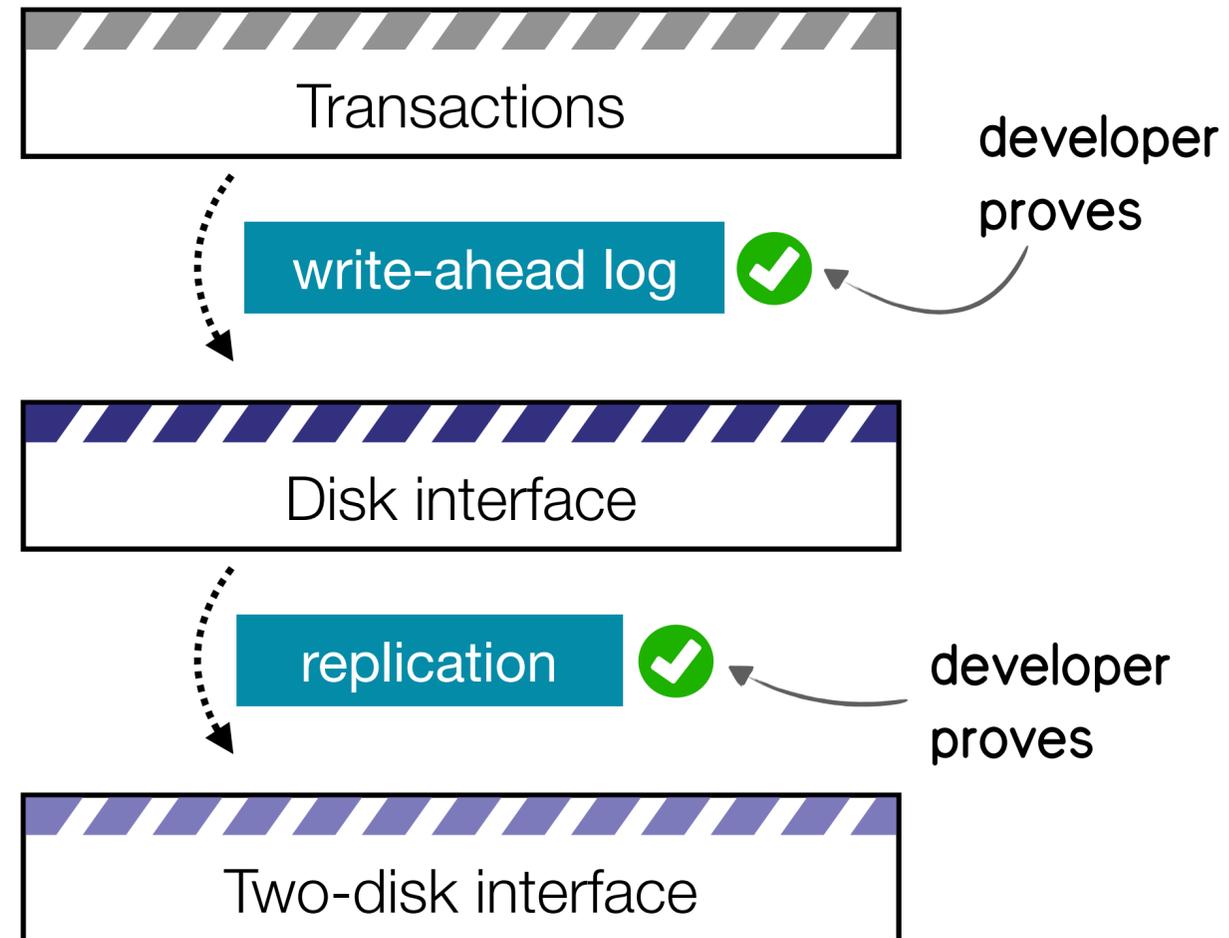
Yggdrasil [OSDI '16]

single recovery

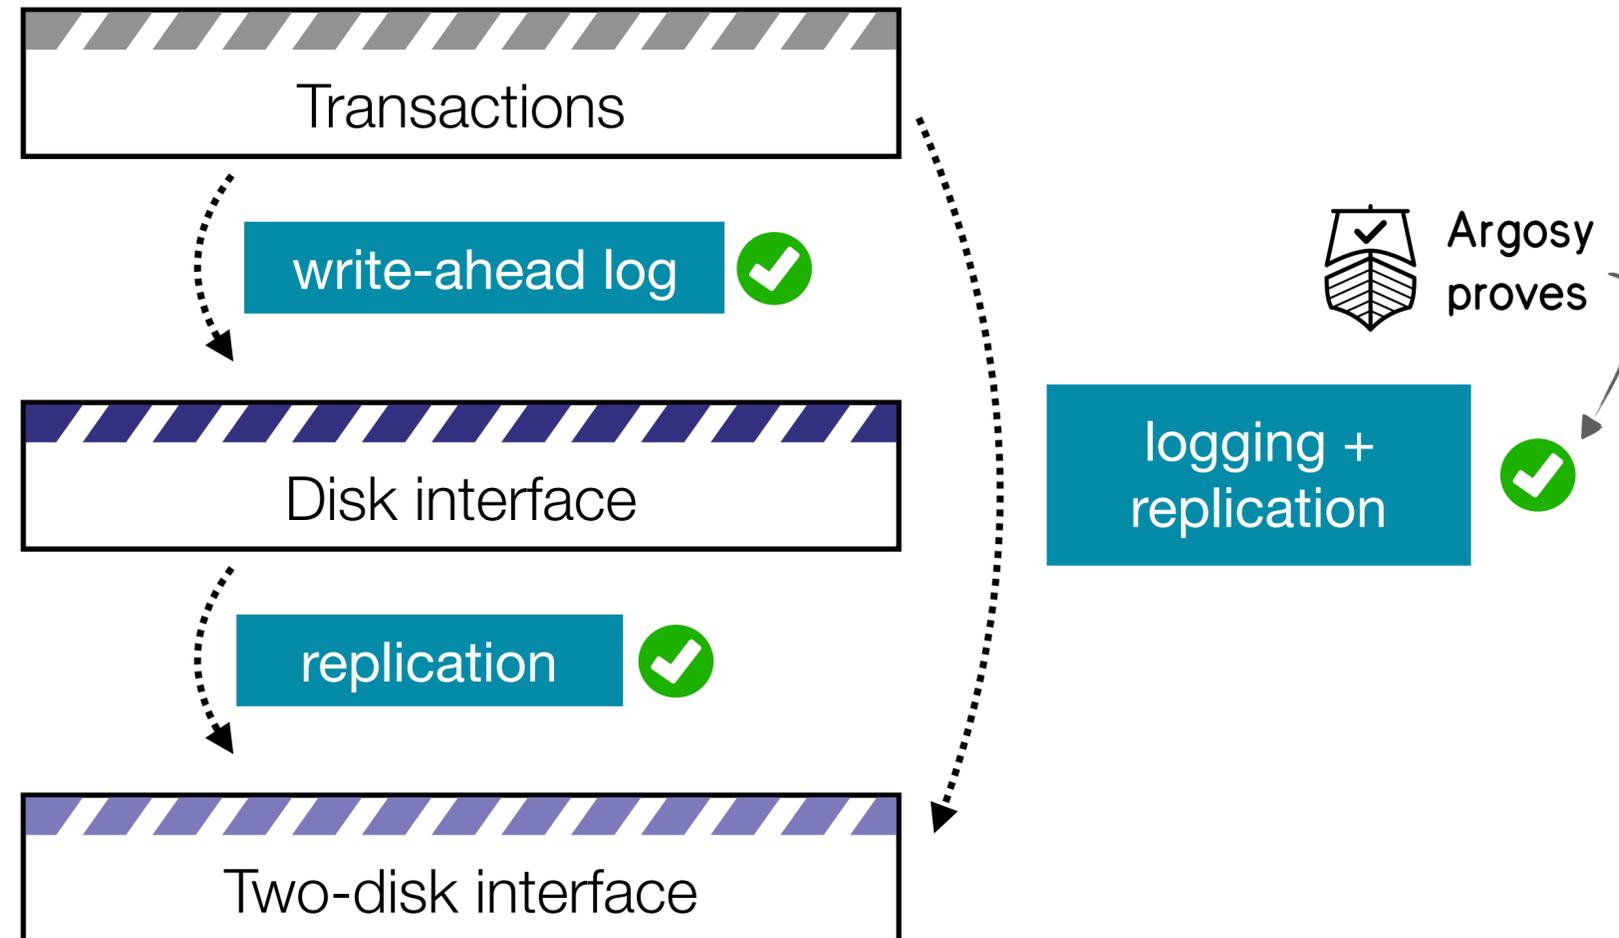
Flashix [SCP '16]

restricted recovery
procedures

Argosy supports modular recovery proofs



Argosy supports modular recovery proofs



Contributions

Recovery refinement for modular proofs

Contributions

Recovery refinement for modular proofs

see paper CHL for proving recovery refinement

see paper Verified example: logging + replication

Contributions

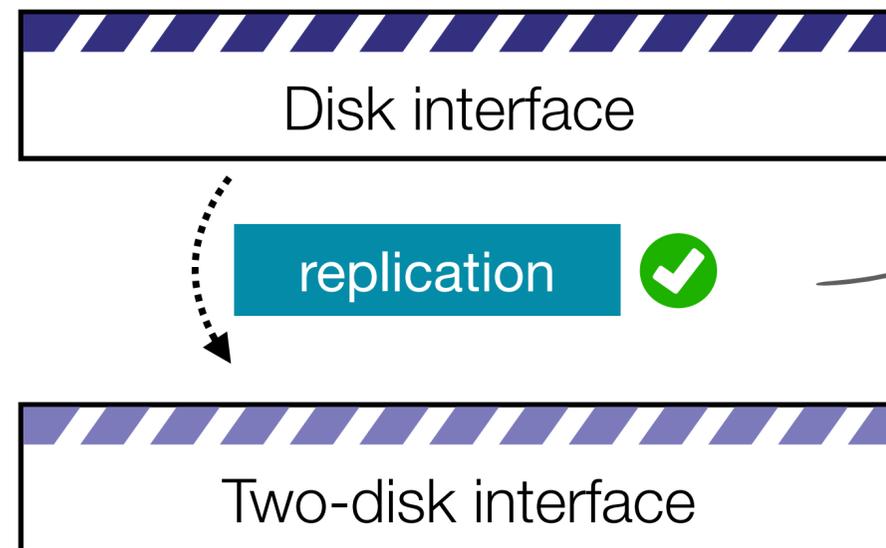
Recovery refinement for modular proofs

see paper CHL for proving recovery refinement

see paper Verified example: logging + replication

see code Machine-checked proofs in Coq 

Preview: recovery refinement



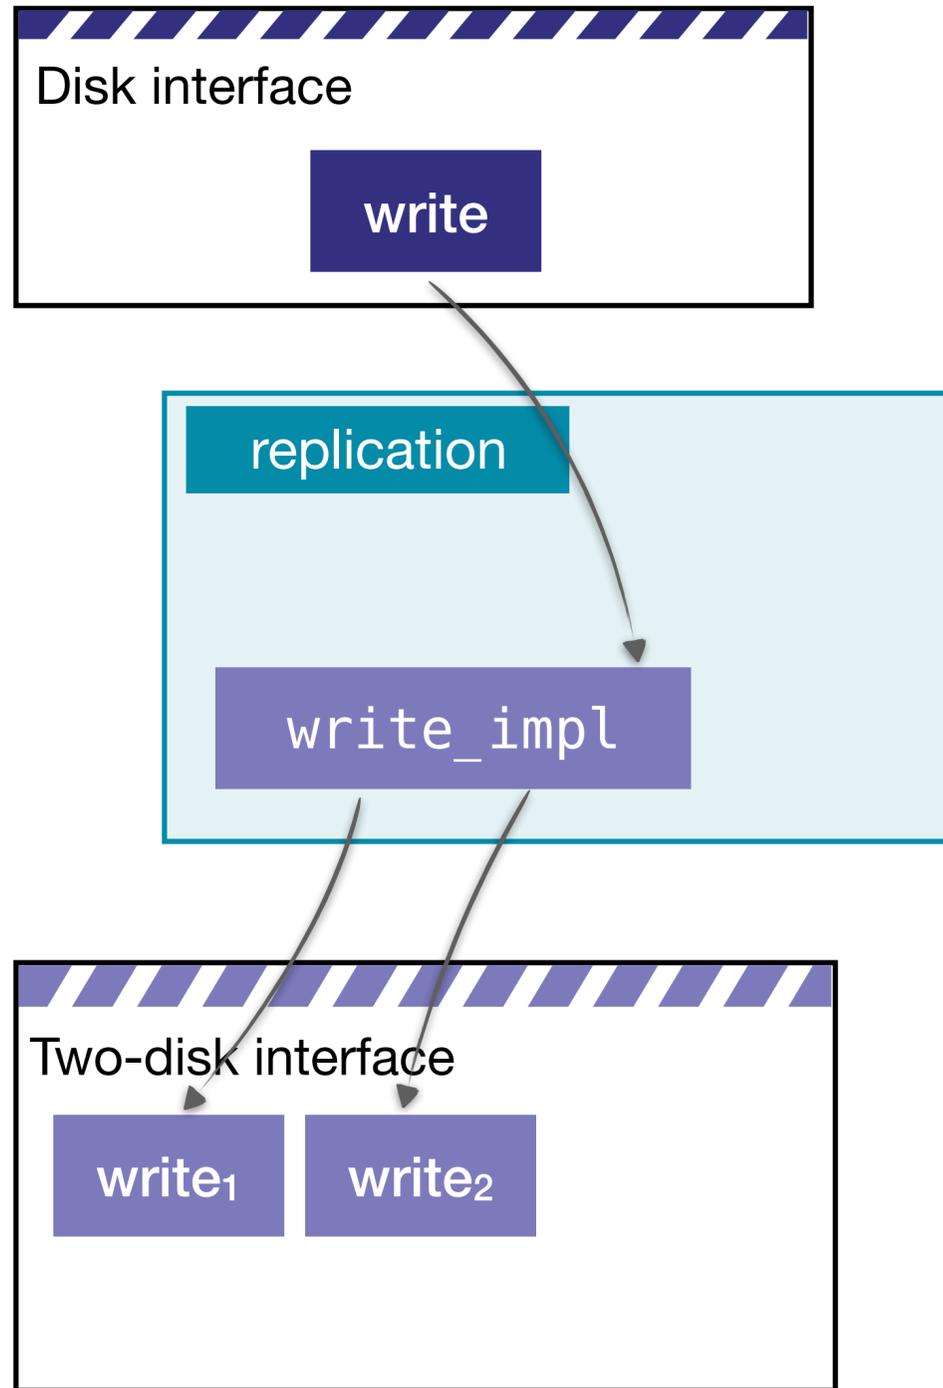
1. Normal execution correctness using *refinement*
2. Crash and recovery correctness using *recovery refinement*

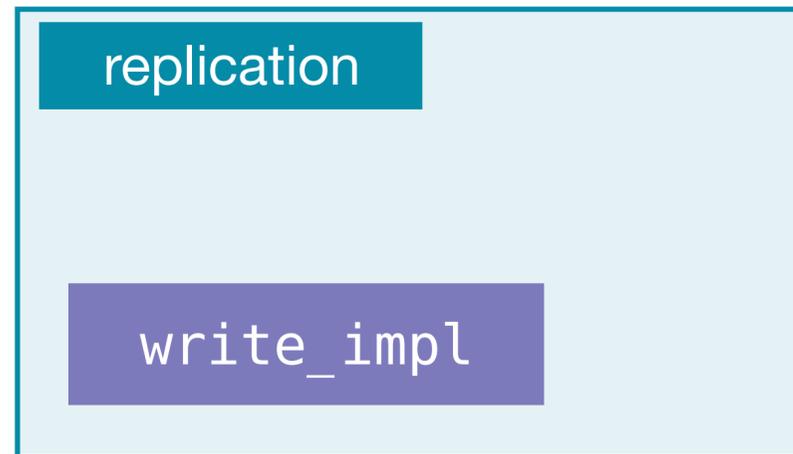
Refinement

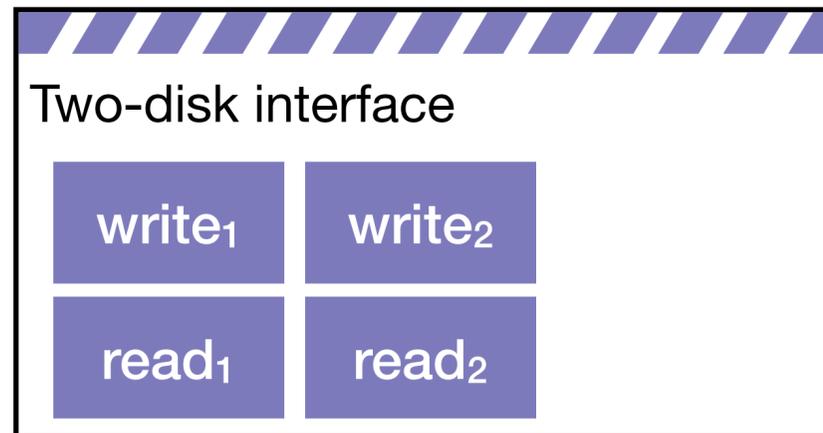
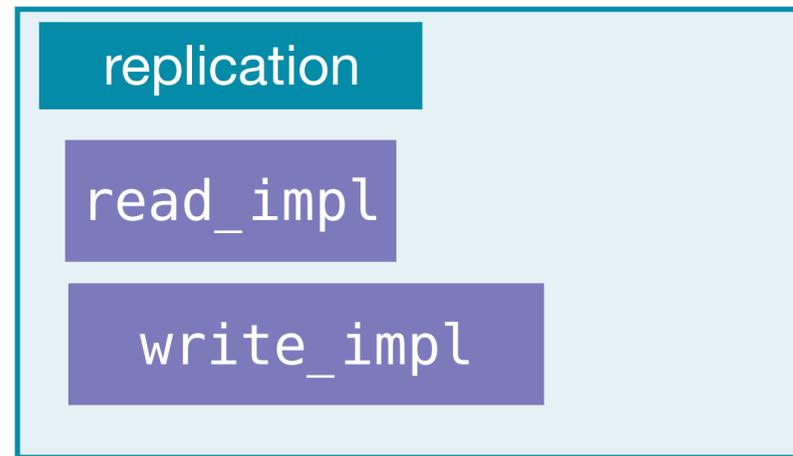
Disk interface

