


# Overcoming Memory Weakness with Unified Fairness

---

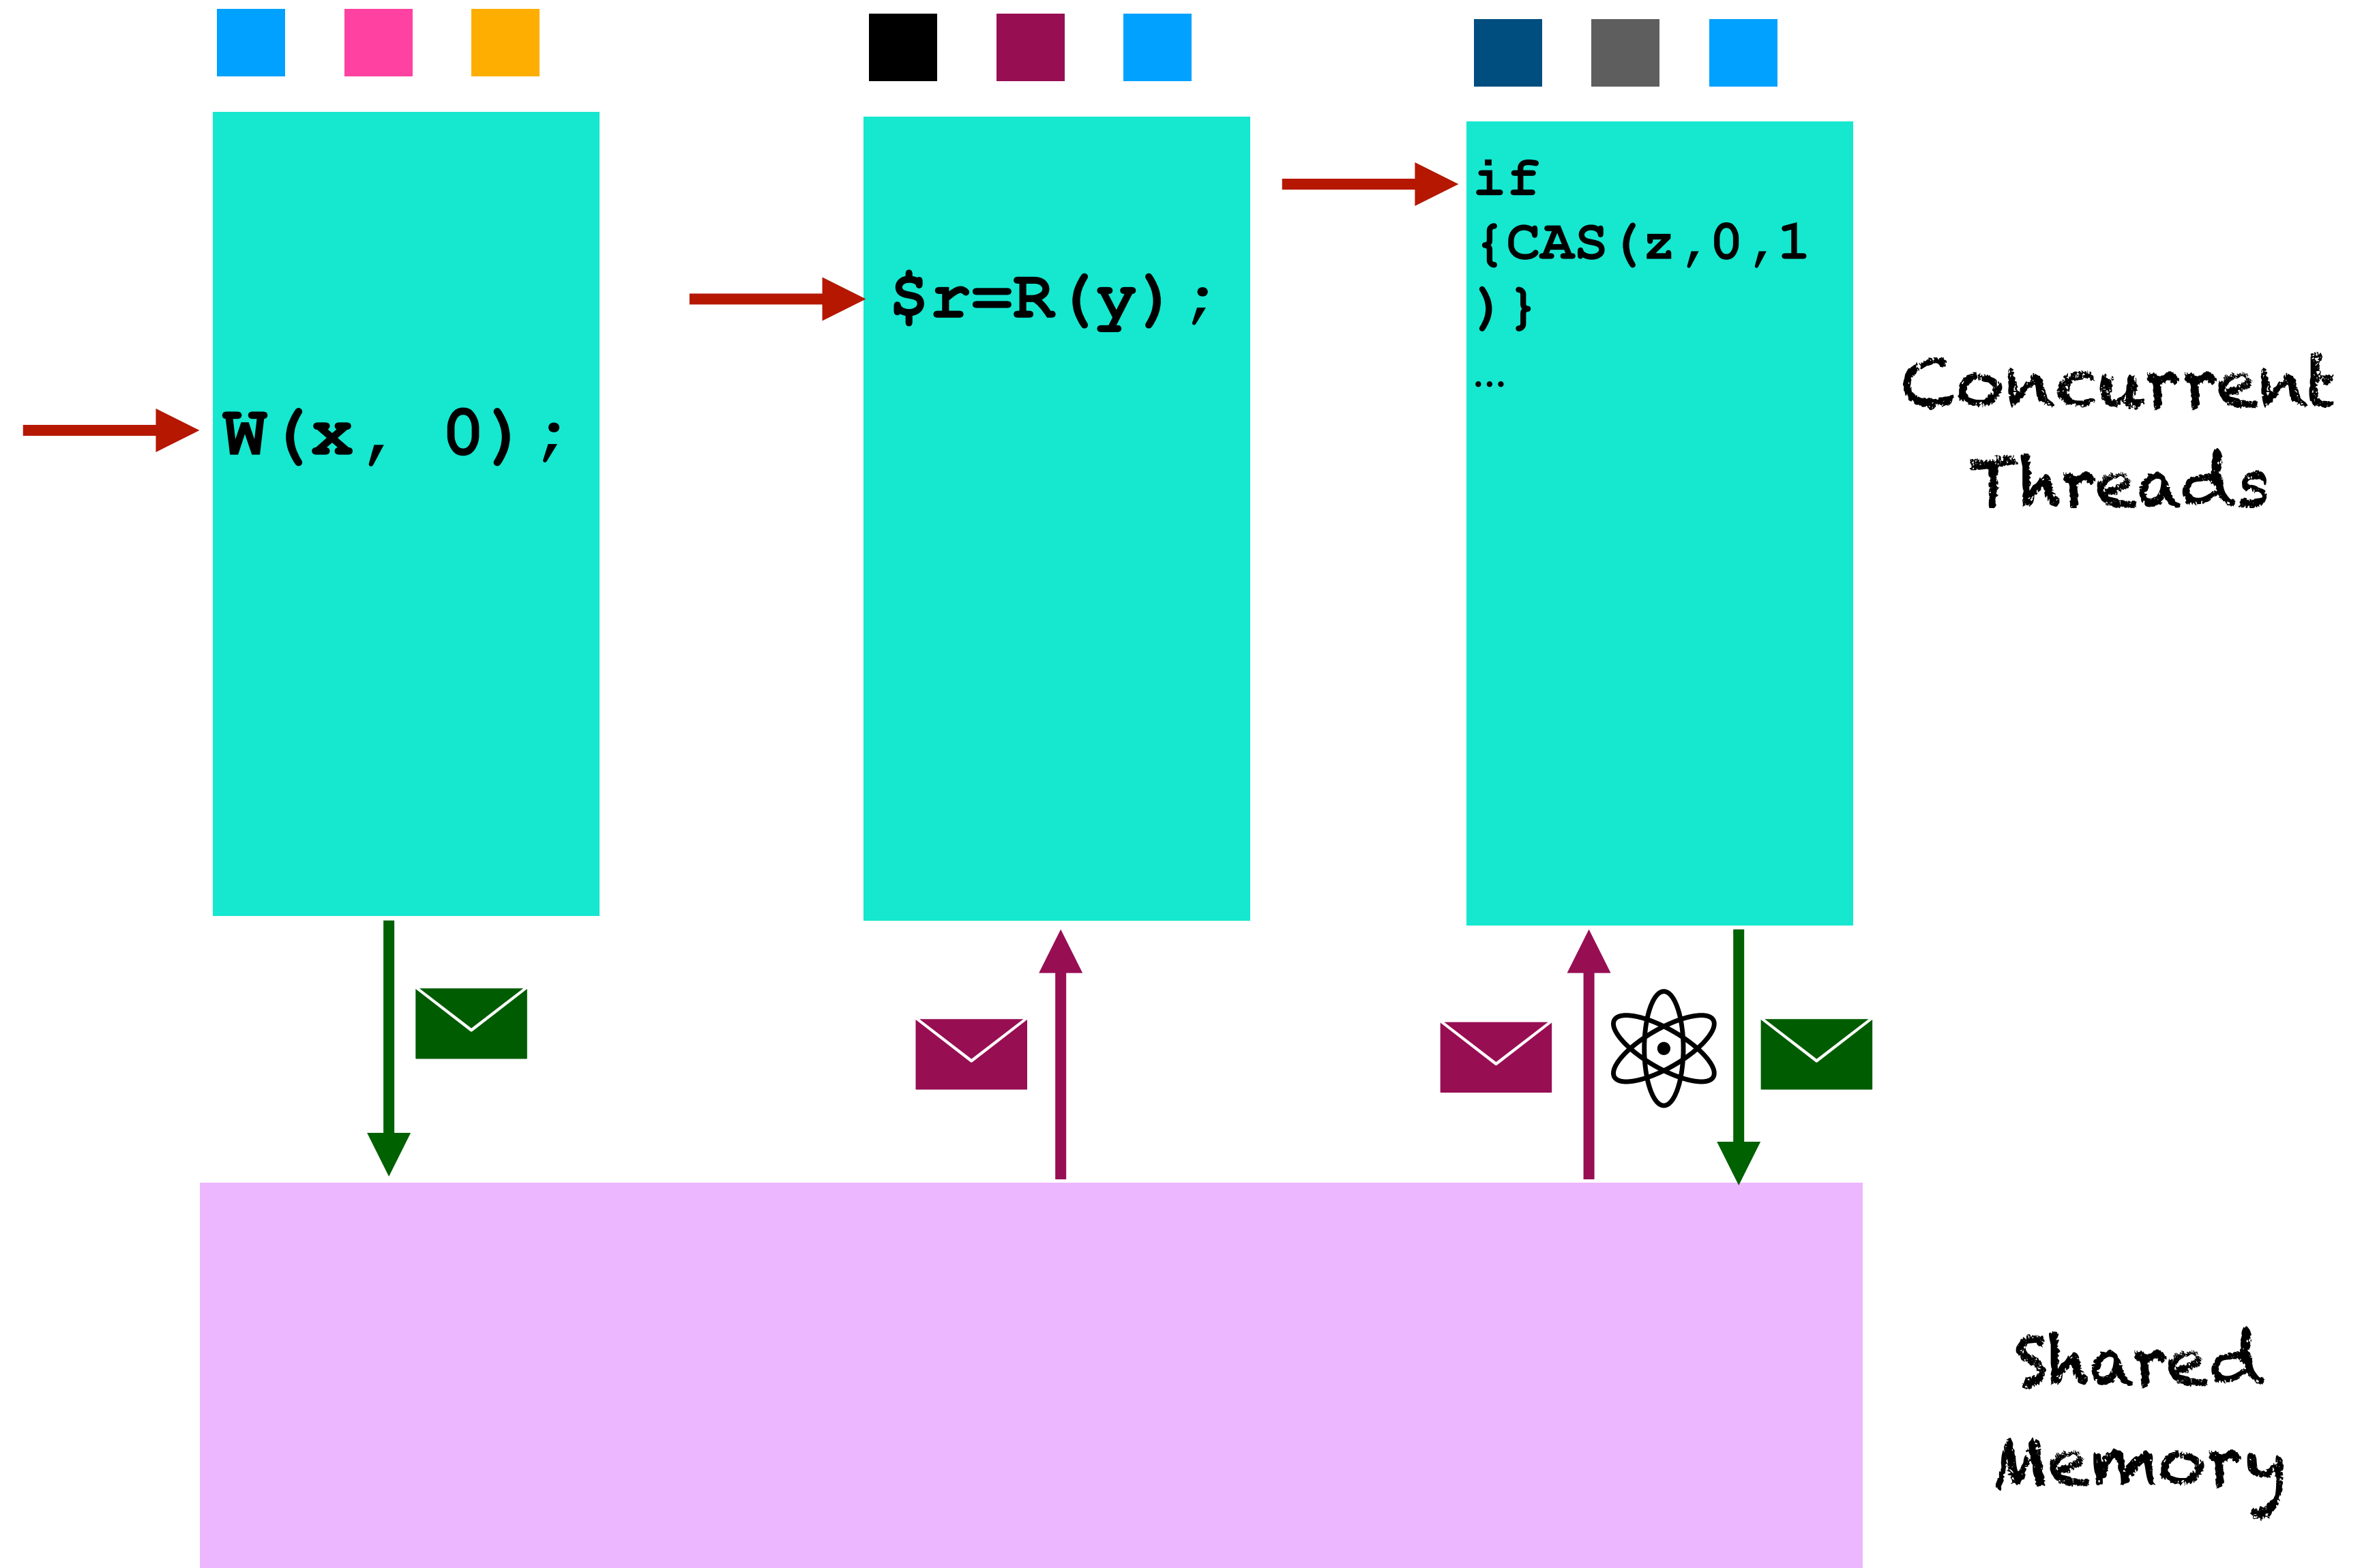
- Parosh Aziz Abdulla (Uppsala University)
- Mohamed Faouzi Atig (Uppsala University)
- Adwait Godbole (University of California, Berkeley)
- Shankaranarayanan Krishna (IIT Bombay)
- **Mihir Vahanwala (MPI-SWS, Saarbrücken)**