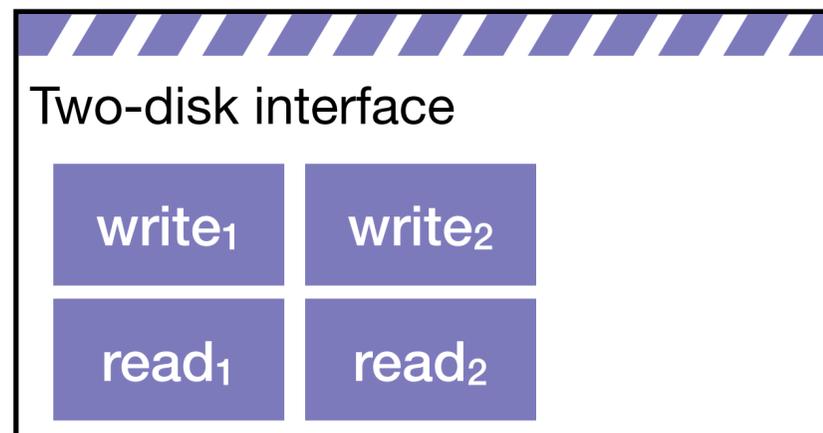
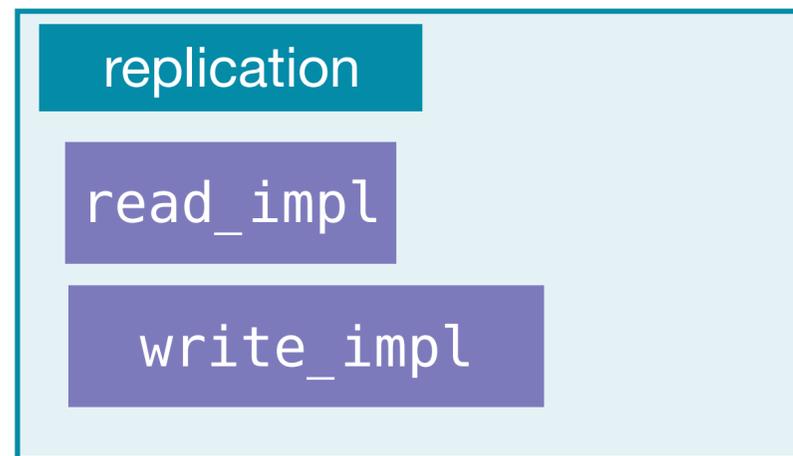
replication

Two-disk interface

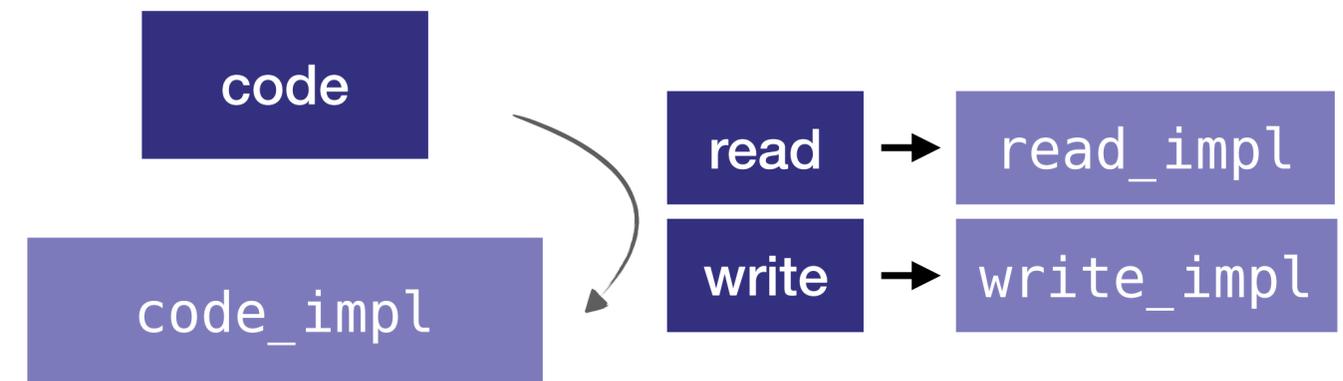




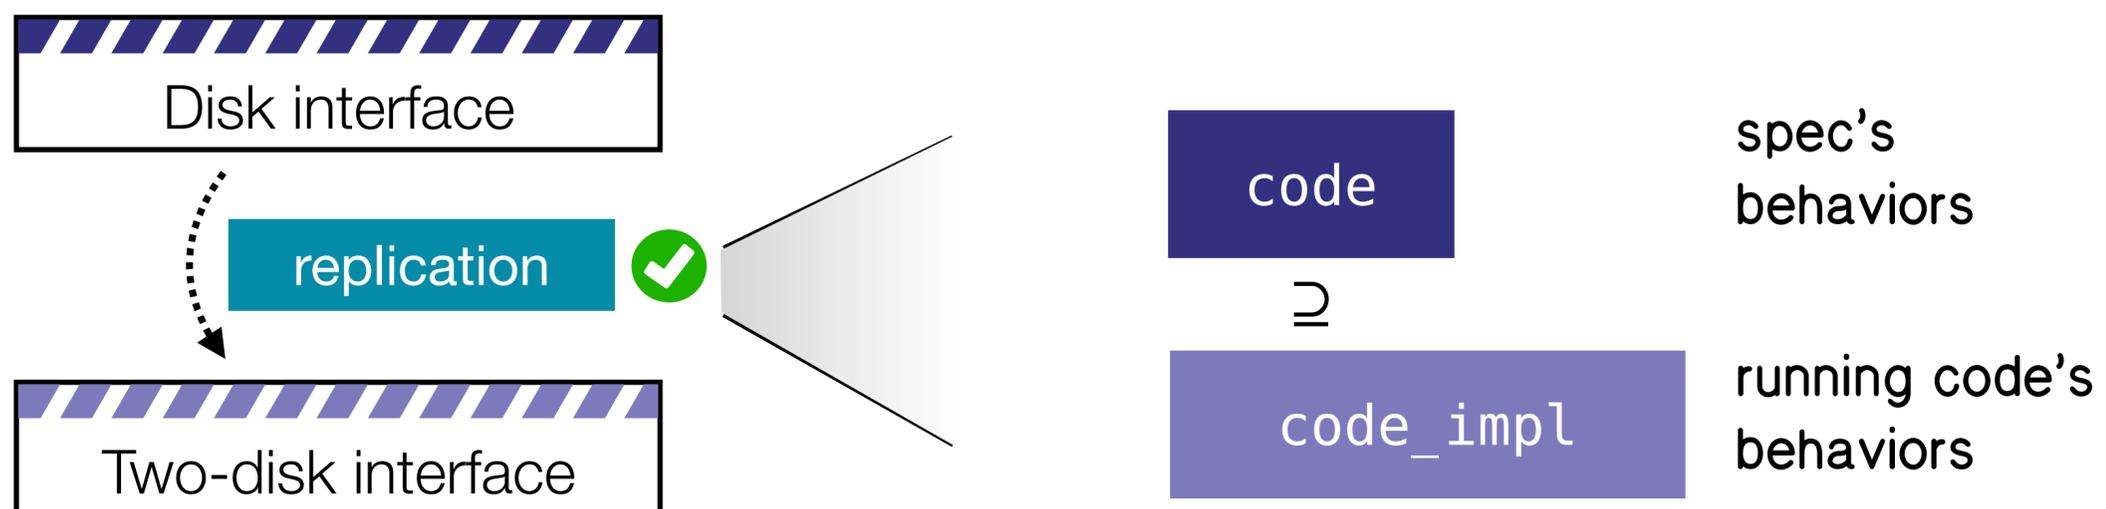




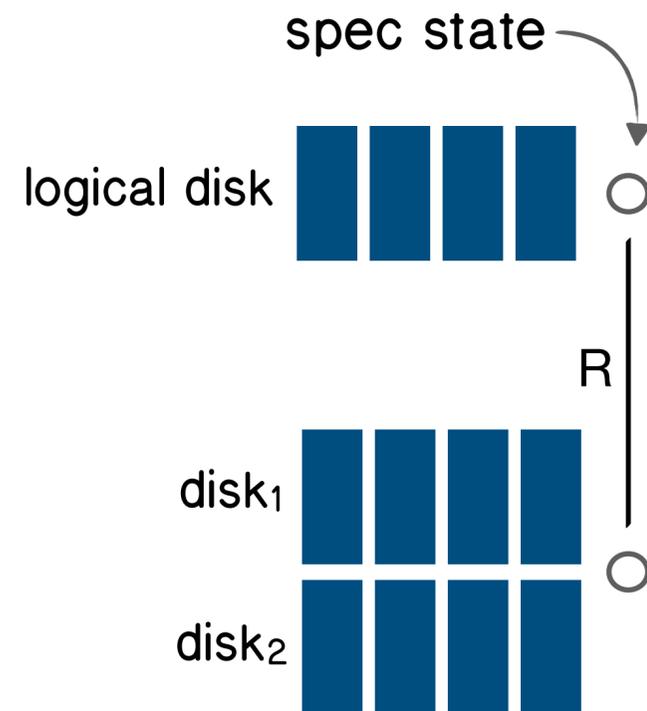
correctness is based on how we use replication:
run code using Disk interface on top of two disks



Correctness: trace inclusion

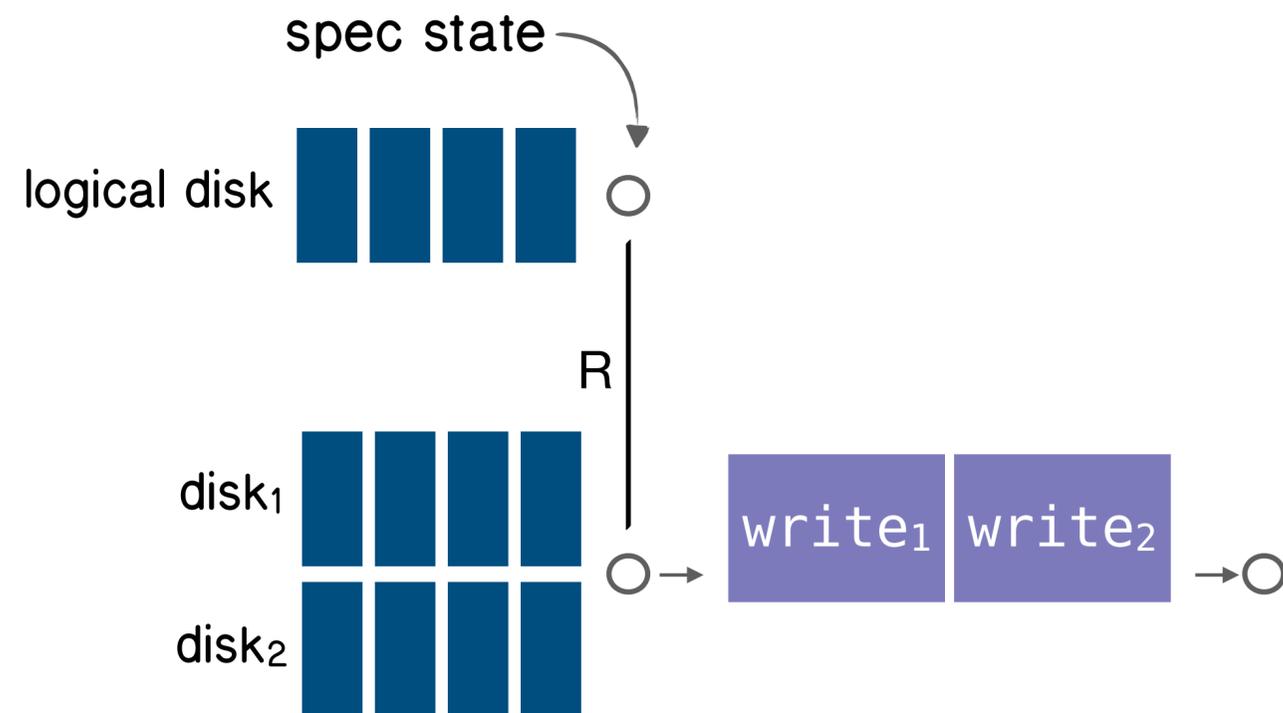


Proving correctness with an abstraction relation



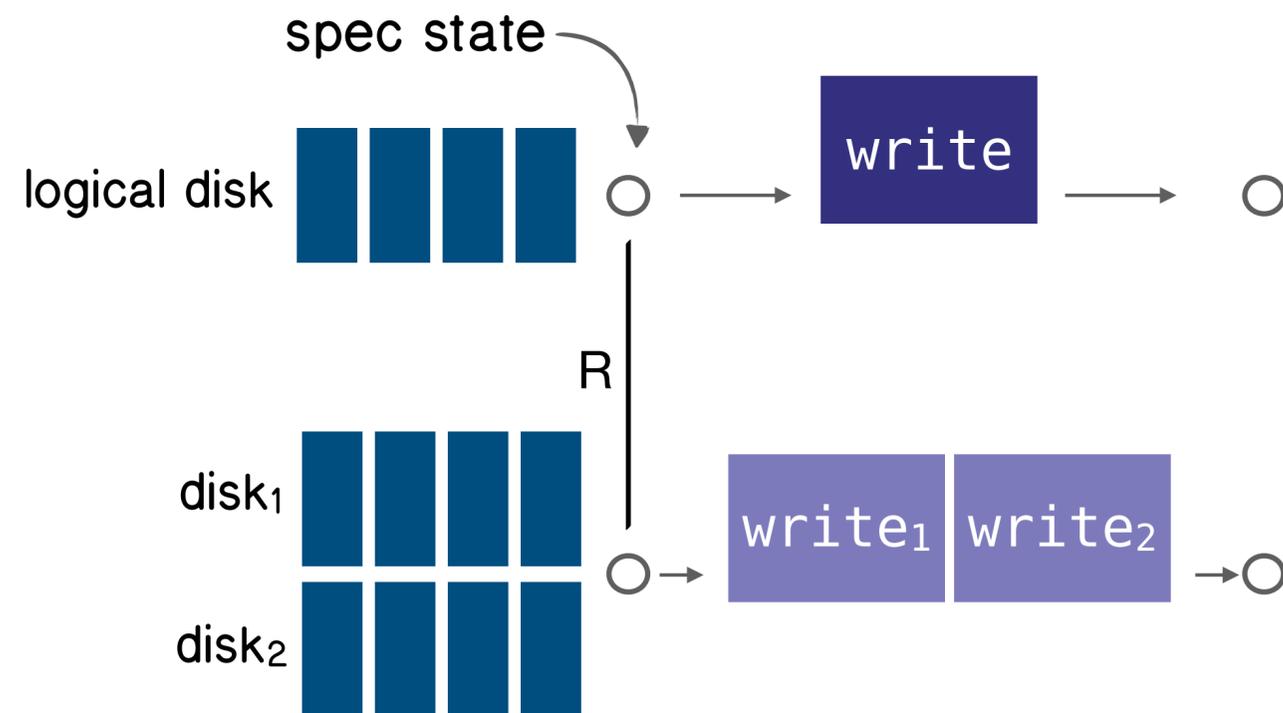
1. developer provides abstraction relation R

Proving correctness with an abstraction relation



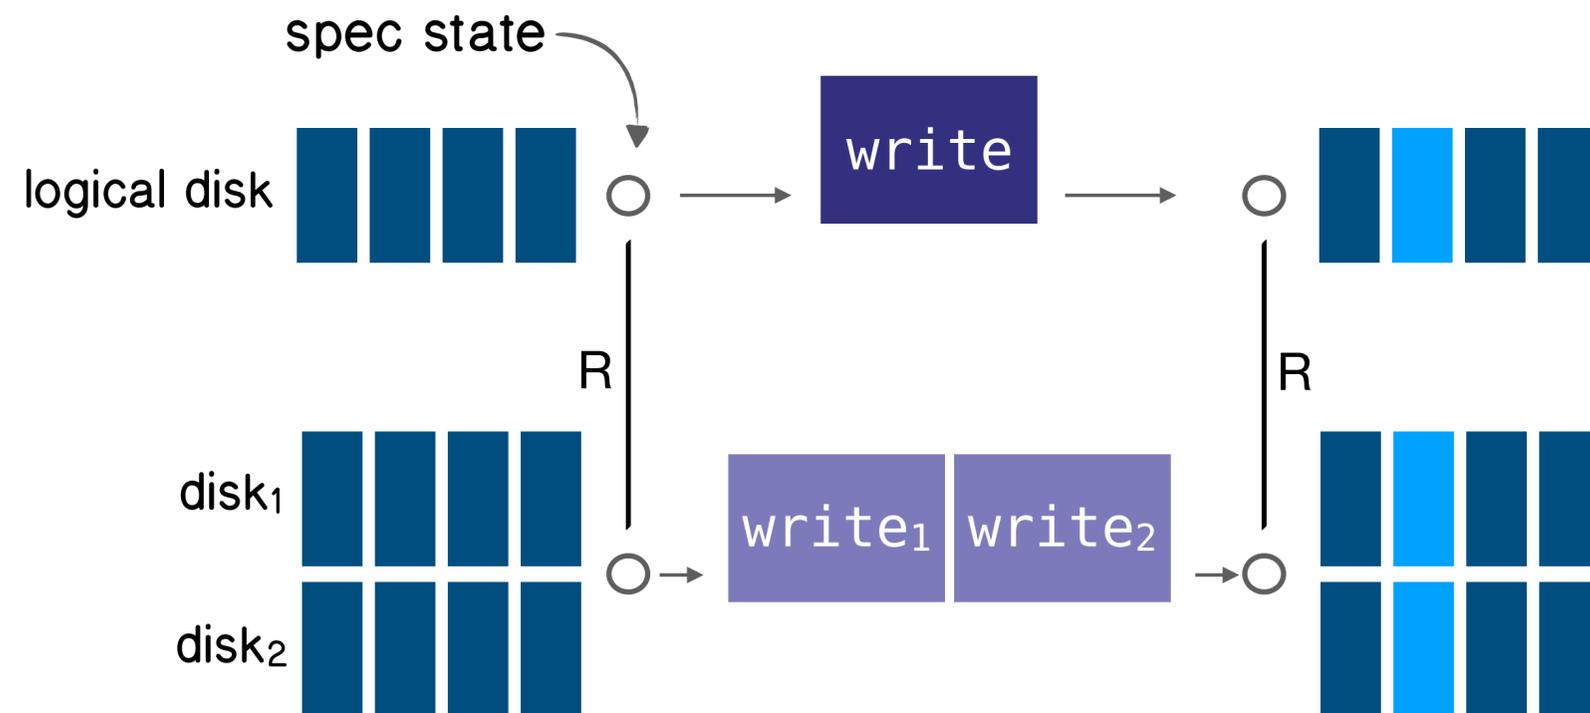
1. developer provides abstraction relation R

Proving correctness with an abstraction relation



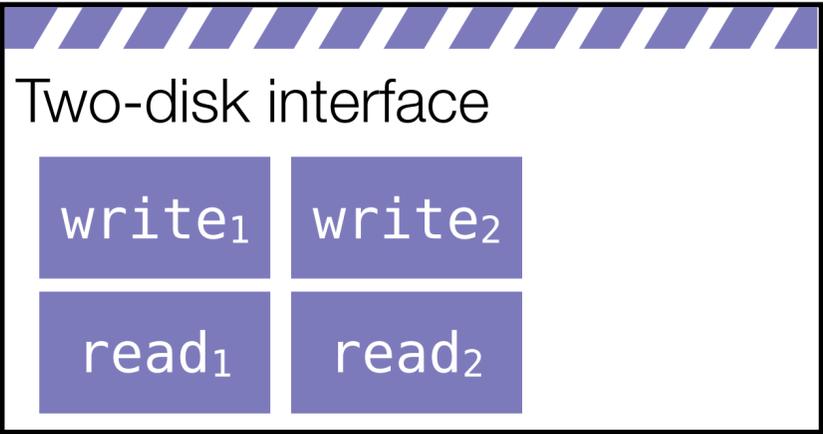
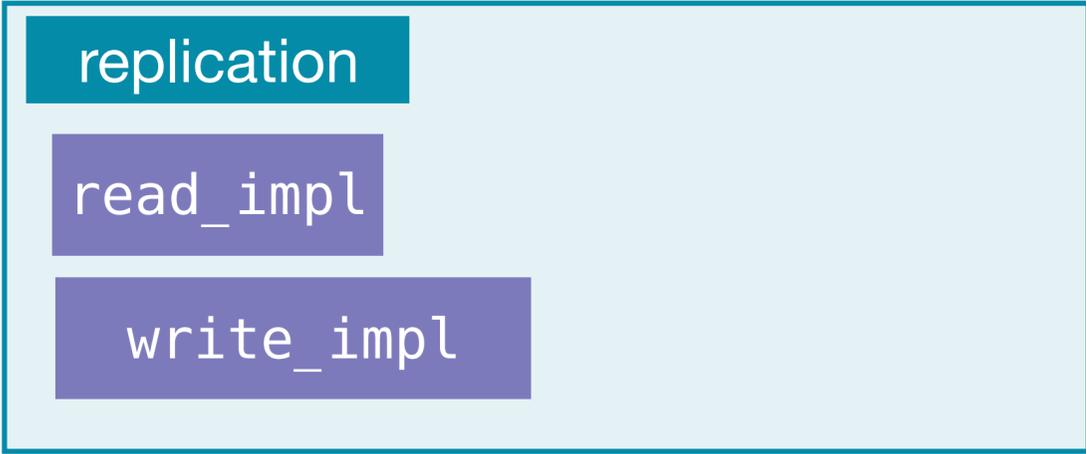
1. developer provides abstraction relation R
2. prove spec execution exists

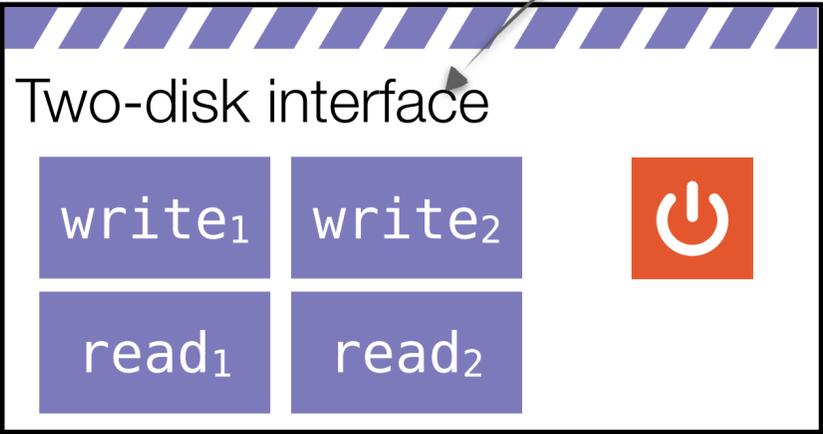
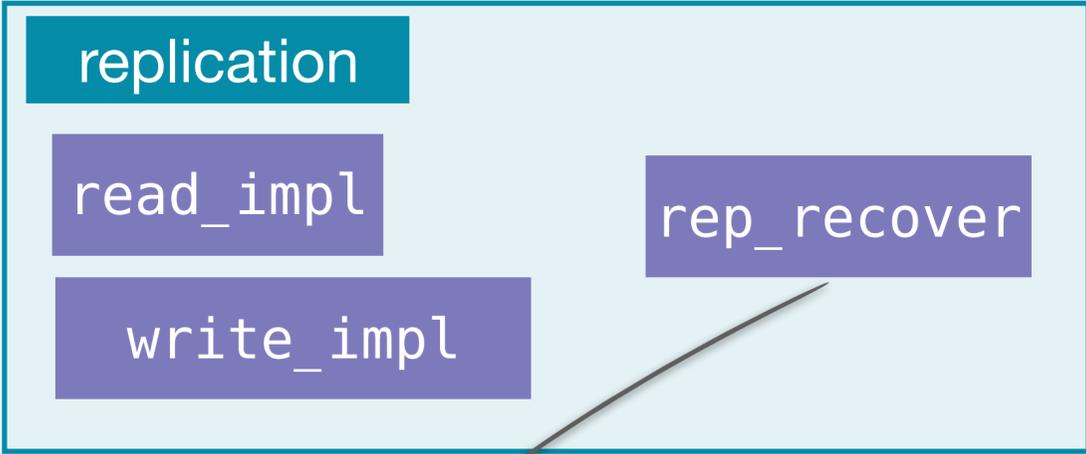
Proving correctness with an abstraction relation

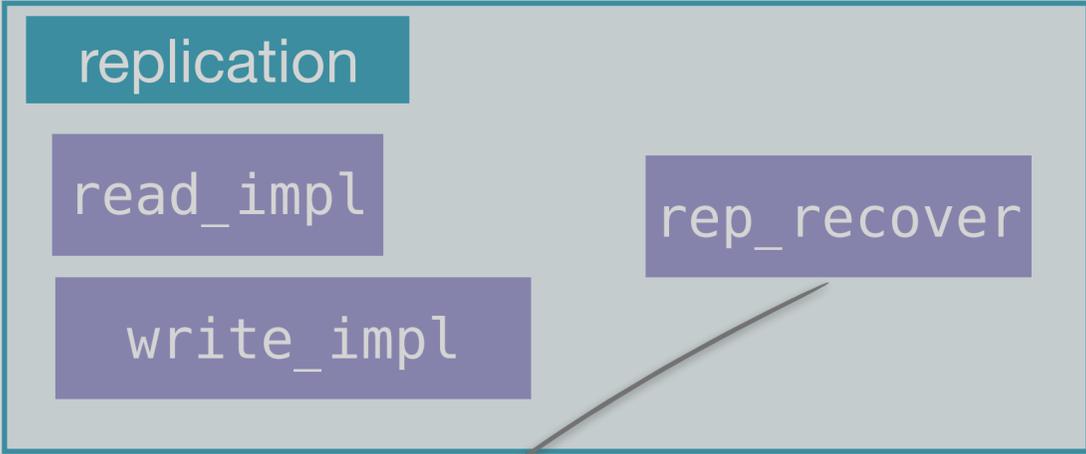


1. developer provides abstraction relation R
2. prove spec execution exists
3. and abstraction relation is preserved

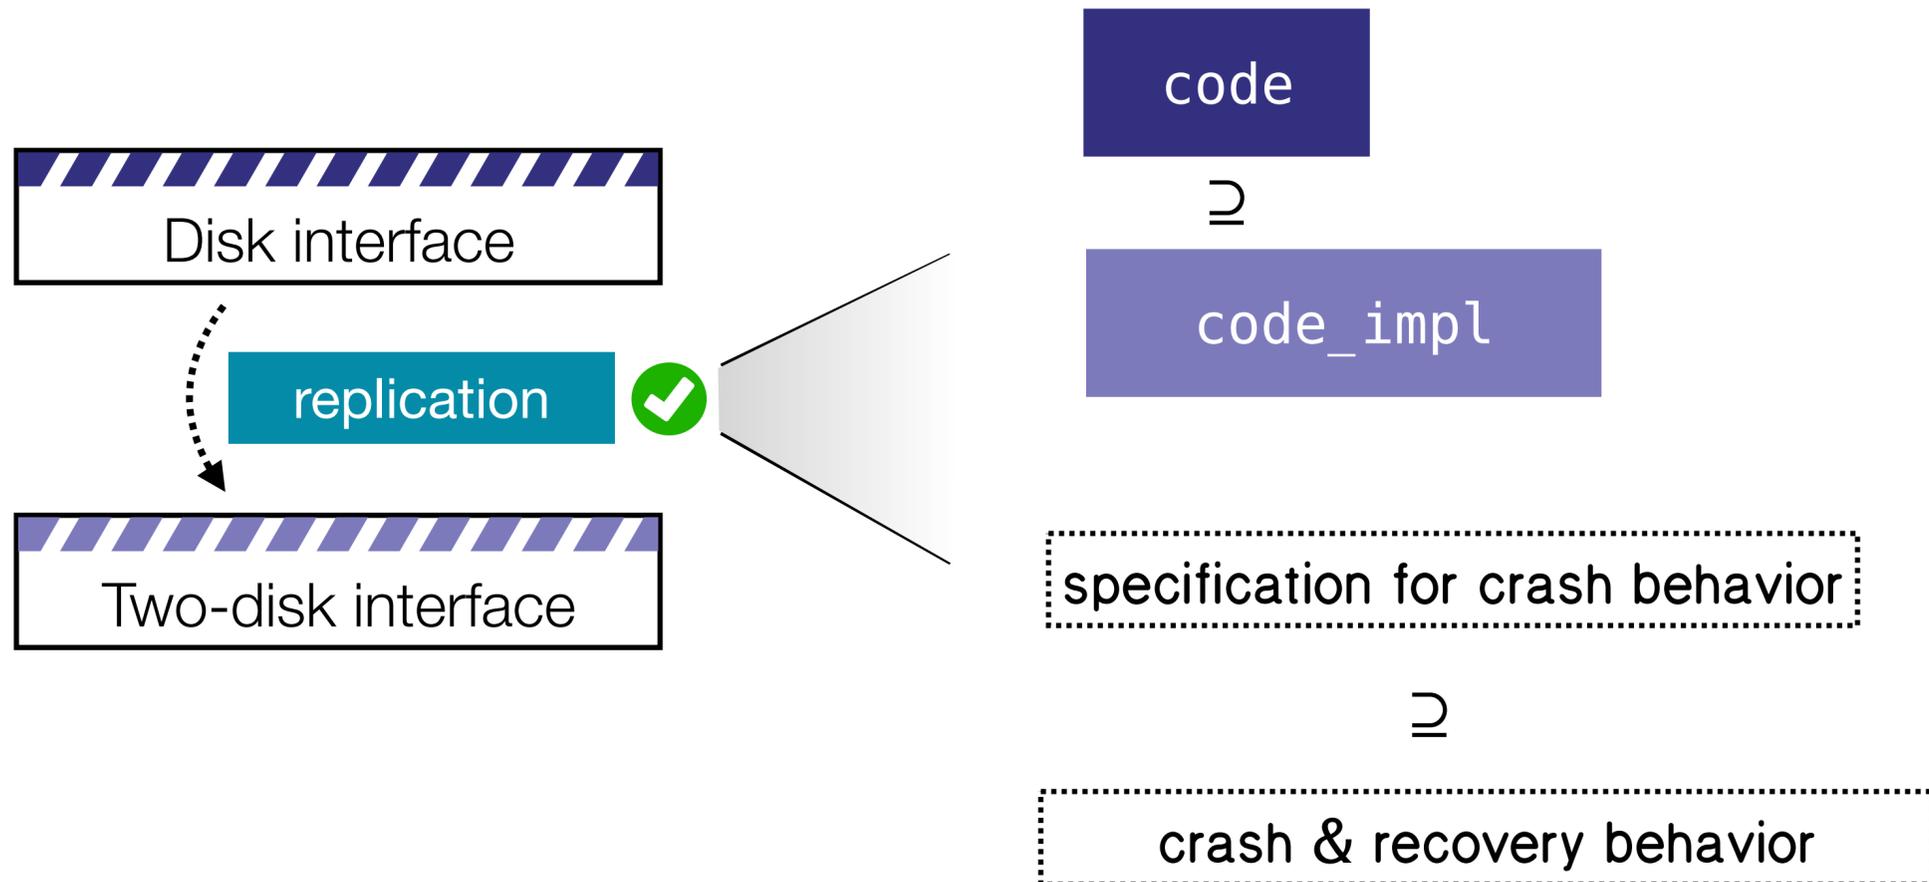
Recovery refinement



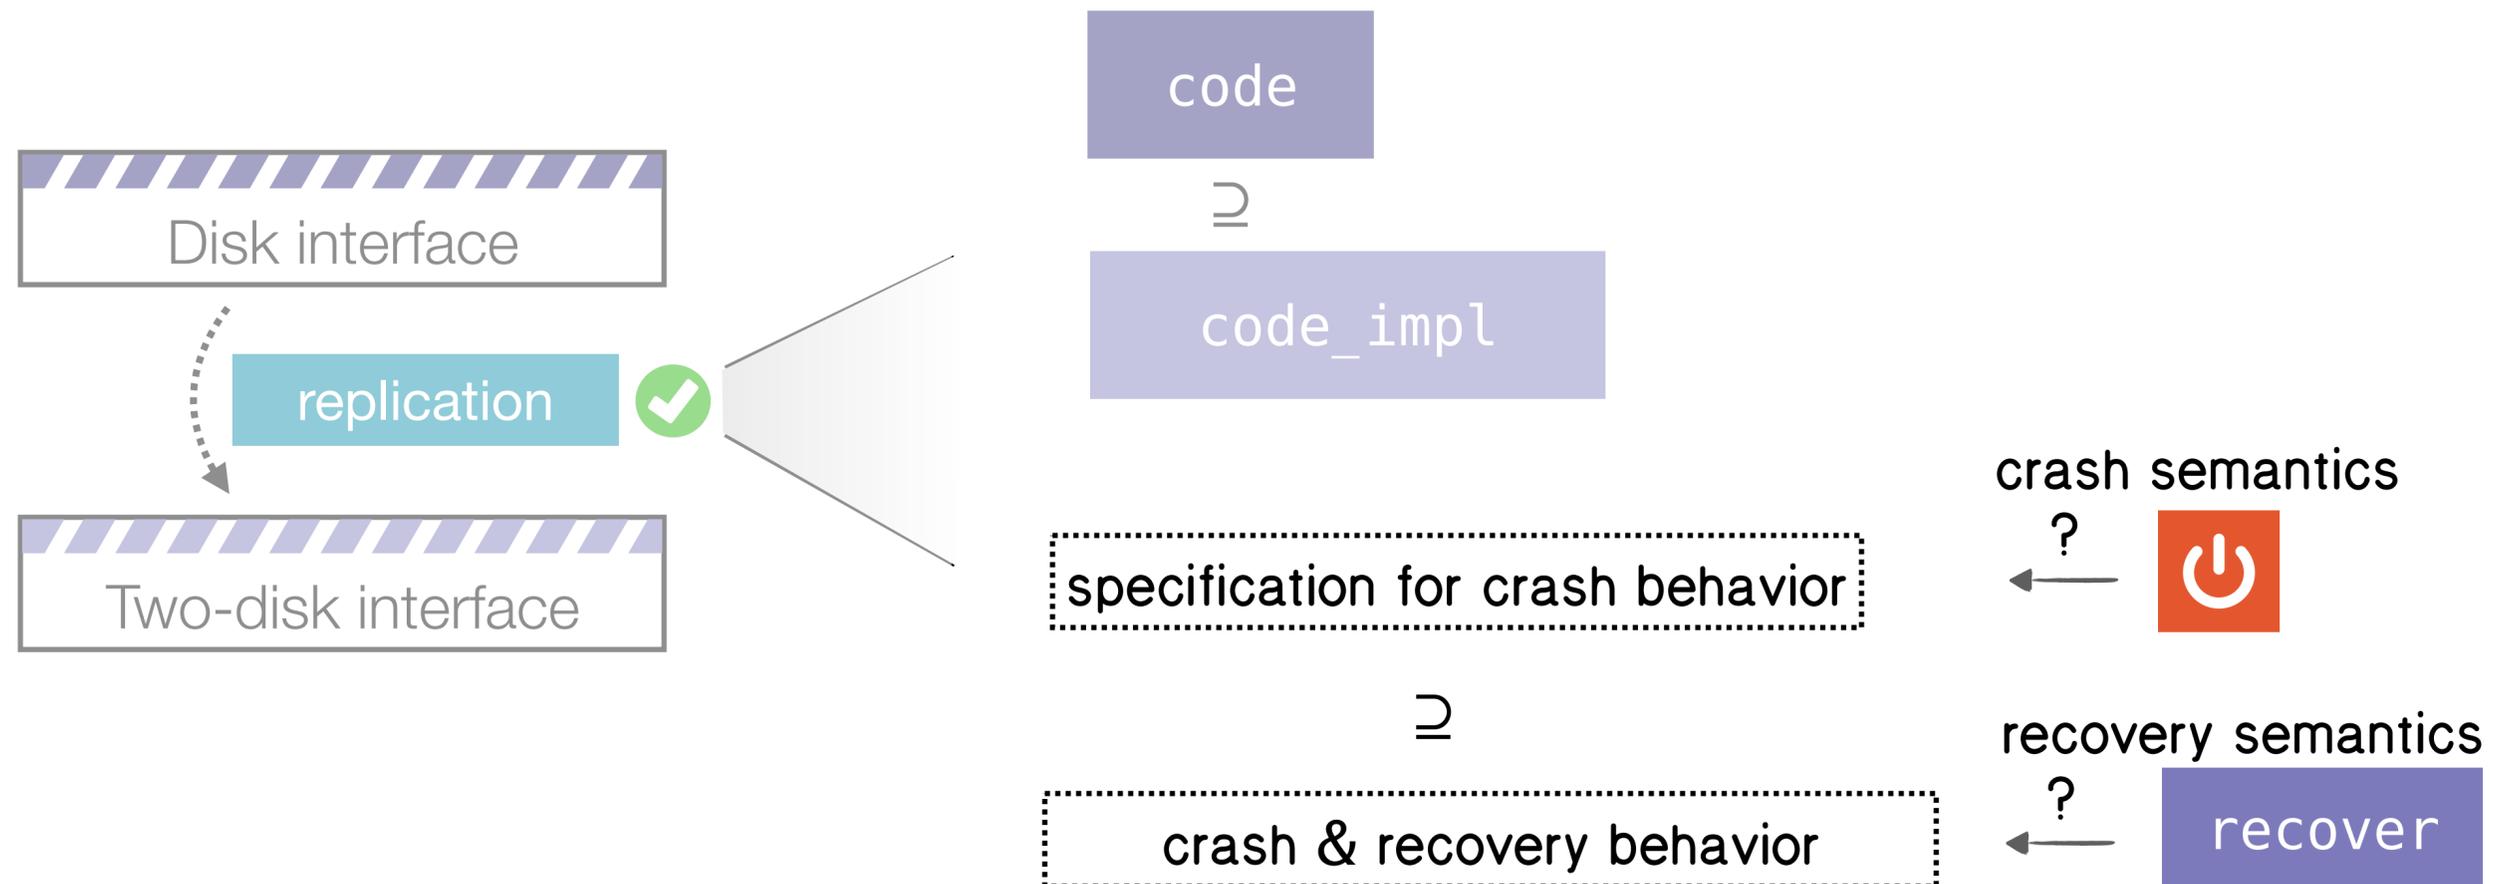


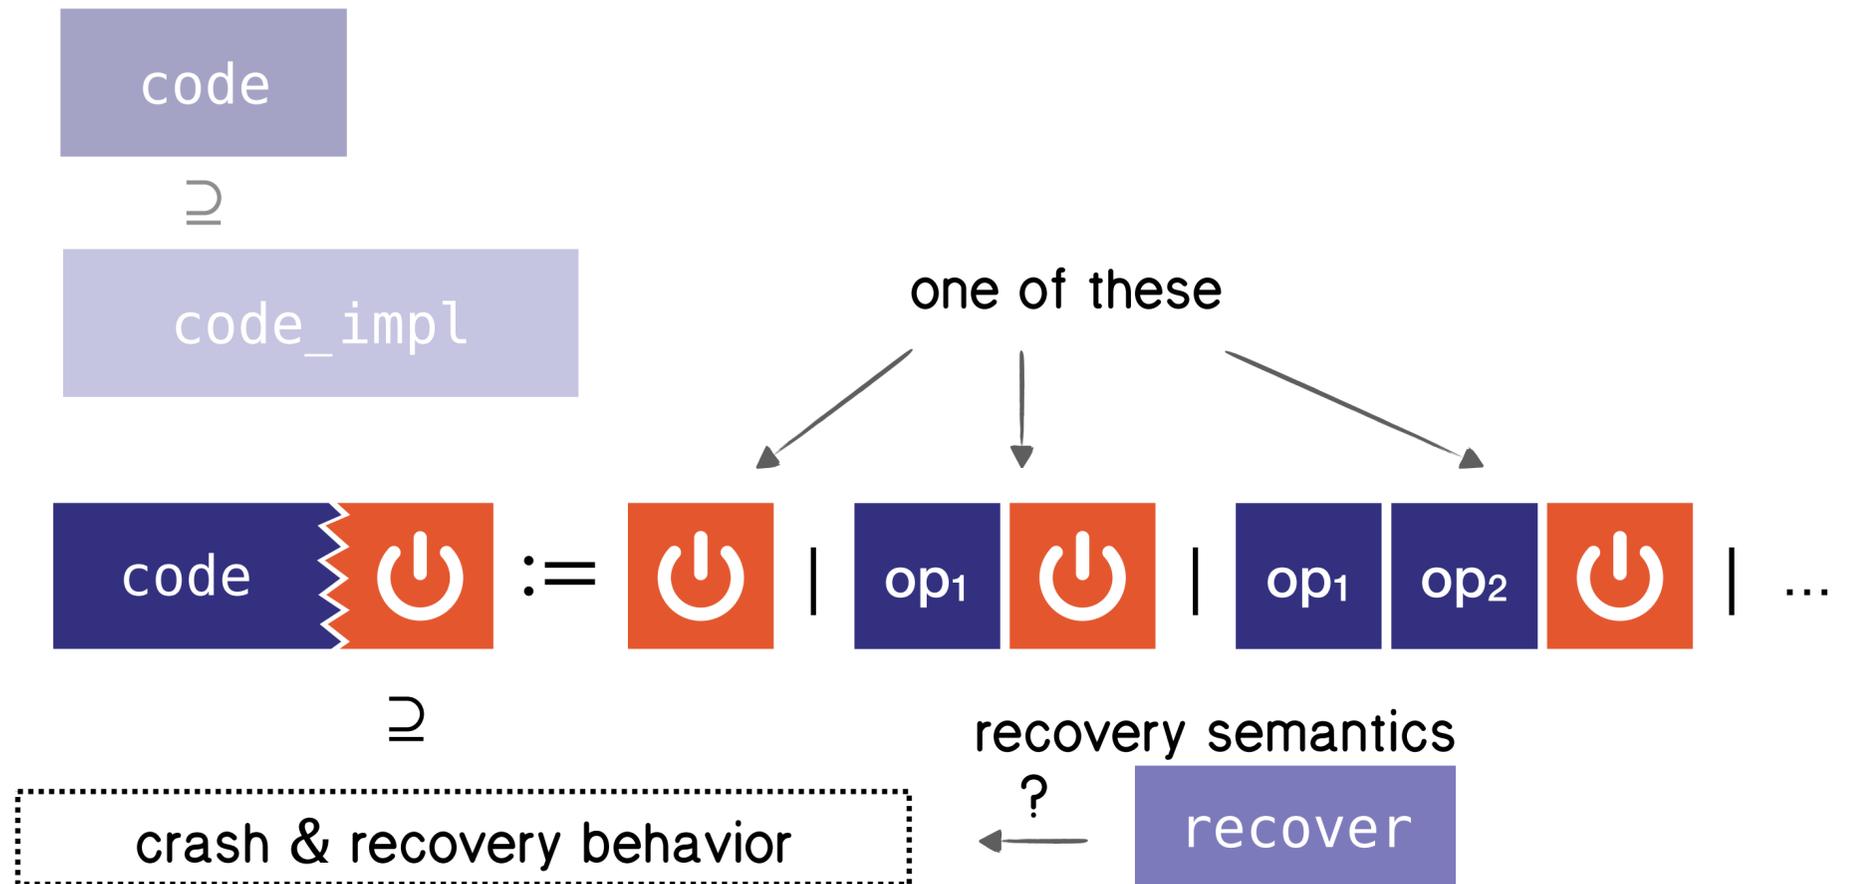
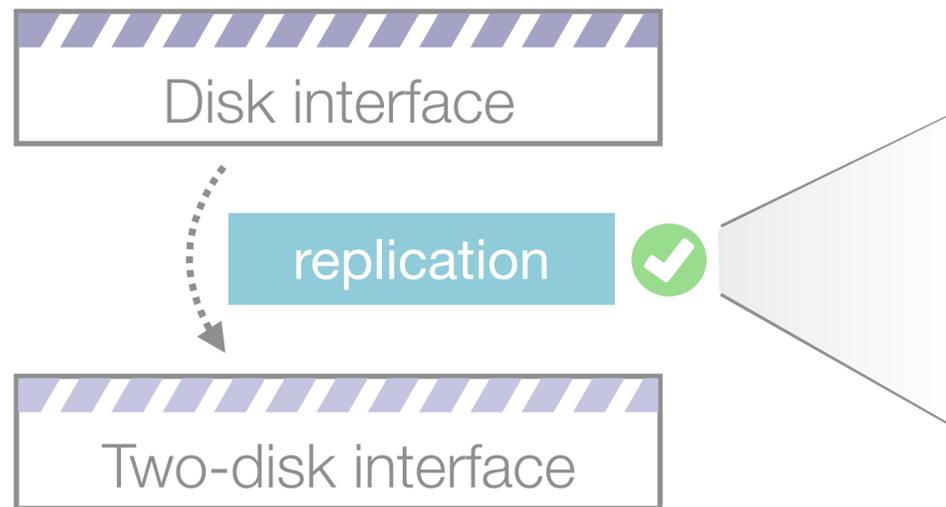


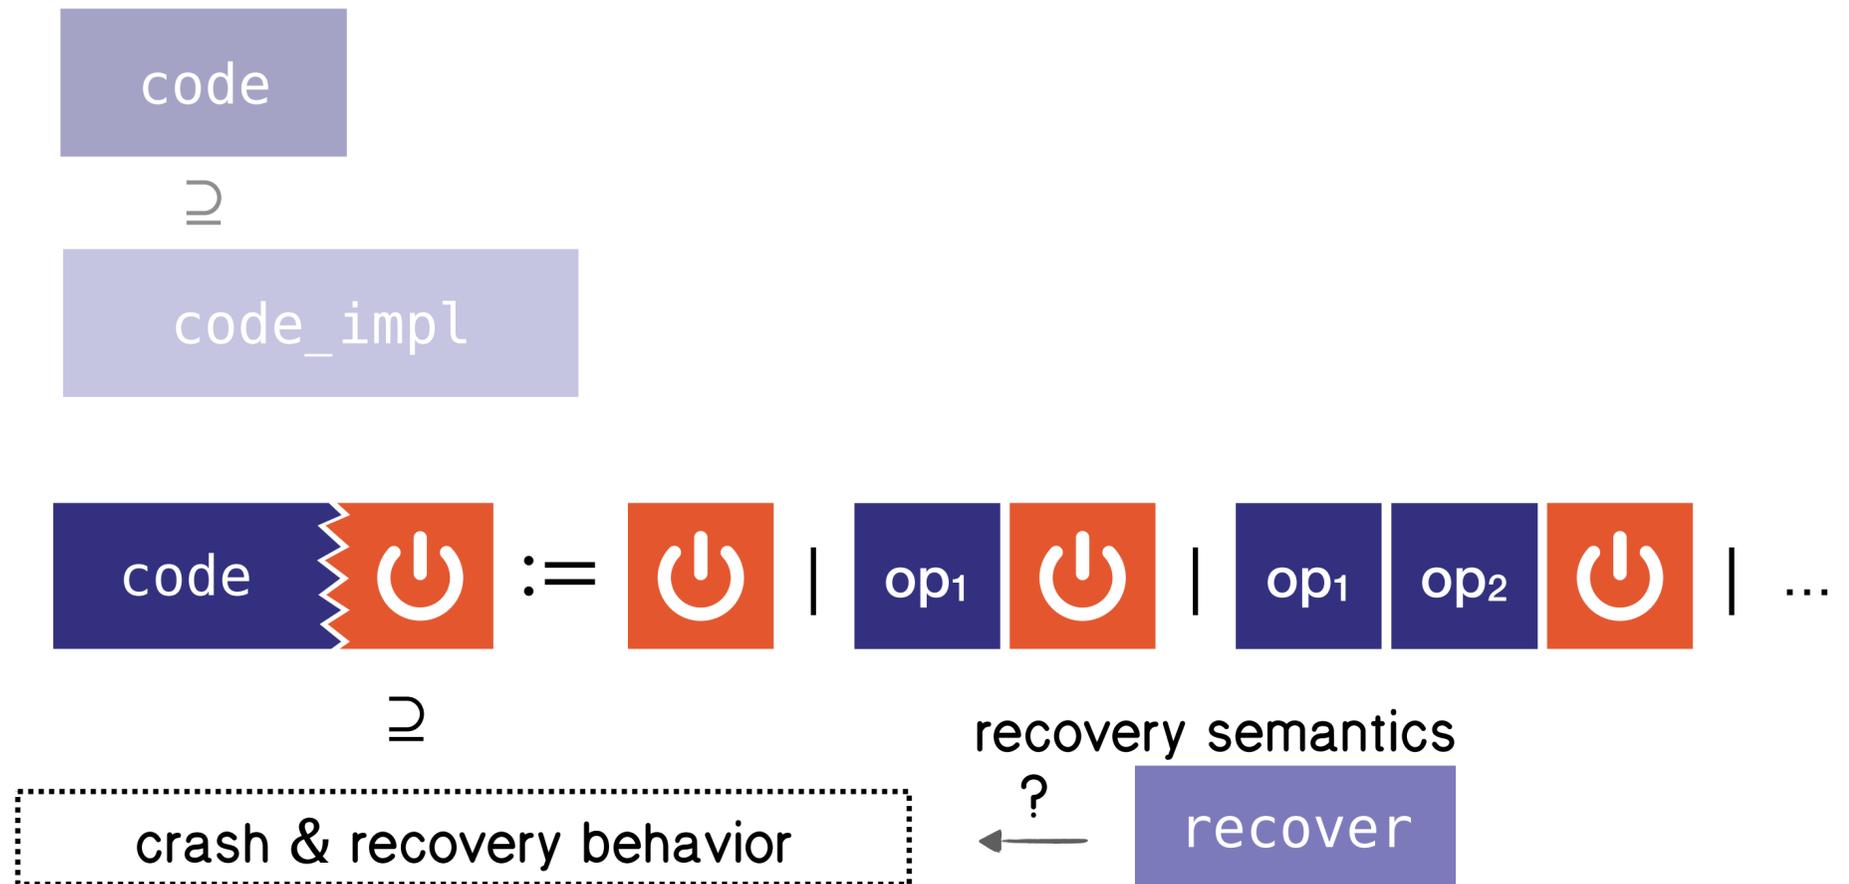
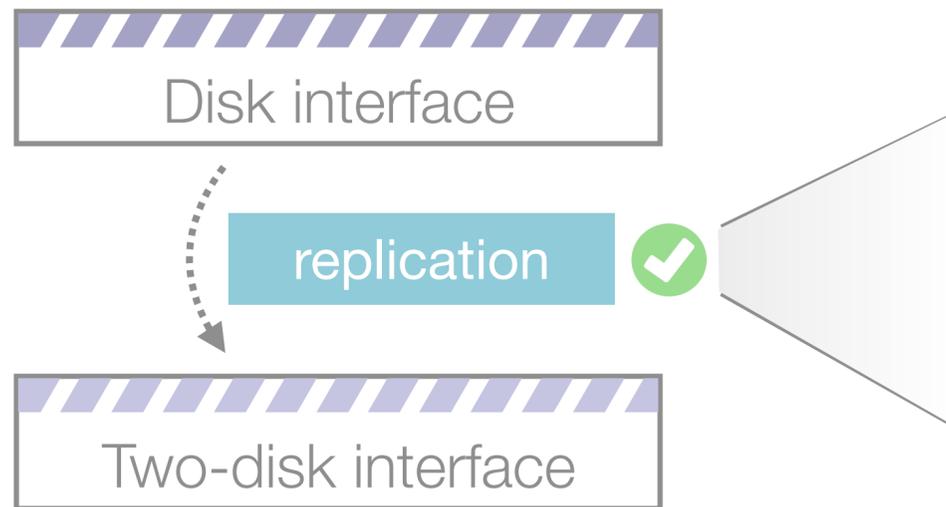
Extending trace inclusion with recovery

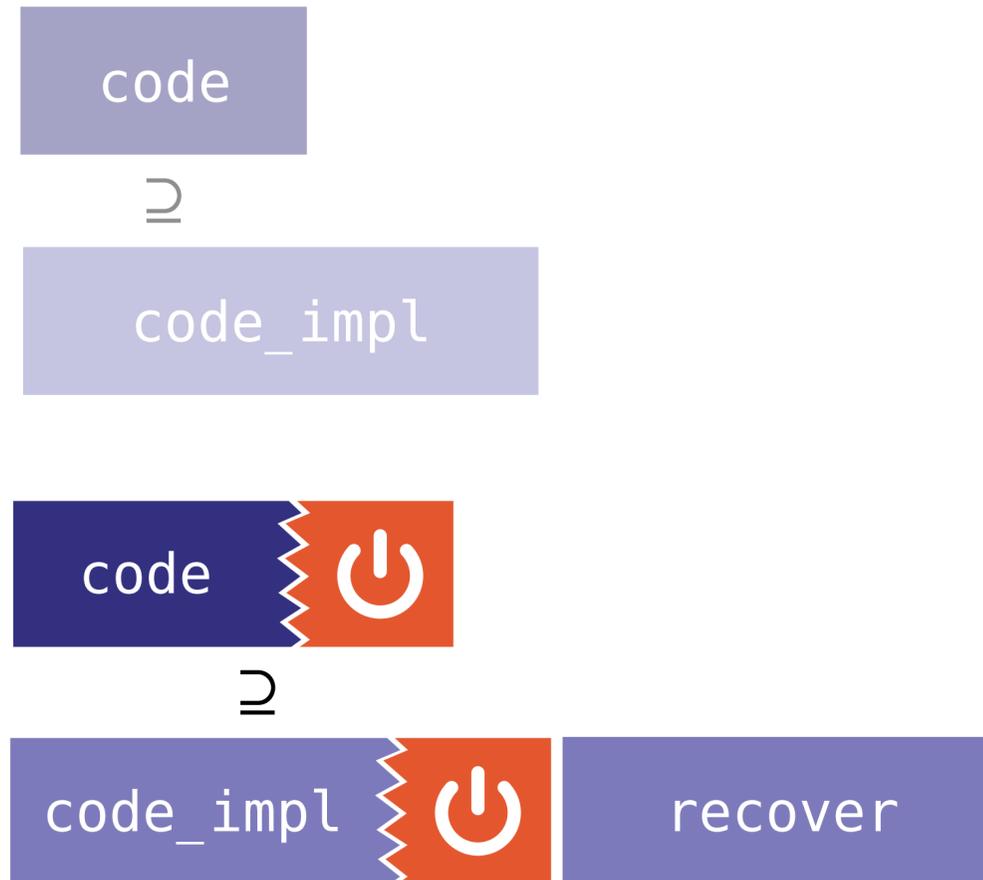
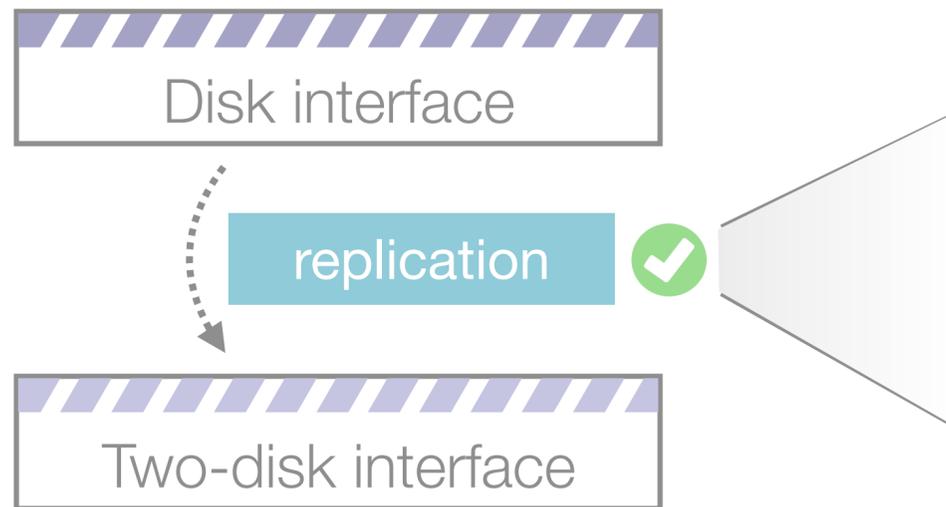


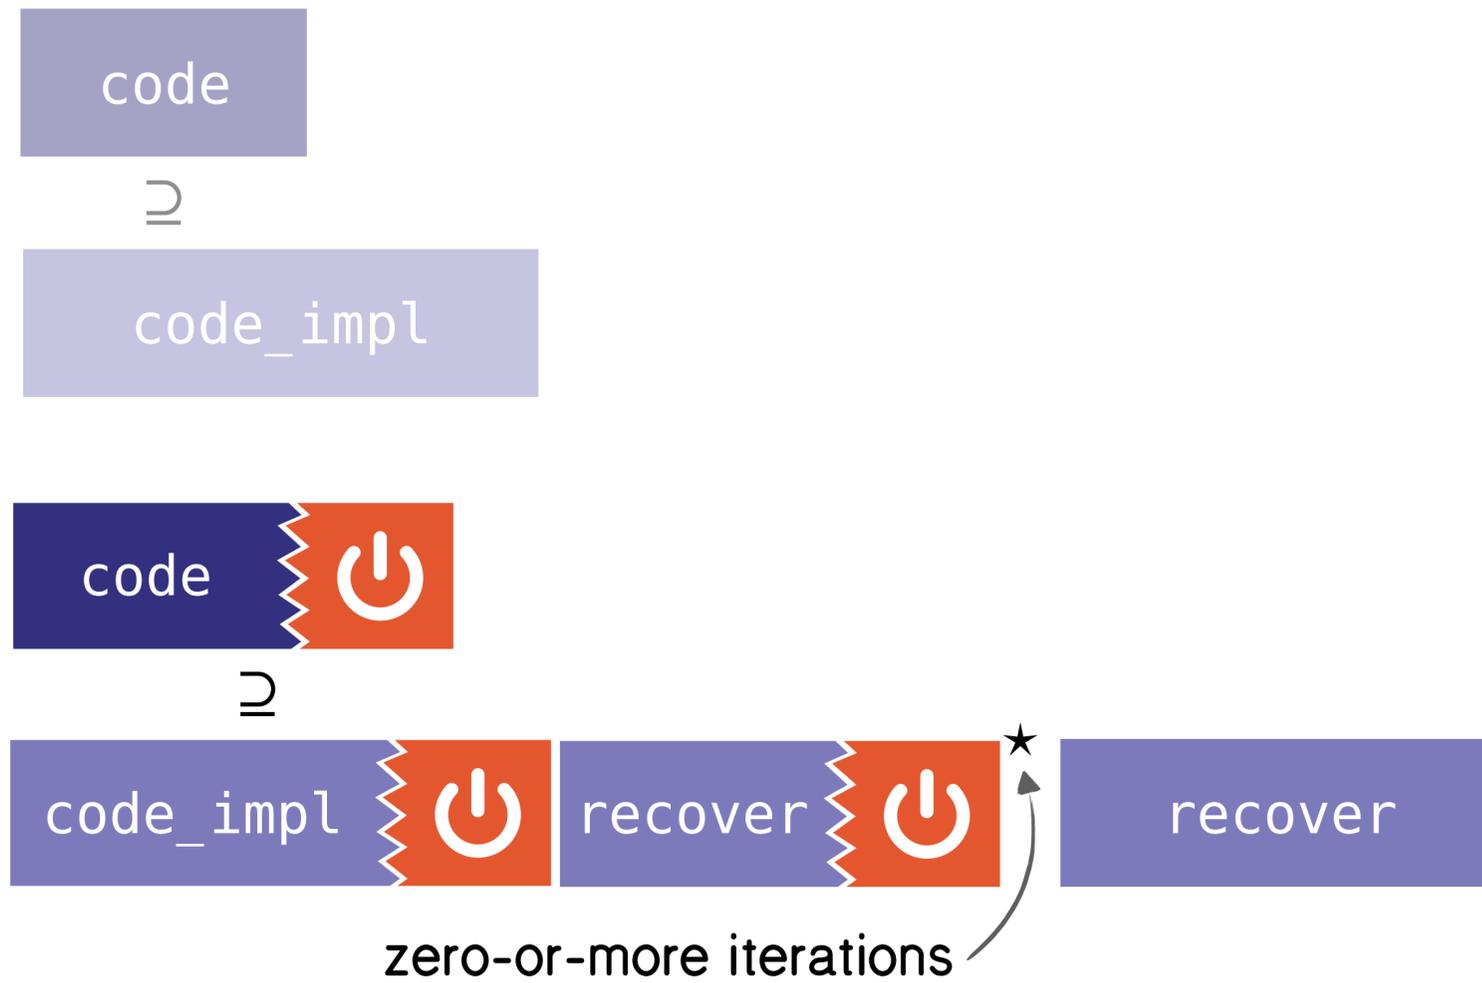
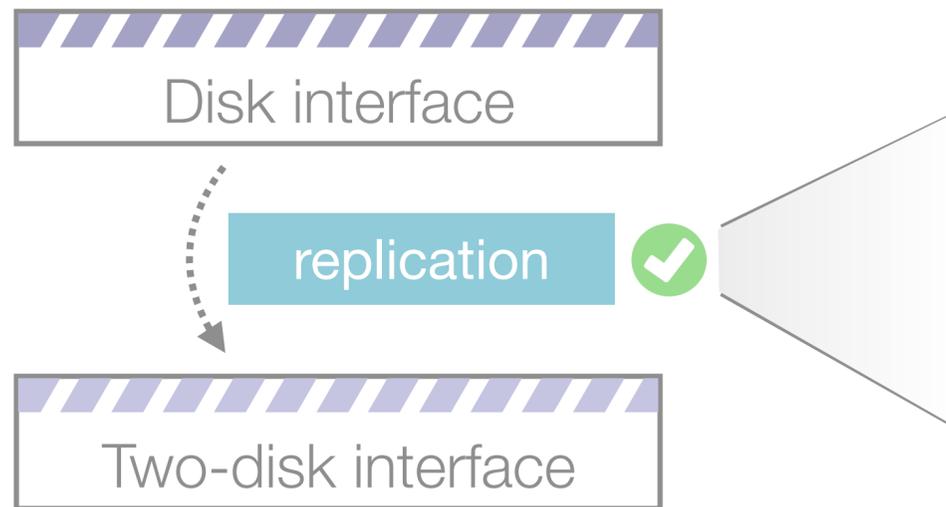
Extending trace inclusion with recovery

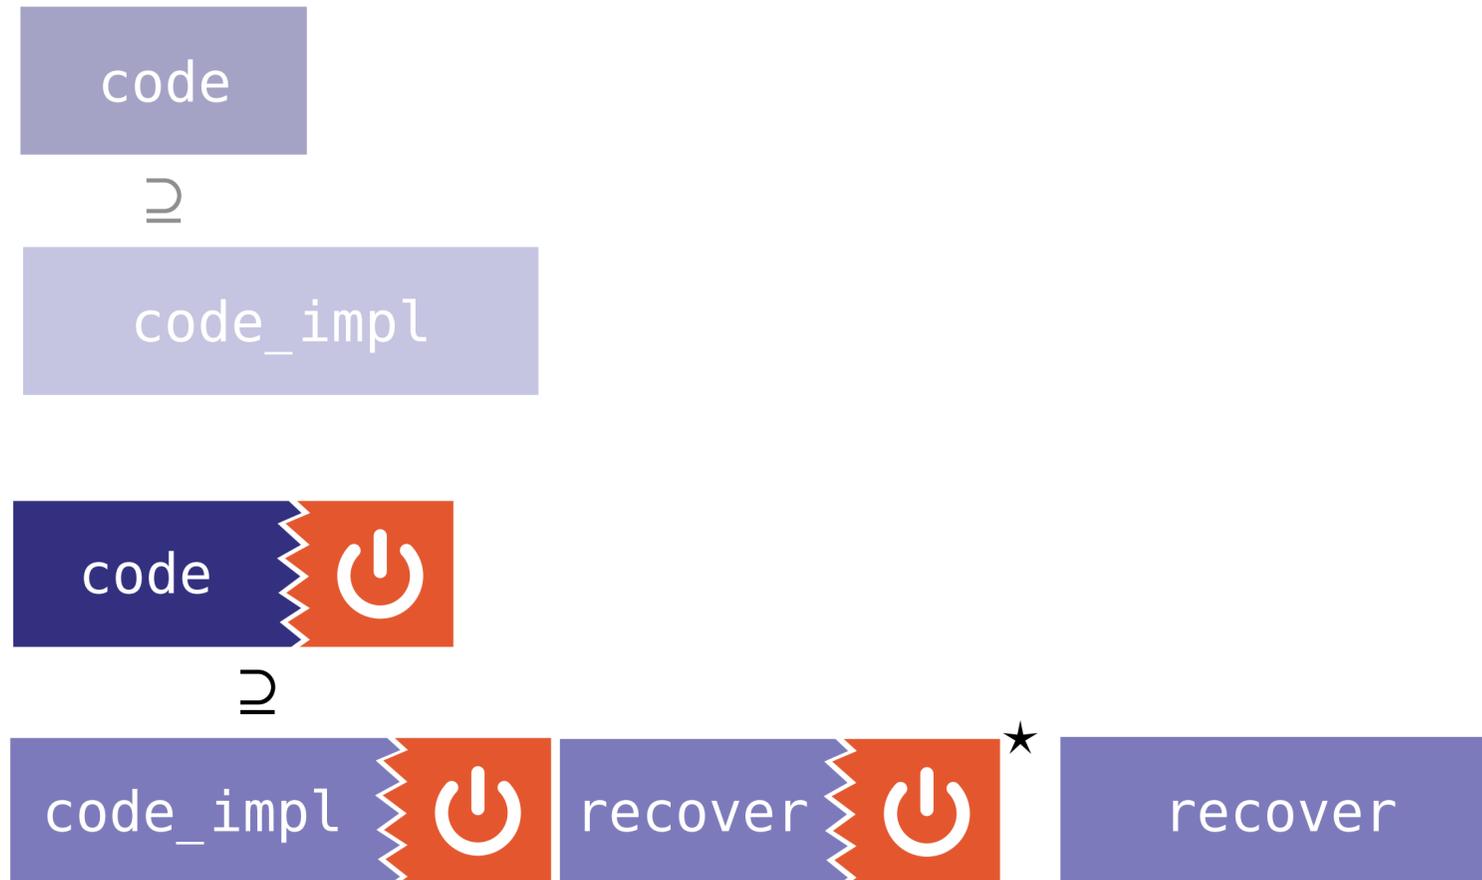
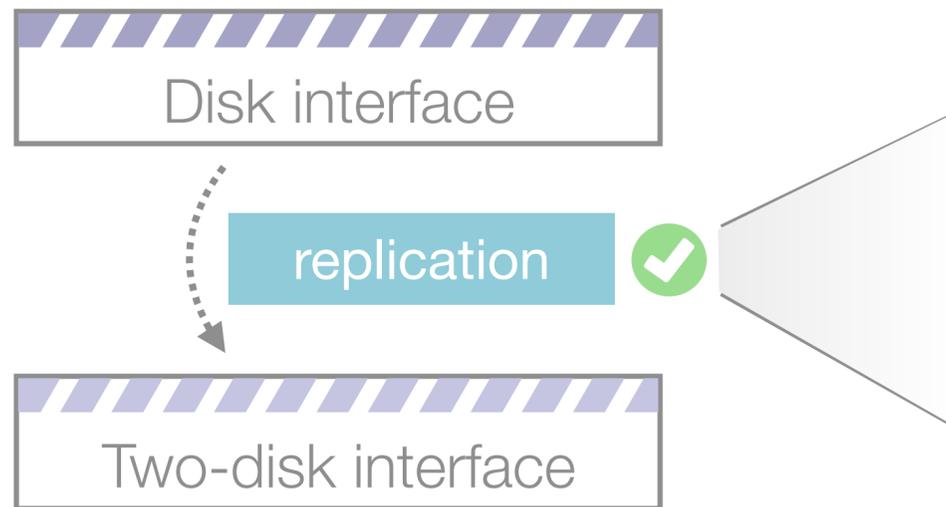




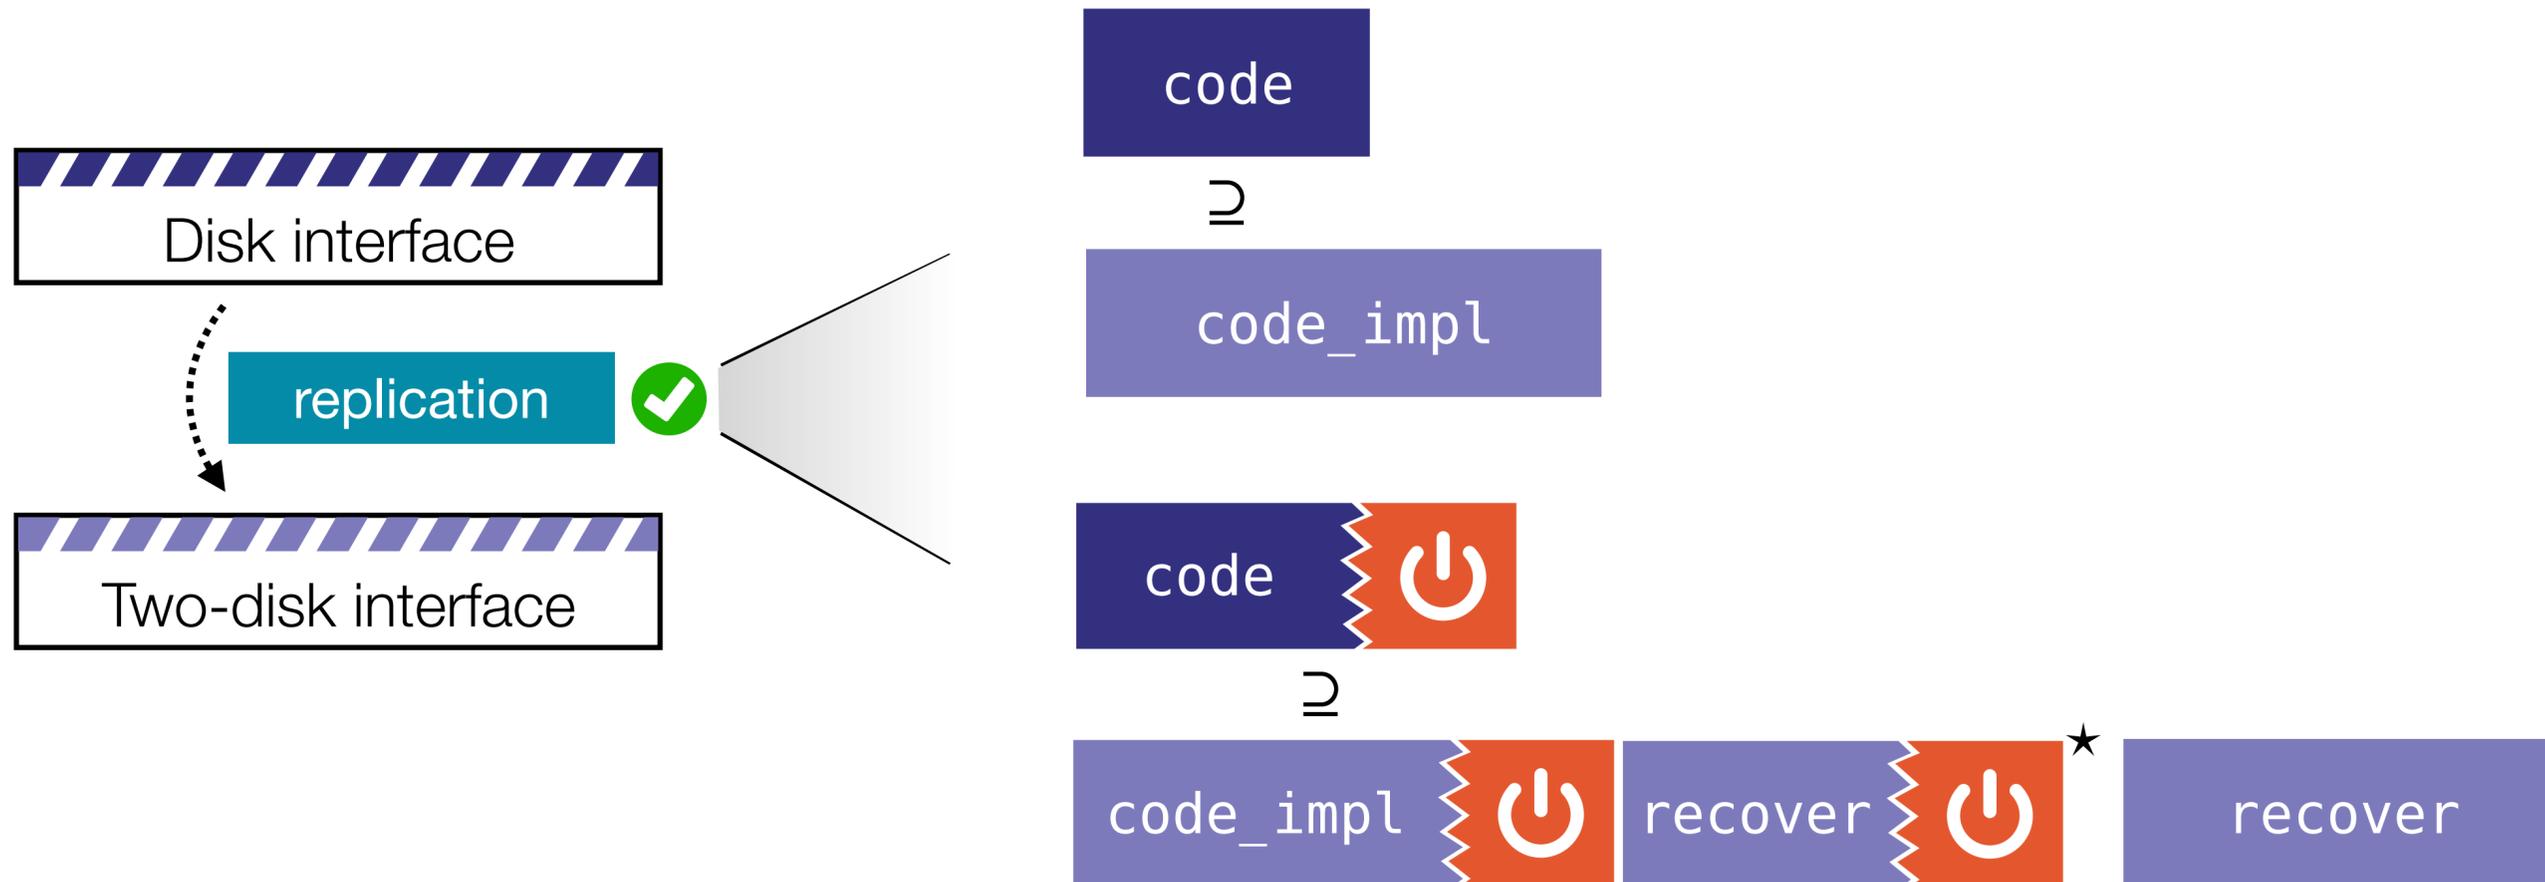








Trace inclusion, with recovery



Proving trace inclusion, with recovery



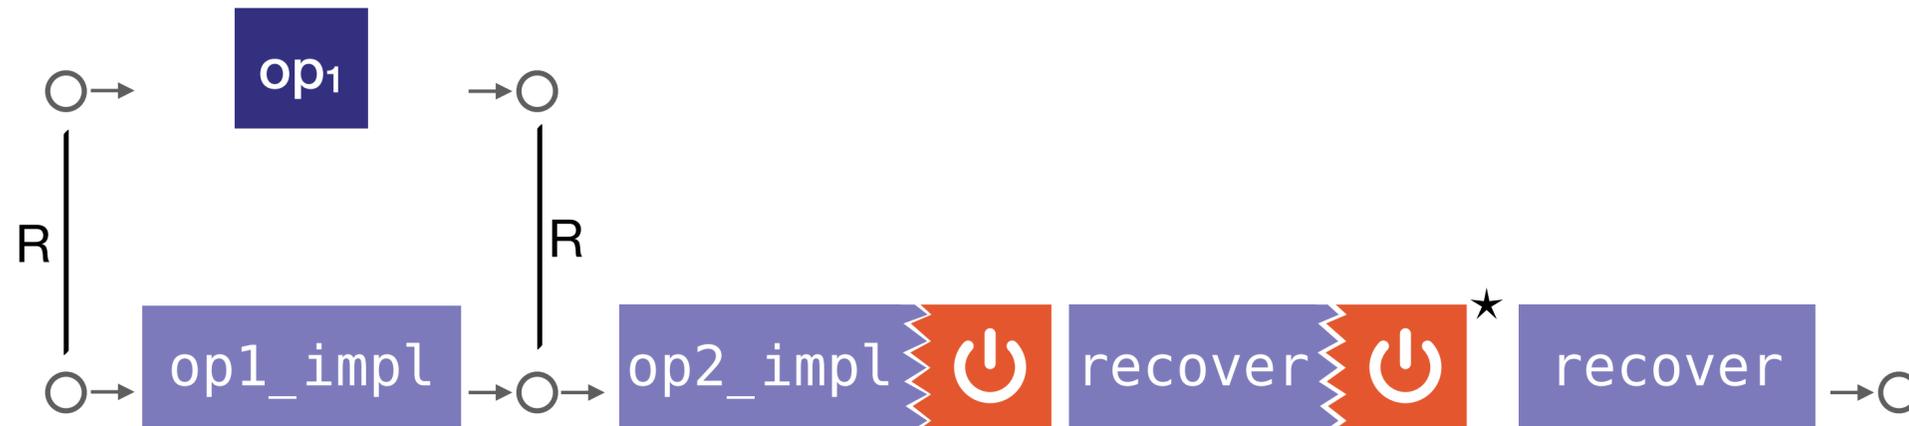
Proving trace inclusion, with recovery



Proving trace inclusion, with recovery



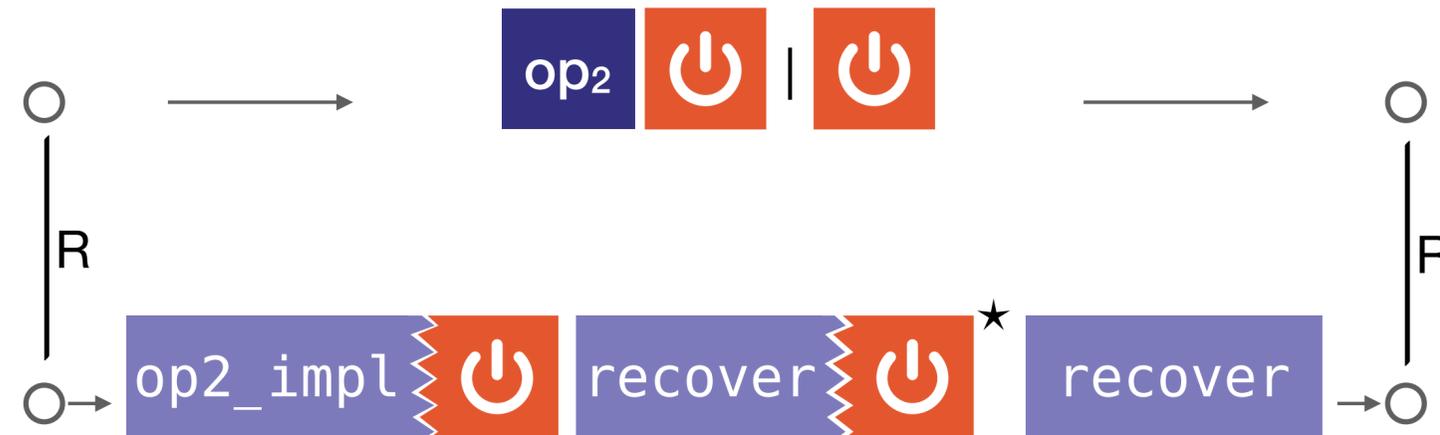
Proving trace inclusion, with recovery



Proving trace inclusion, with recovery

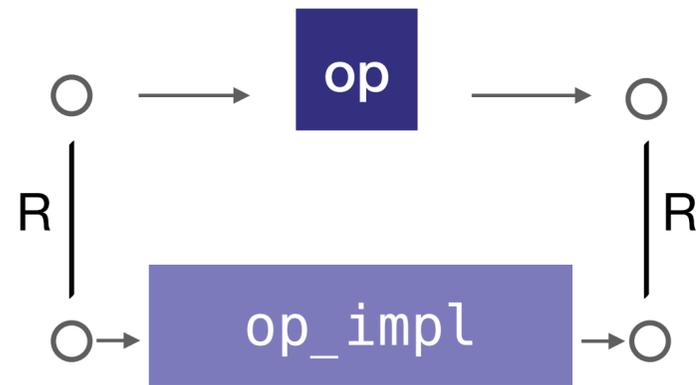


Proving trace inclusion, with recovery



Recovery refinement

non-crash execution

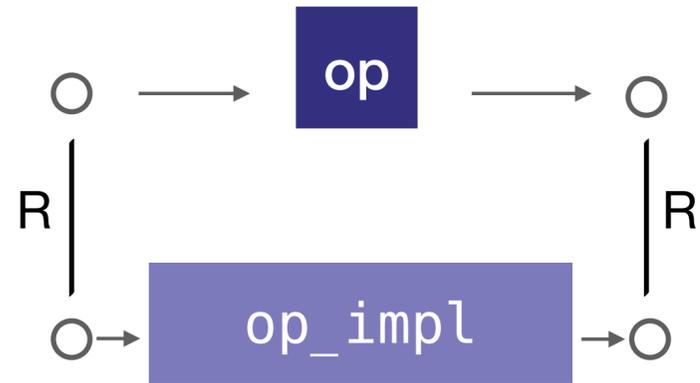


crash and recovery execution



Recovery refinement

non-crash execution



crash and recovery execution



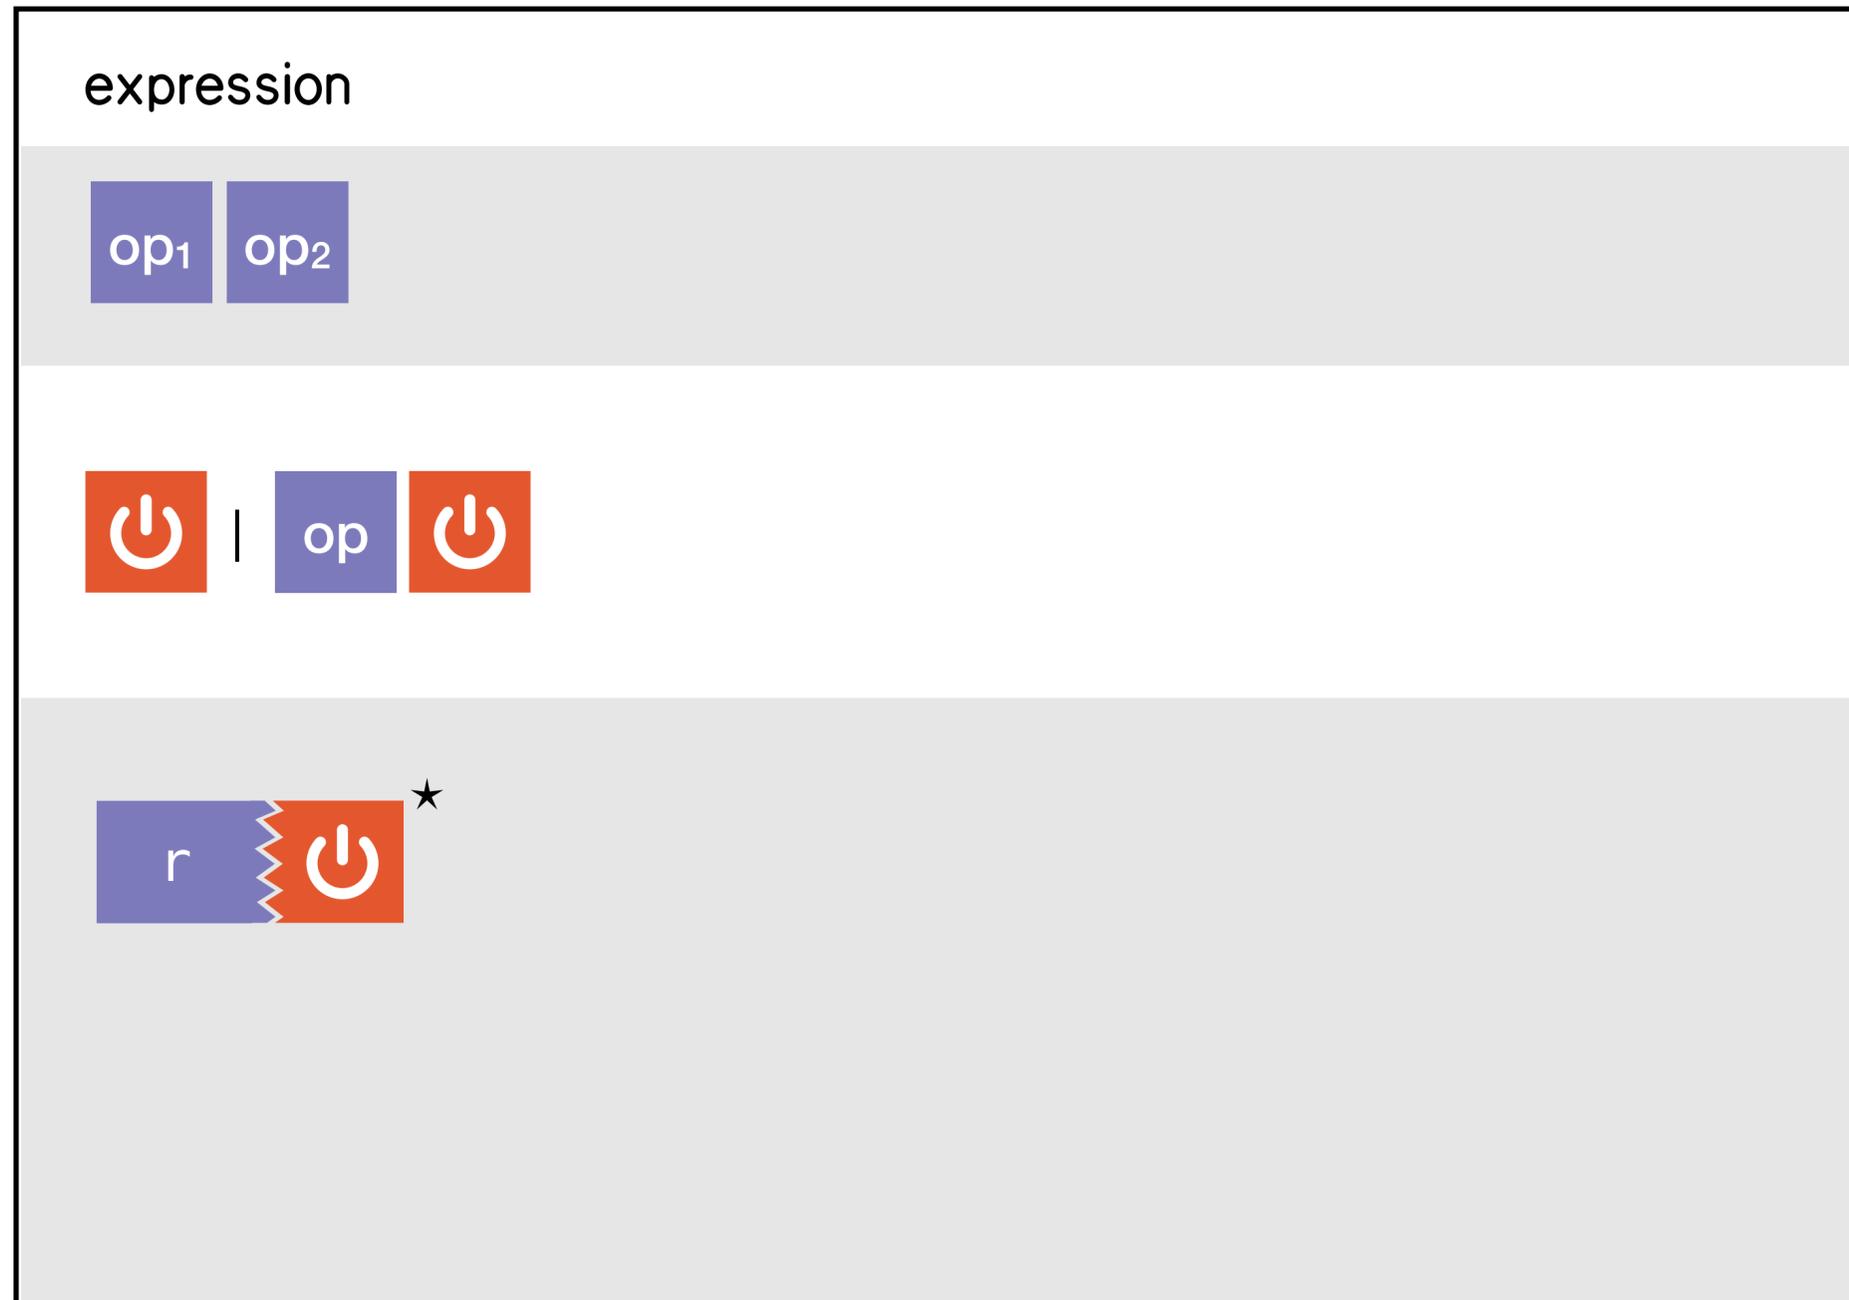
implies

Trace inclusion

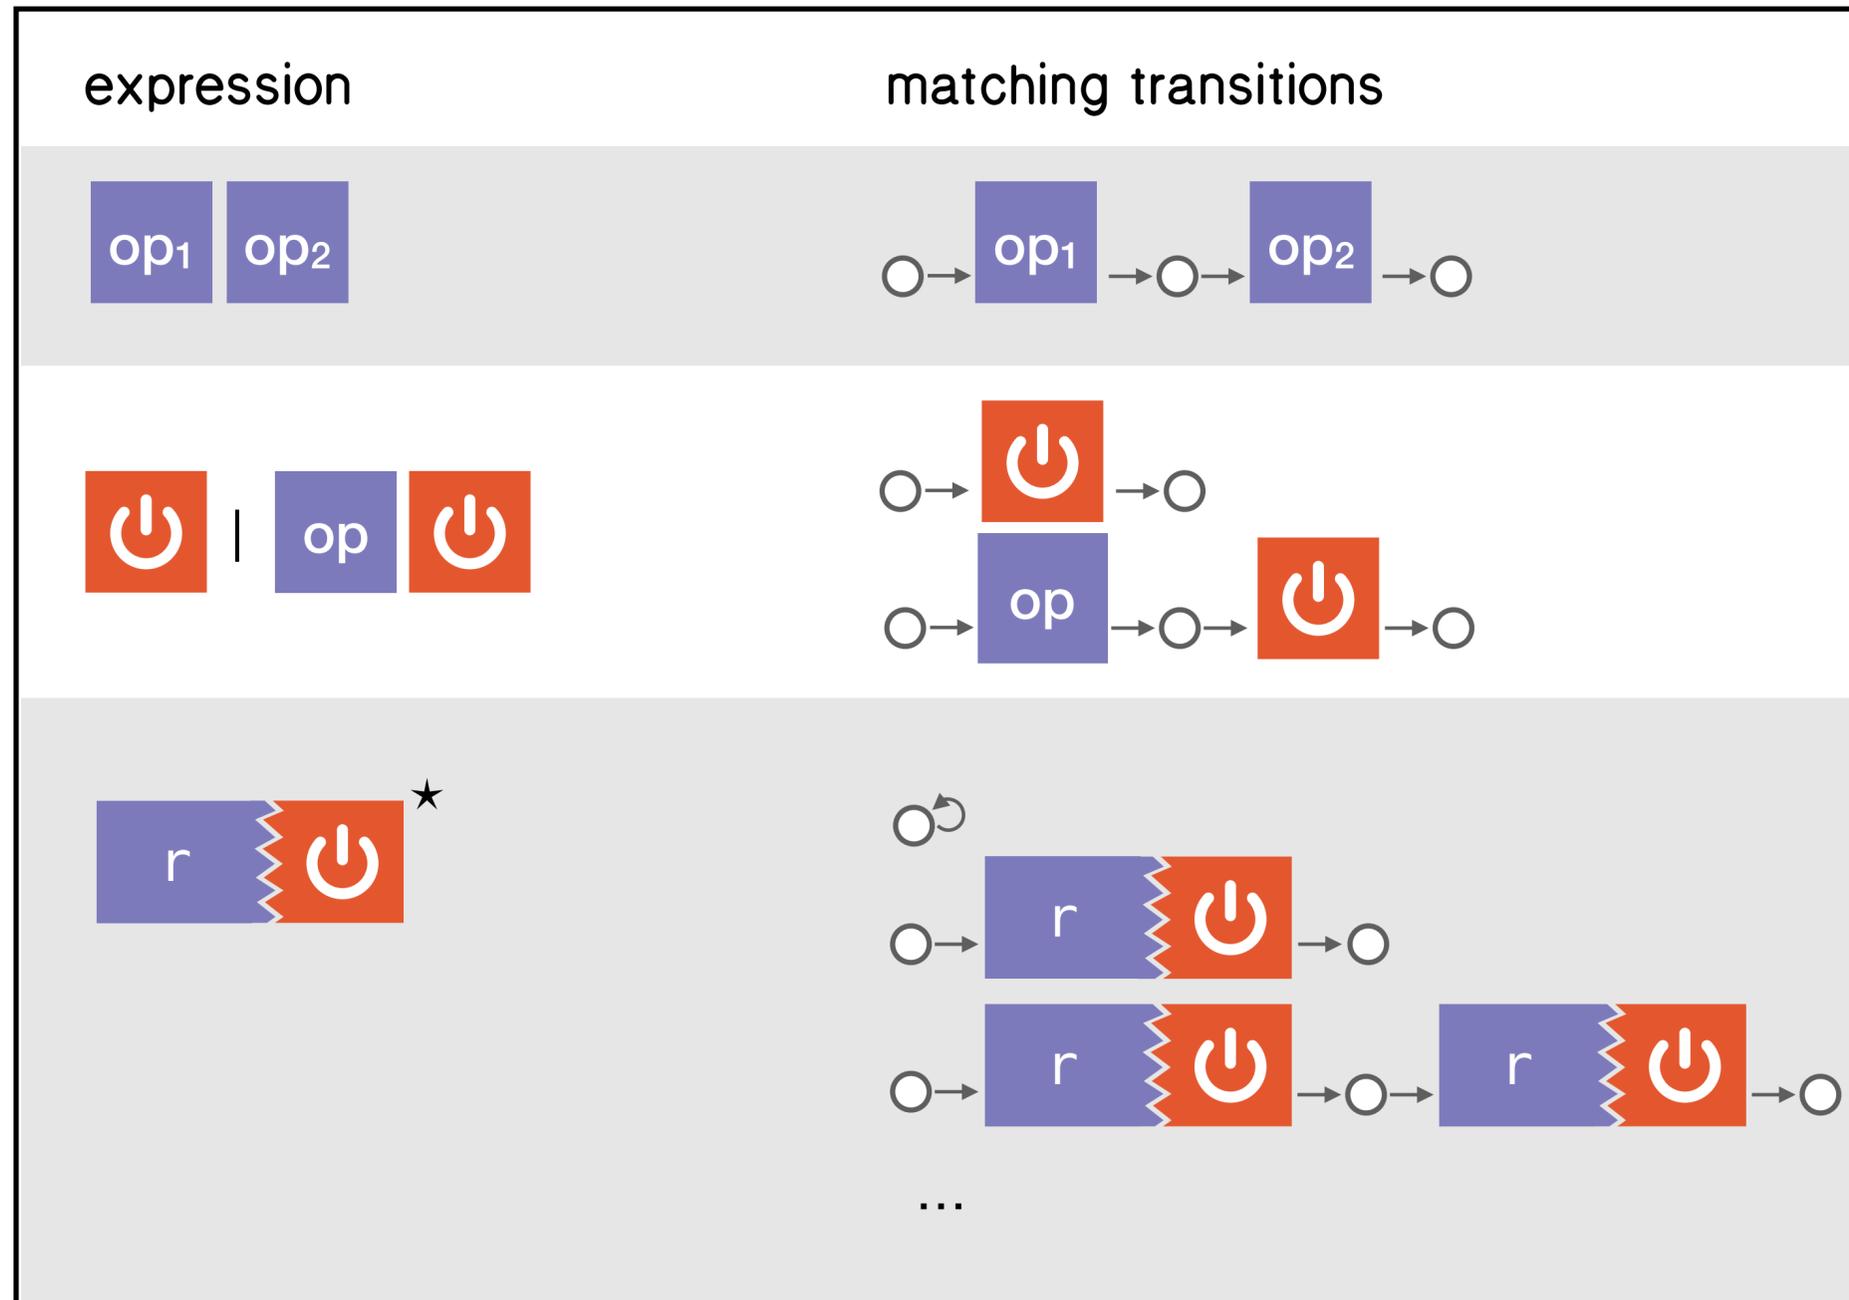
specification behavior
 \supseteq
running code behavior

Composition theorem

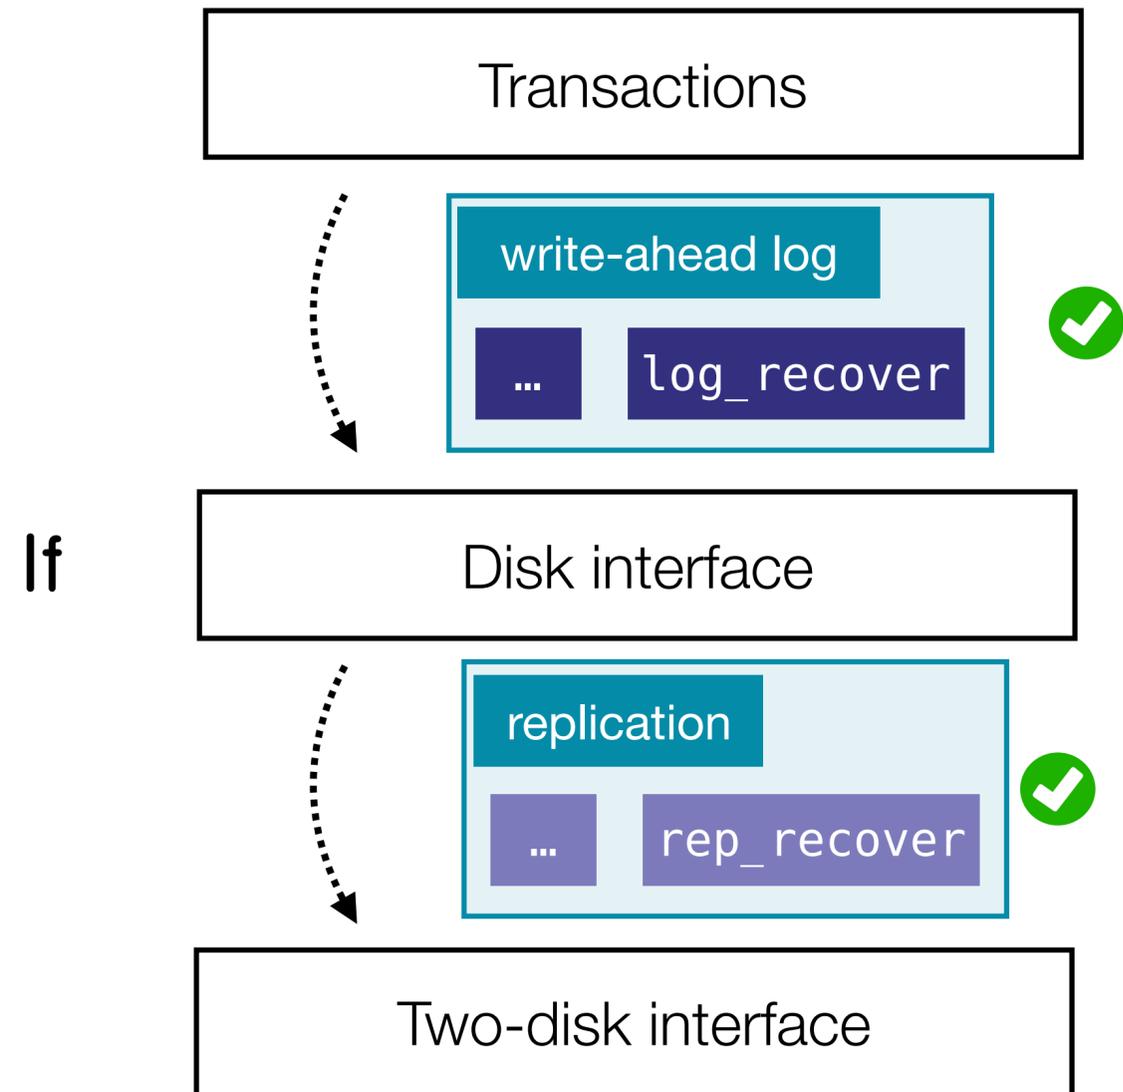
Kleene algebra for transition relations



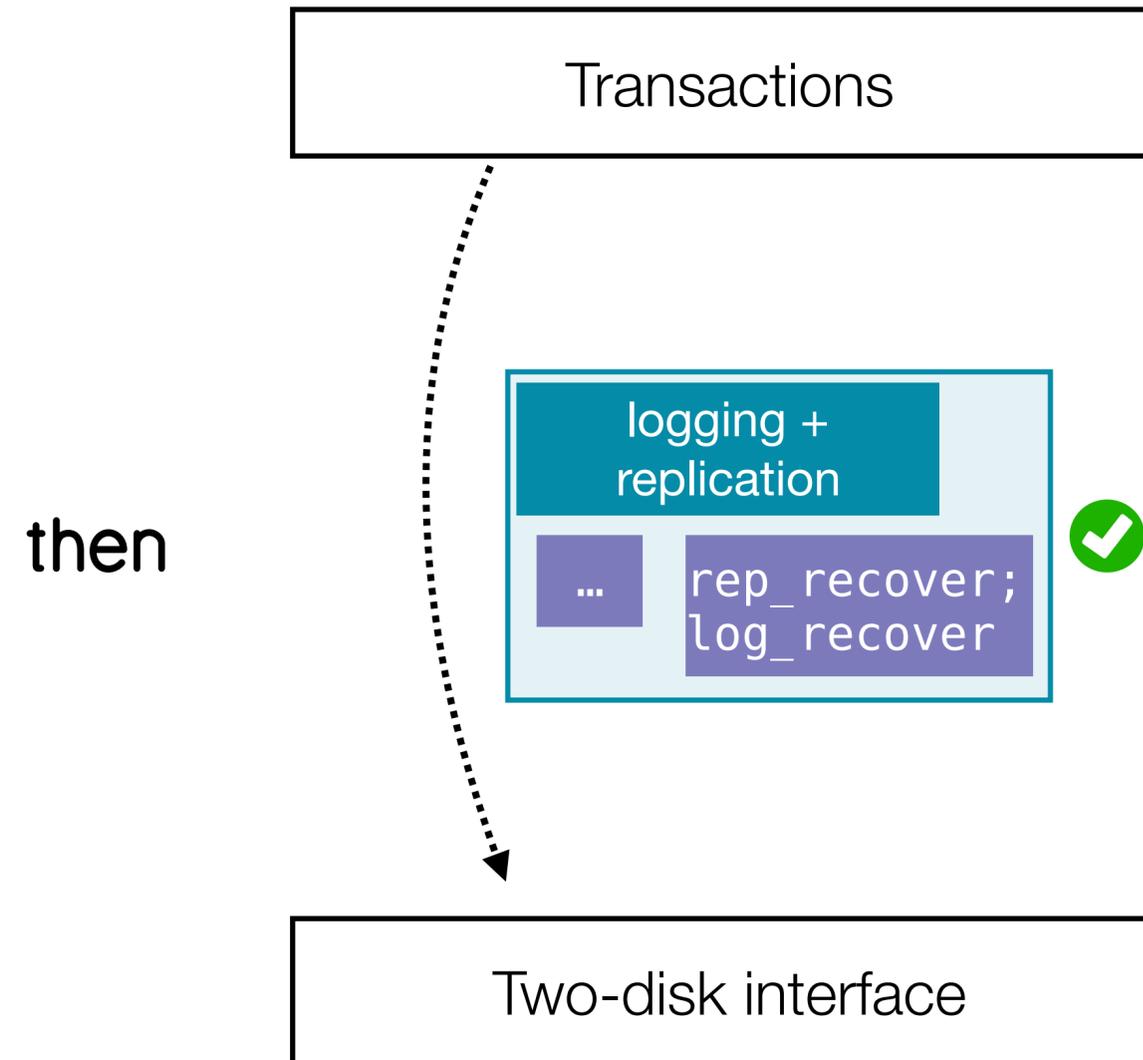
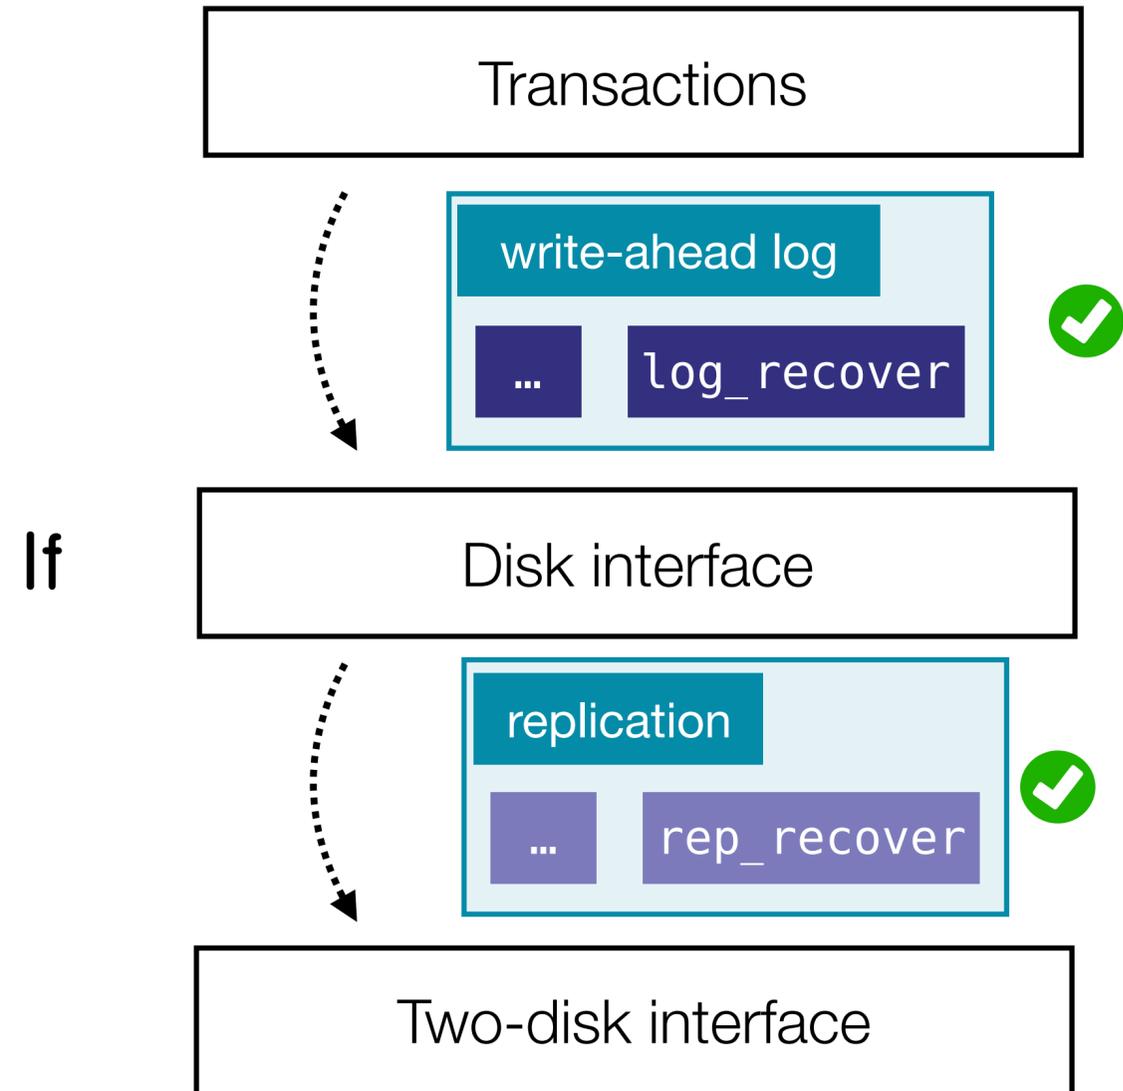
Kleene algebra for transition relations



Theorem: recovery refinements compose



Theorem: recovery refinements compose



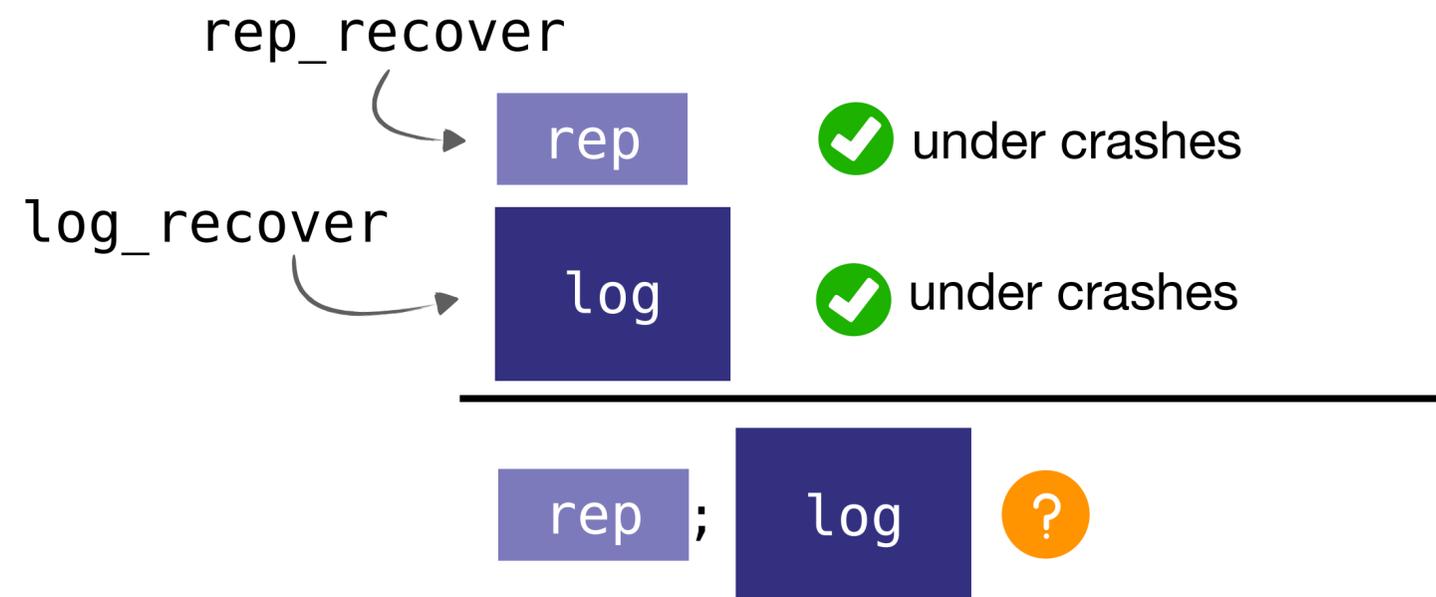
Goal: prove composed recovery correct

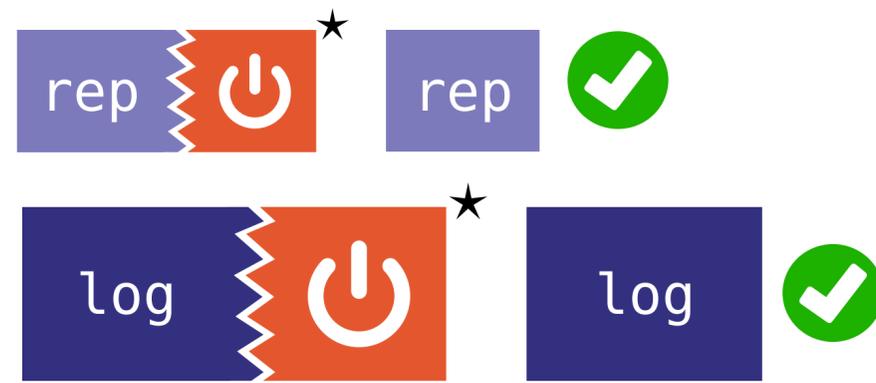
rep_recover ✓ under crashes

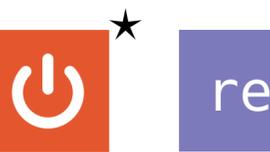
log_recover ✓ under crashes

rep_recover ; log_recover ?

Goal: prove composed recovery correct





rep * rep 

log * log 

(rep  | rep log )* rep log



how to re-use recovery proofs here?

Using Kleene algebra for reasoning

$$\left(\text{rep} \text{ } \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \text{ } \mid \text{rep} \text{ } \begin{array}{|c|} \hline \text{log} \\ \hline \end{array} \text{ } \begin{array}{|c|} \hline \text{ } \\ \hline \end{array} \right)^* \text{rep} \text{ } \begin{array}{|c|} \hline \text{log} \\ \hline \end{array}$$

Using Kleene algebra for reasoning



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

Using Kleene algebra for reasoning

$$\left(\text{rep} \text{ } \left| \text{rep} \text{ log} \right. \right)^* \text{rep} \text{ log}$$

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

$$= \text{rep} \left| \text{rep} \text{ log} \text{ rep} \right|^* \text{rep} \text{ log}$$

Using Kleene algebra for reasoning

$$\left(\text{rep} \text{ } \begin{array}{|c} \text{ } \\ \text{ } \end{array} \text{ } \text{ } \mid \text{rep} \text{ } \begin{array}{|c} \text{log} \\ \text{ } \end{array} \text{ } \text{ } \right)^* \text{rep} \text{ } \begin{array}{|c} \text{log} \\ \text{ } \end{array}$$

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

$$= \text{rep} \text{ } \begin{array}{|c} \text{ } \\ \text{ } \end{array} \text{ } \text{ }^* \left(\text{rep} \text{ } \begin{array}{|c} \text{log} \\ \text{ } \end{array} \text{ } \text{ } \begin{array}{|c} \text{rep} \text{ } \begin{array}{|c} \text{ } \\ \text{ } \end{array} \text{ } \text{ }^* \end{array} \right)^* \text{rep} \text{ } \begin{array}{|c} \text{log} \\ \text{ } \end{array}$$

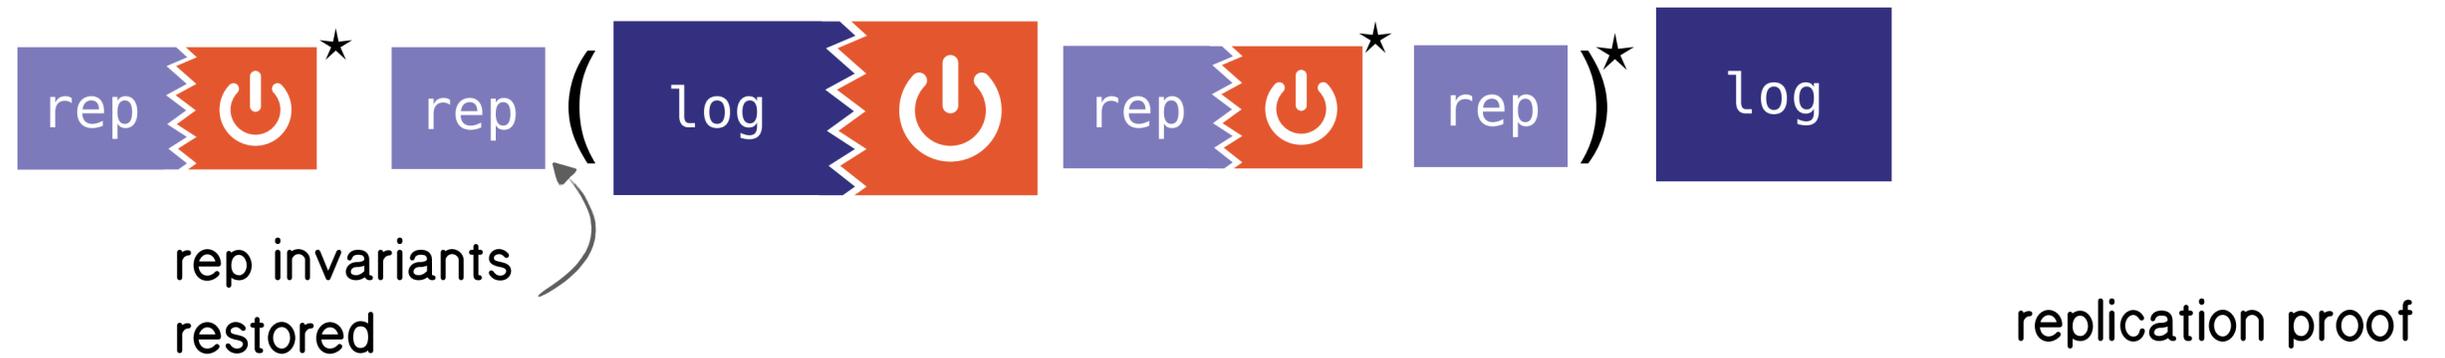
after sliding $(pq)^*p = p(qp)^*$

$$= \text{rep} \text{ } \begin{array}{|c} \text{ } \\ \text{ } \end{array} \text{ } \text{ }^* \text{rep} \left(\begin{array}{|c} \text{log} \\ \text{ } \end{array} \text{ } \text{ } \begin{array}{|c} \text{rep} \text{ } \begin{array}{|c} \text{ } \\ \text{ } \end{array} \text{ } \text{ }^* \end{array} \text{rep} \right)^* \text{rep} \text{ } \begin{array}{|c} \text{log} \\ \text{ } \end{array}$$

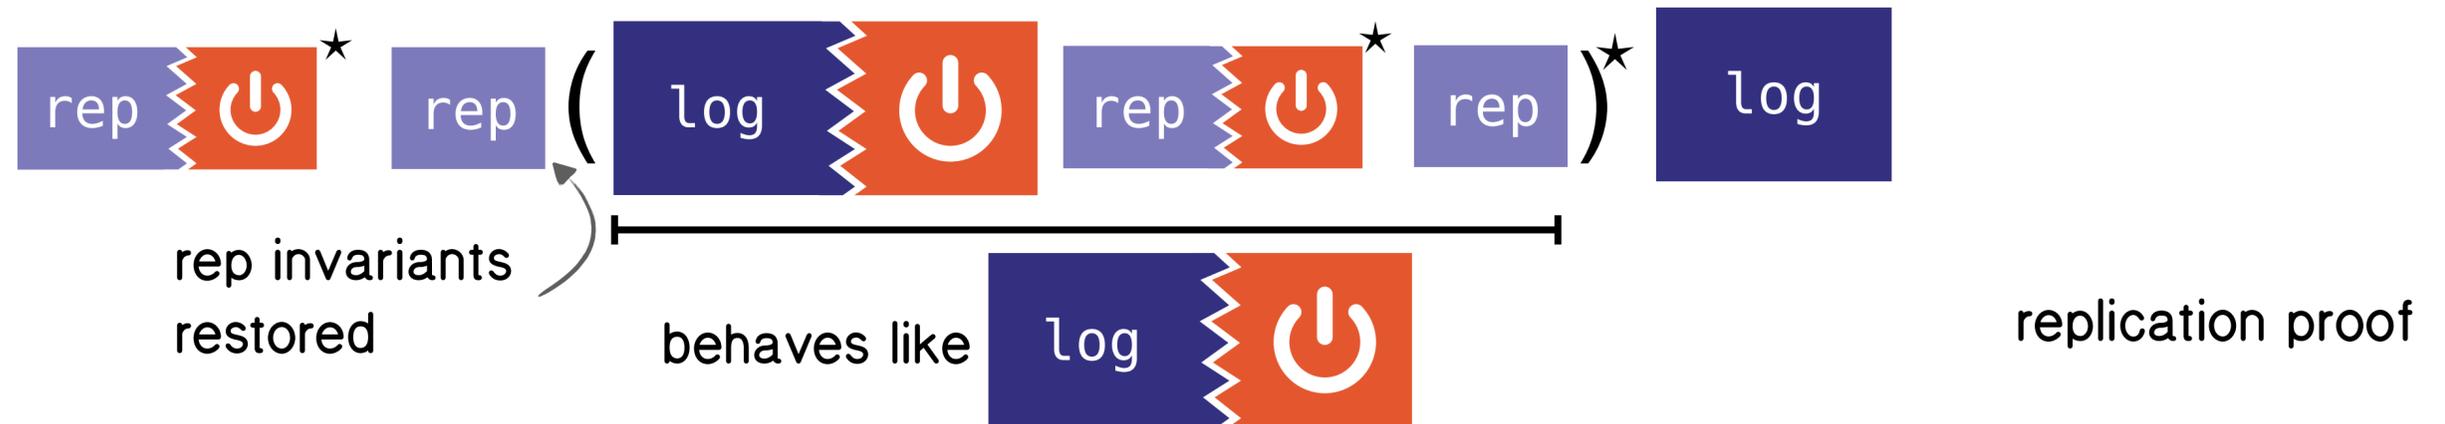
After rewrite both proofs apply



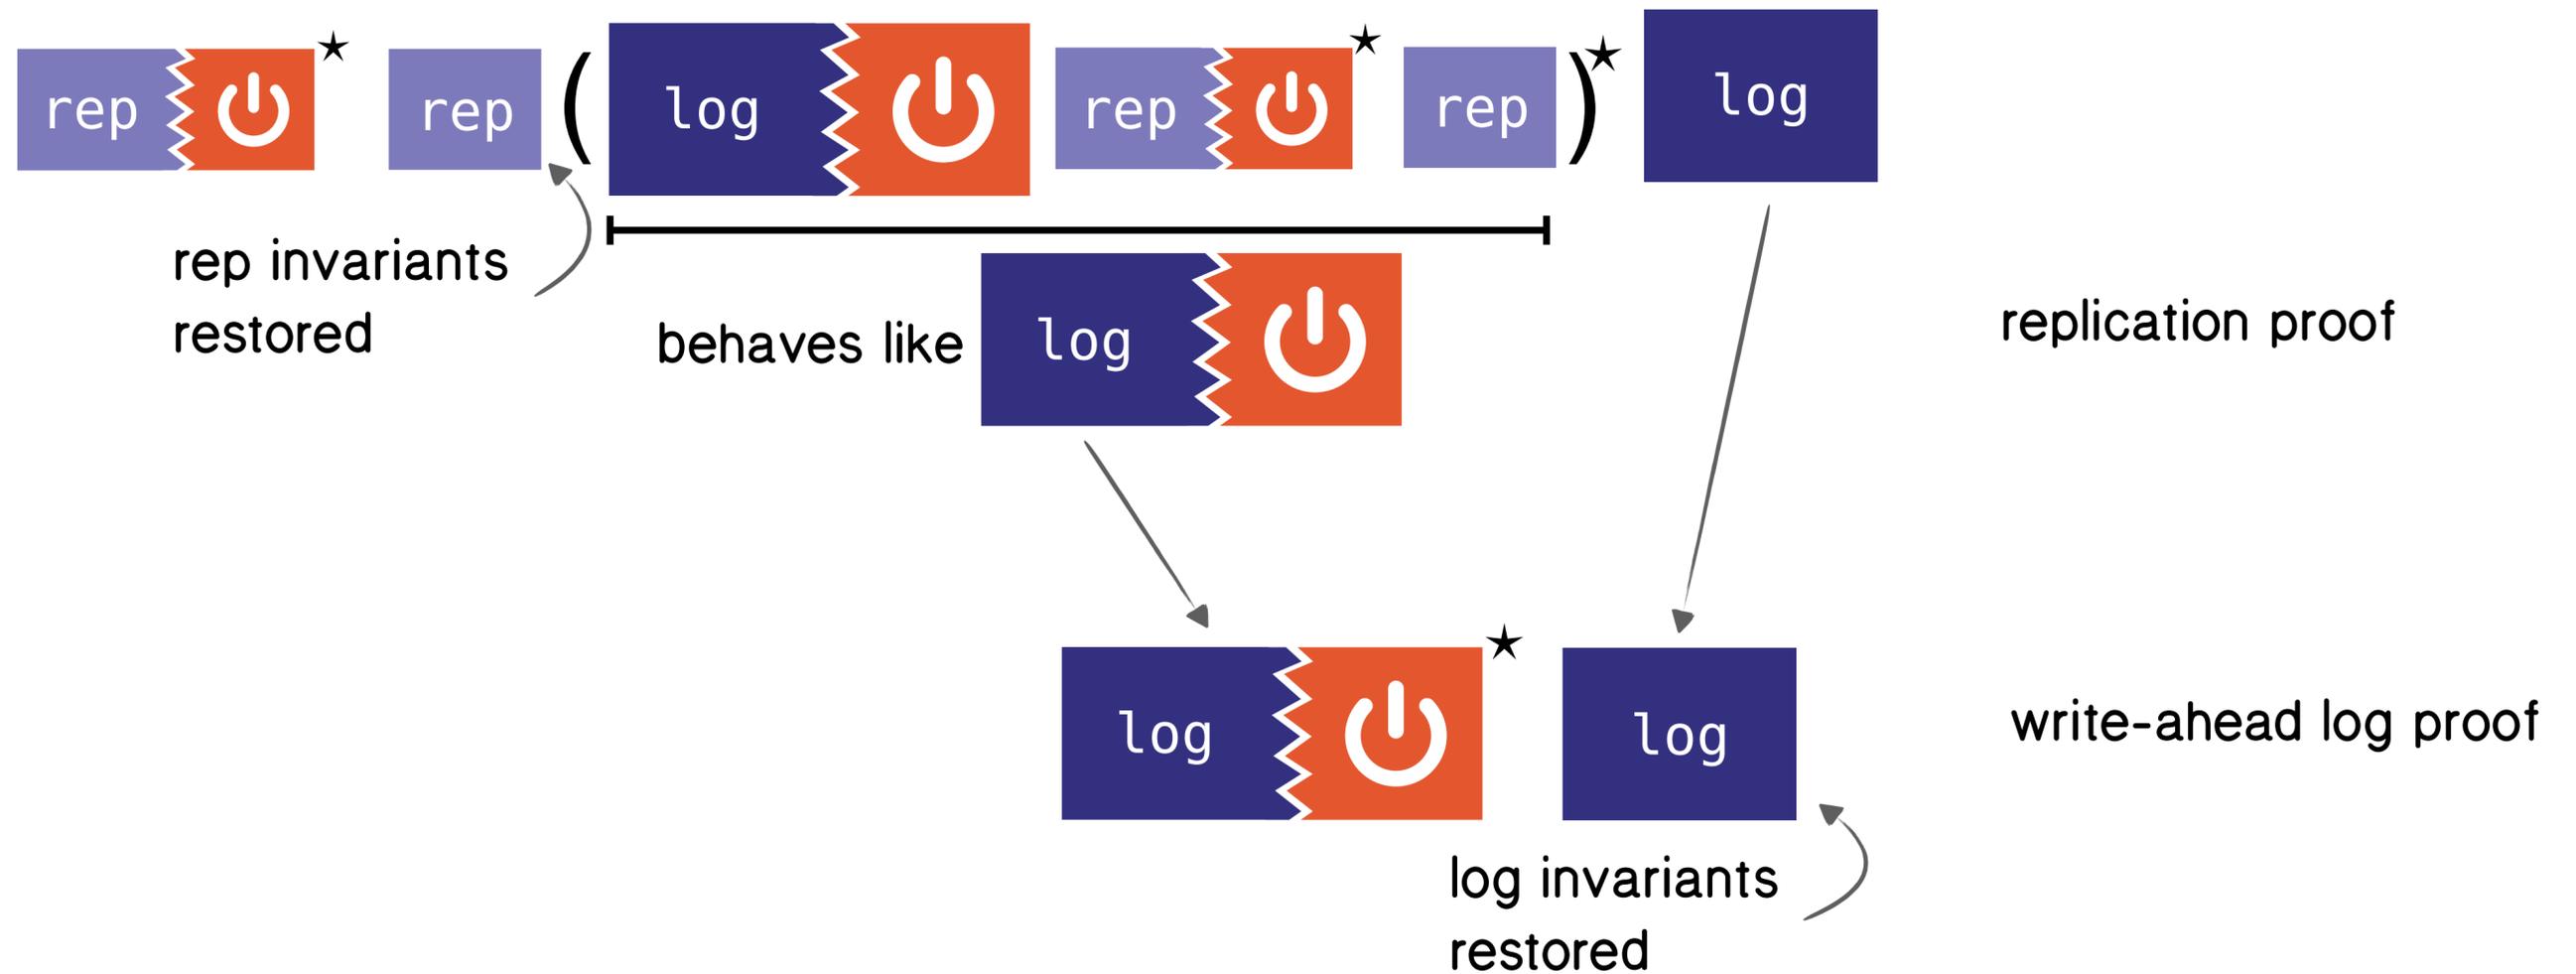
After rewrite both proofs apply



After rewrite both proofs apply



After rewrite both proofs apply



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

Argosy: modular proofs of layered storage systems



Argosy: modular proofs of layered storage systems

Kleene algebra

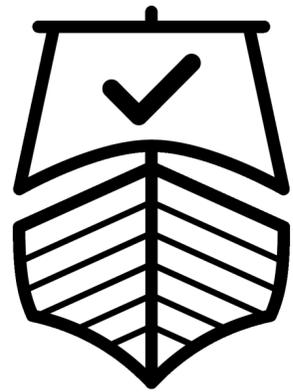
$$\left(\text{rep} \begin{array}{|c|} \hline \text{⚡} \\ \hline \end{array} \mid \text{rep} \begin{array}{|c|} \hline \text{log} \\ \hline \end{array} \begin{array}{|c|} \hline \text{⚡} \\ \hline \end{array} \right)^*$$



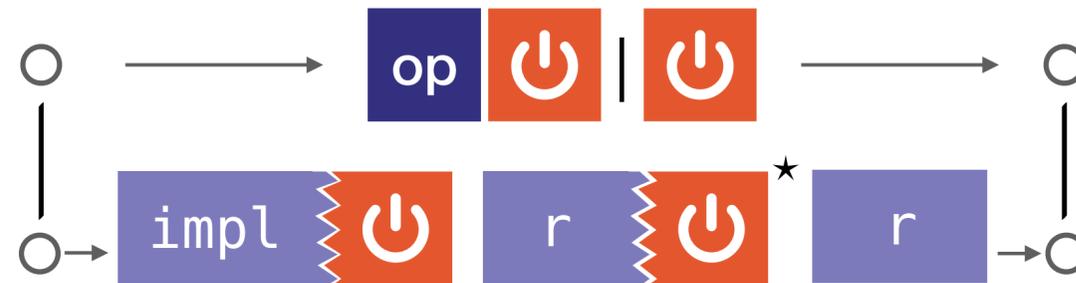
Argosy: modular proofs of layered storage systems

Kleene algebra

$$\left(\text{rep} \text{ } \left| \text{ } \right. \text{rep} \text{ } \text{log} \right)^*$$



recovery refinement



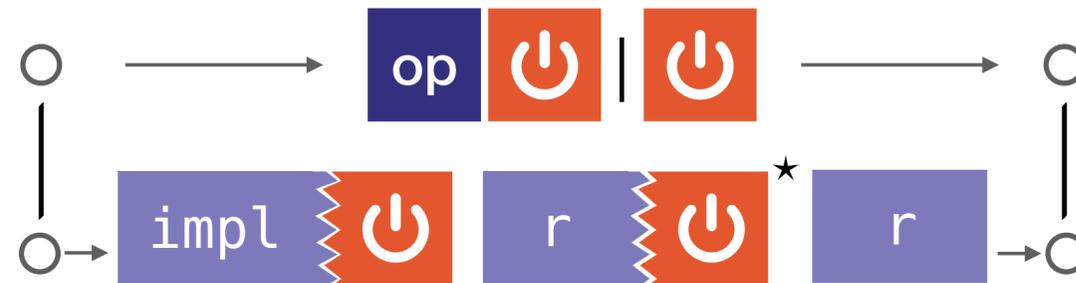
Argosy: modular proofs of layered storage systems



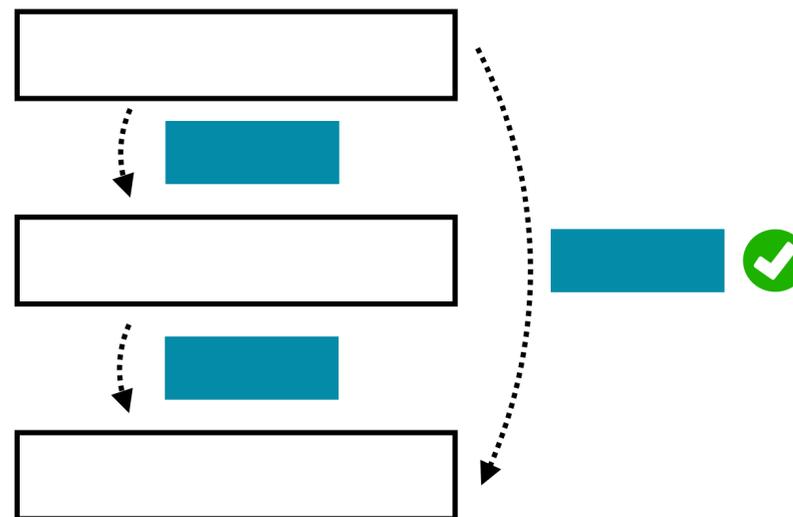
Kleene algebra

$$\left(\text{rep} \text{ } \left| \text{ } \text{rep} \text{ log} \right. \right)^*$$

recovery refinement



modular proofs



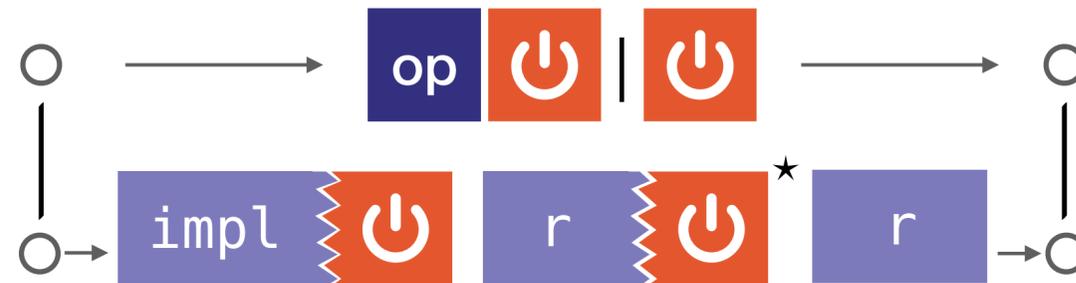
Argosy: modular proofs of layered storage systems

Kleene algebra

$$\left(\text{rep} \text{ } \left| \text{ } \text{rep} \text{ } \log \right. \right)^*$$



recovery refinement



come find us after!

Tej and Joe

modular proofs