**RHPL 2023, Hyderabad**

**December 20**

# Concurrency and Memory: The Setup



# Weak memory: an abstract idea




# Weak Memory: The challenge


Demonic, impractical  
Non-determinism



??





Algorithms: all at sea  
Whole host of models




Liveness Verification:  
-Program termination  
-Repeated control state reachability

Memory Fairness,  
rooted in practicality





Unified framework




# Systematic verification of liveness in Weak Memory Models

Liveness, verified

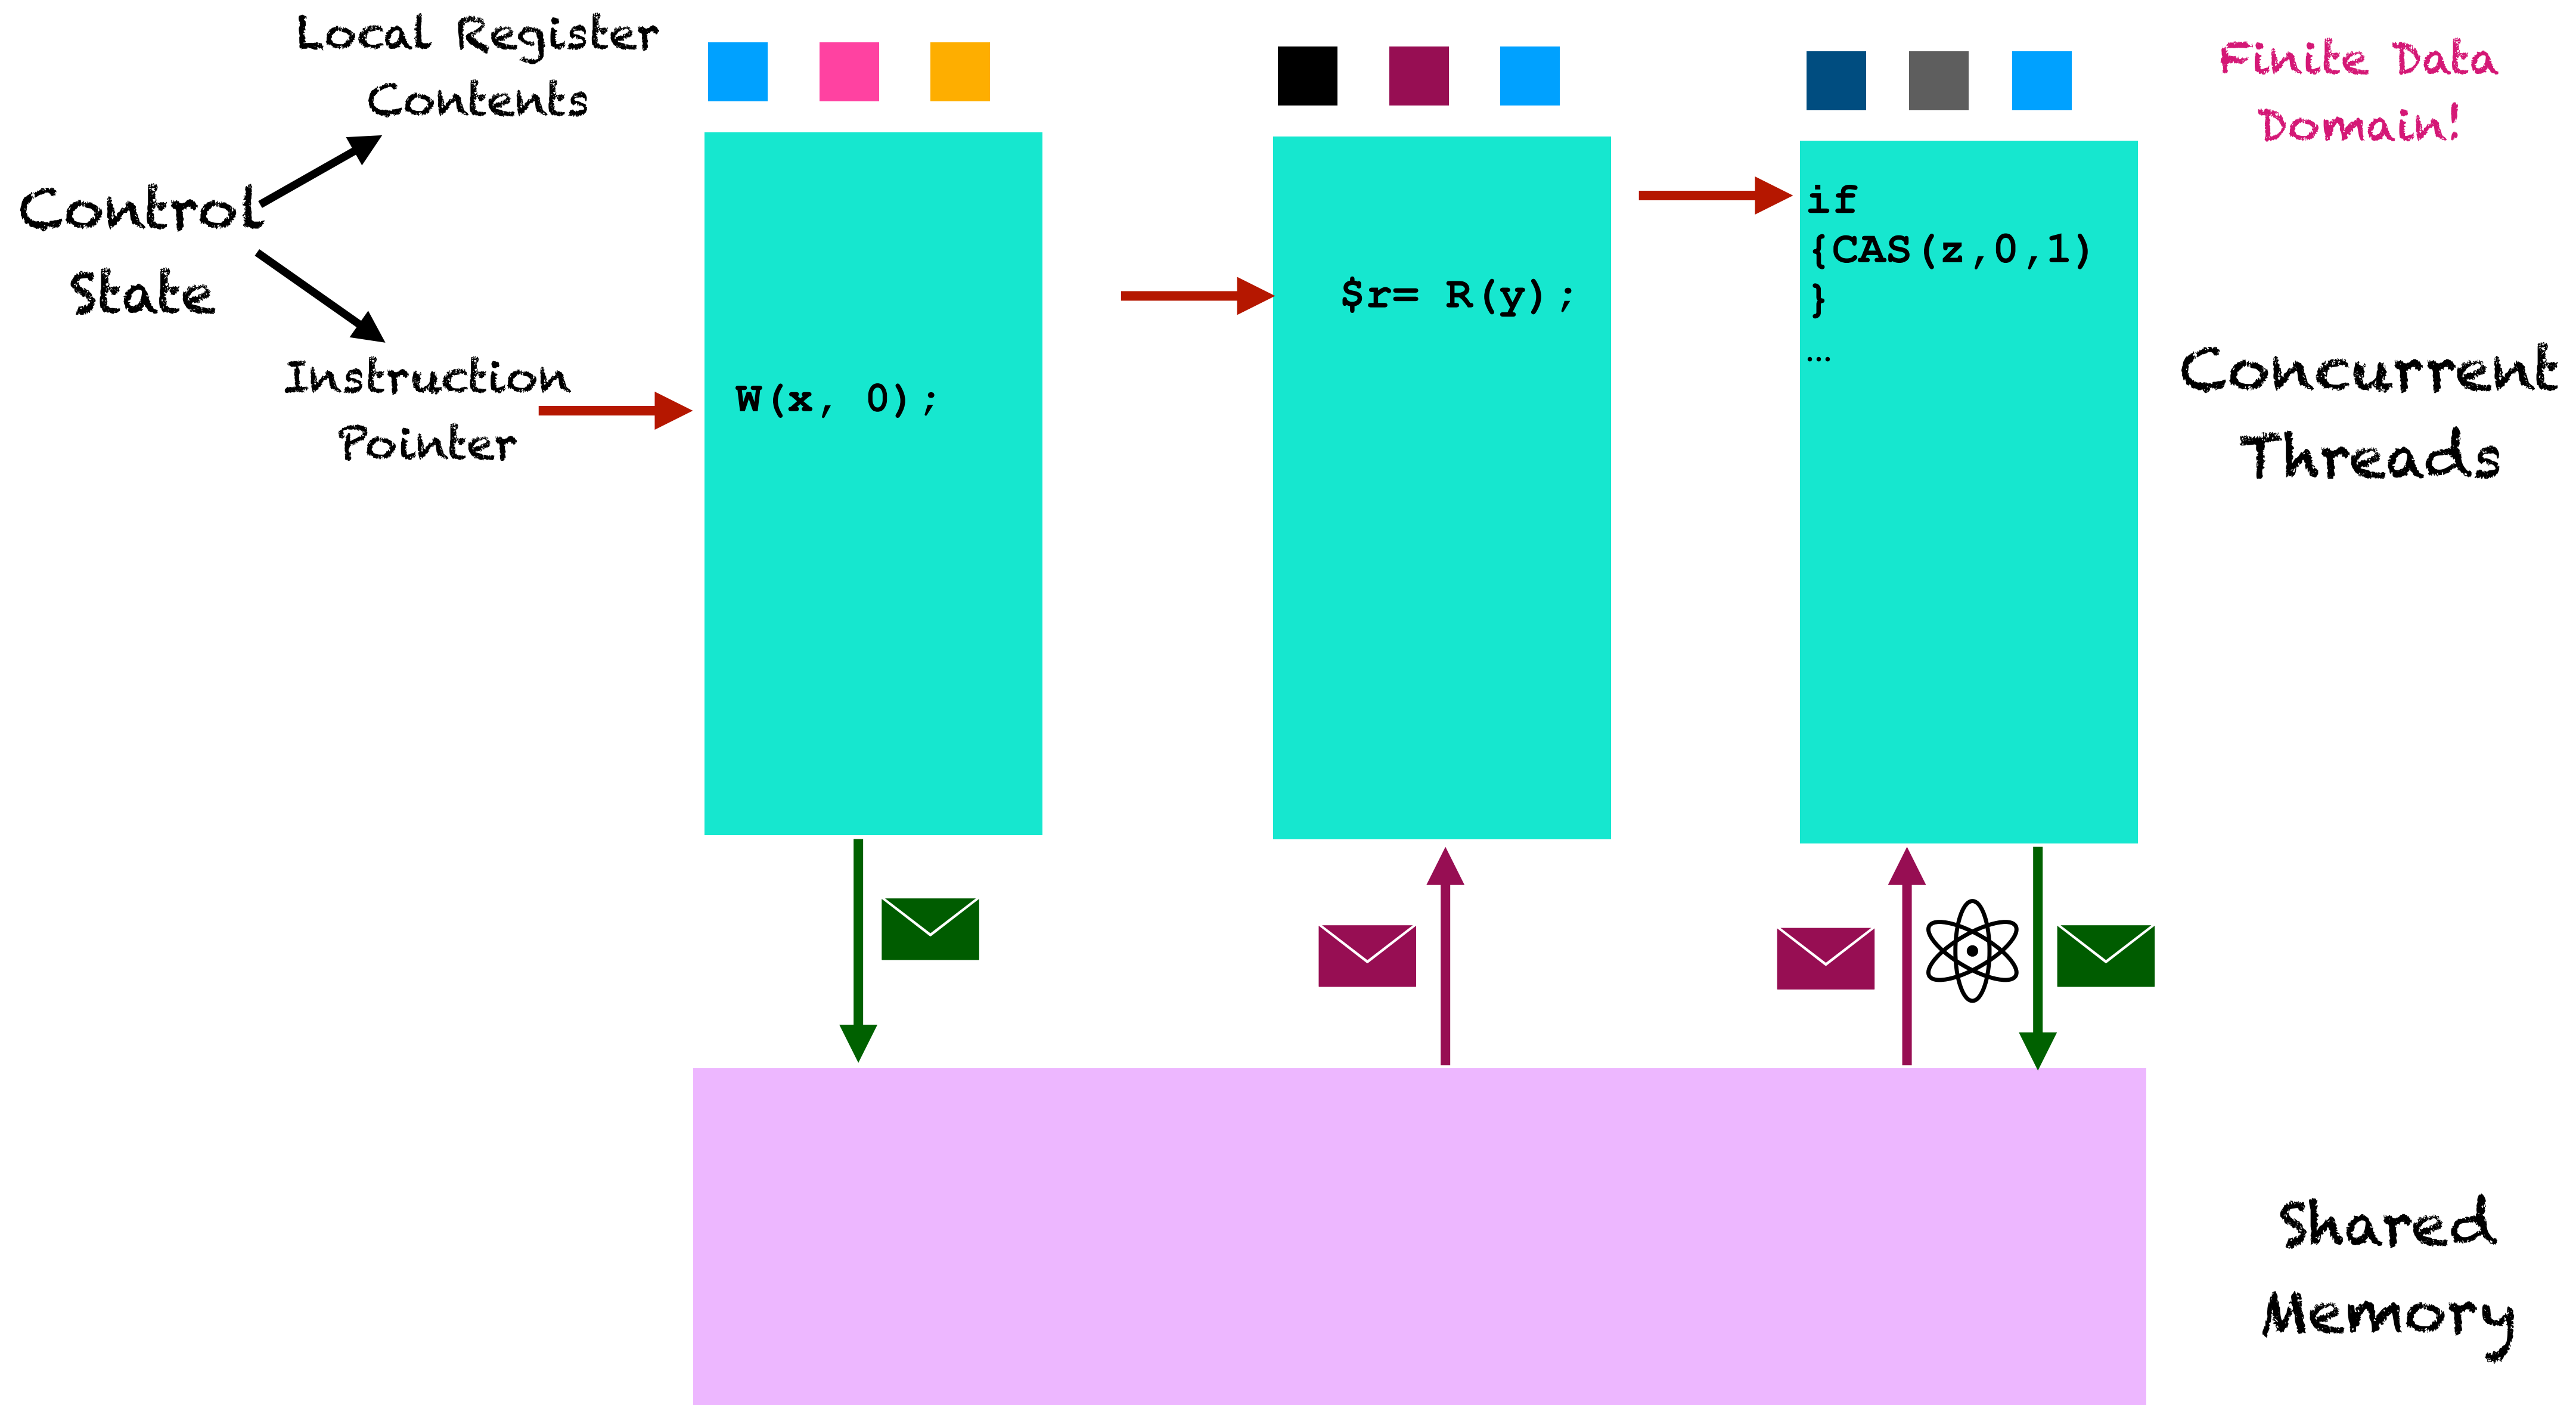


Established techniques

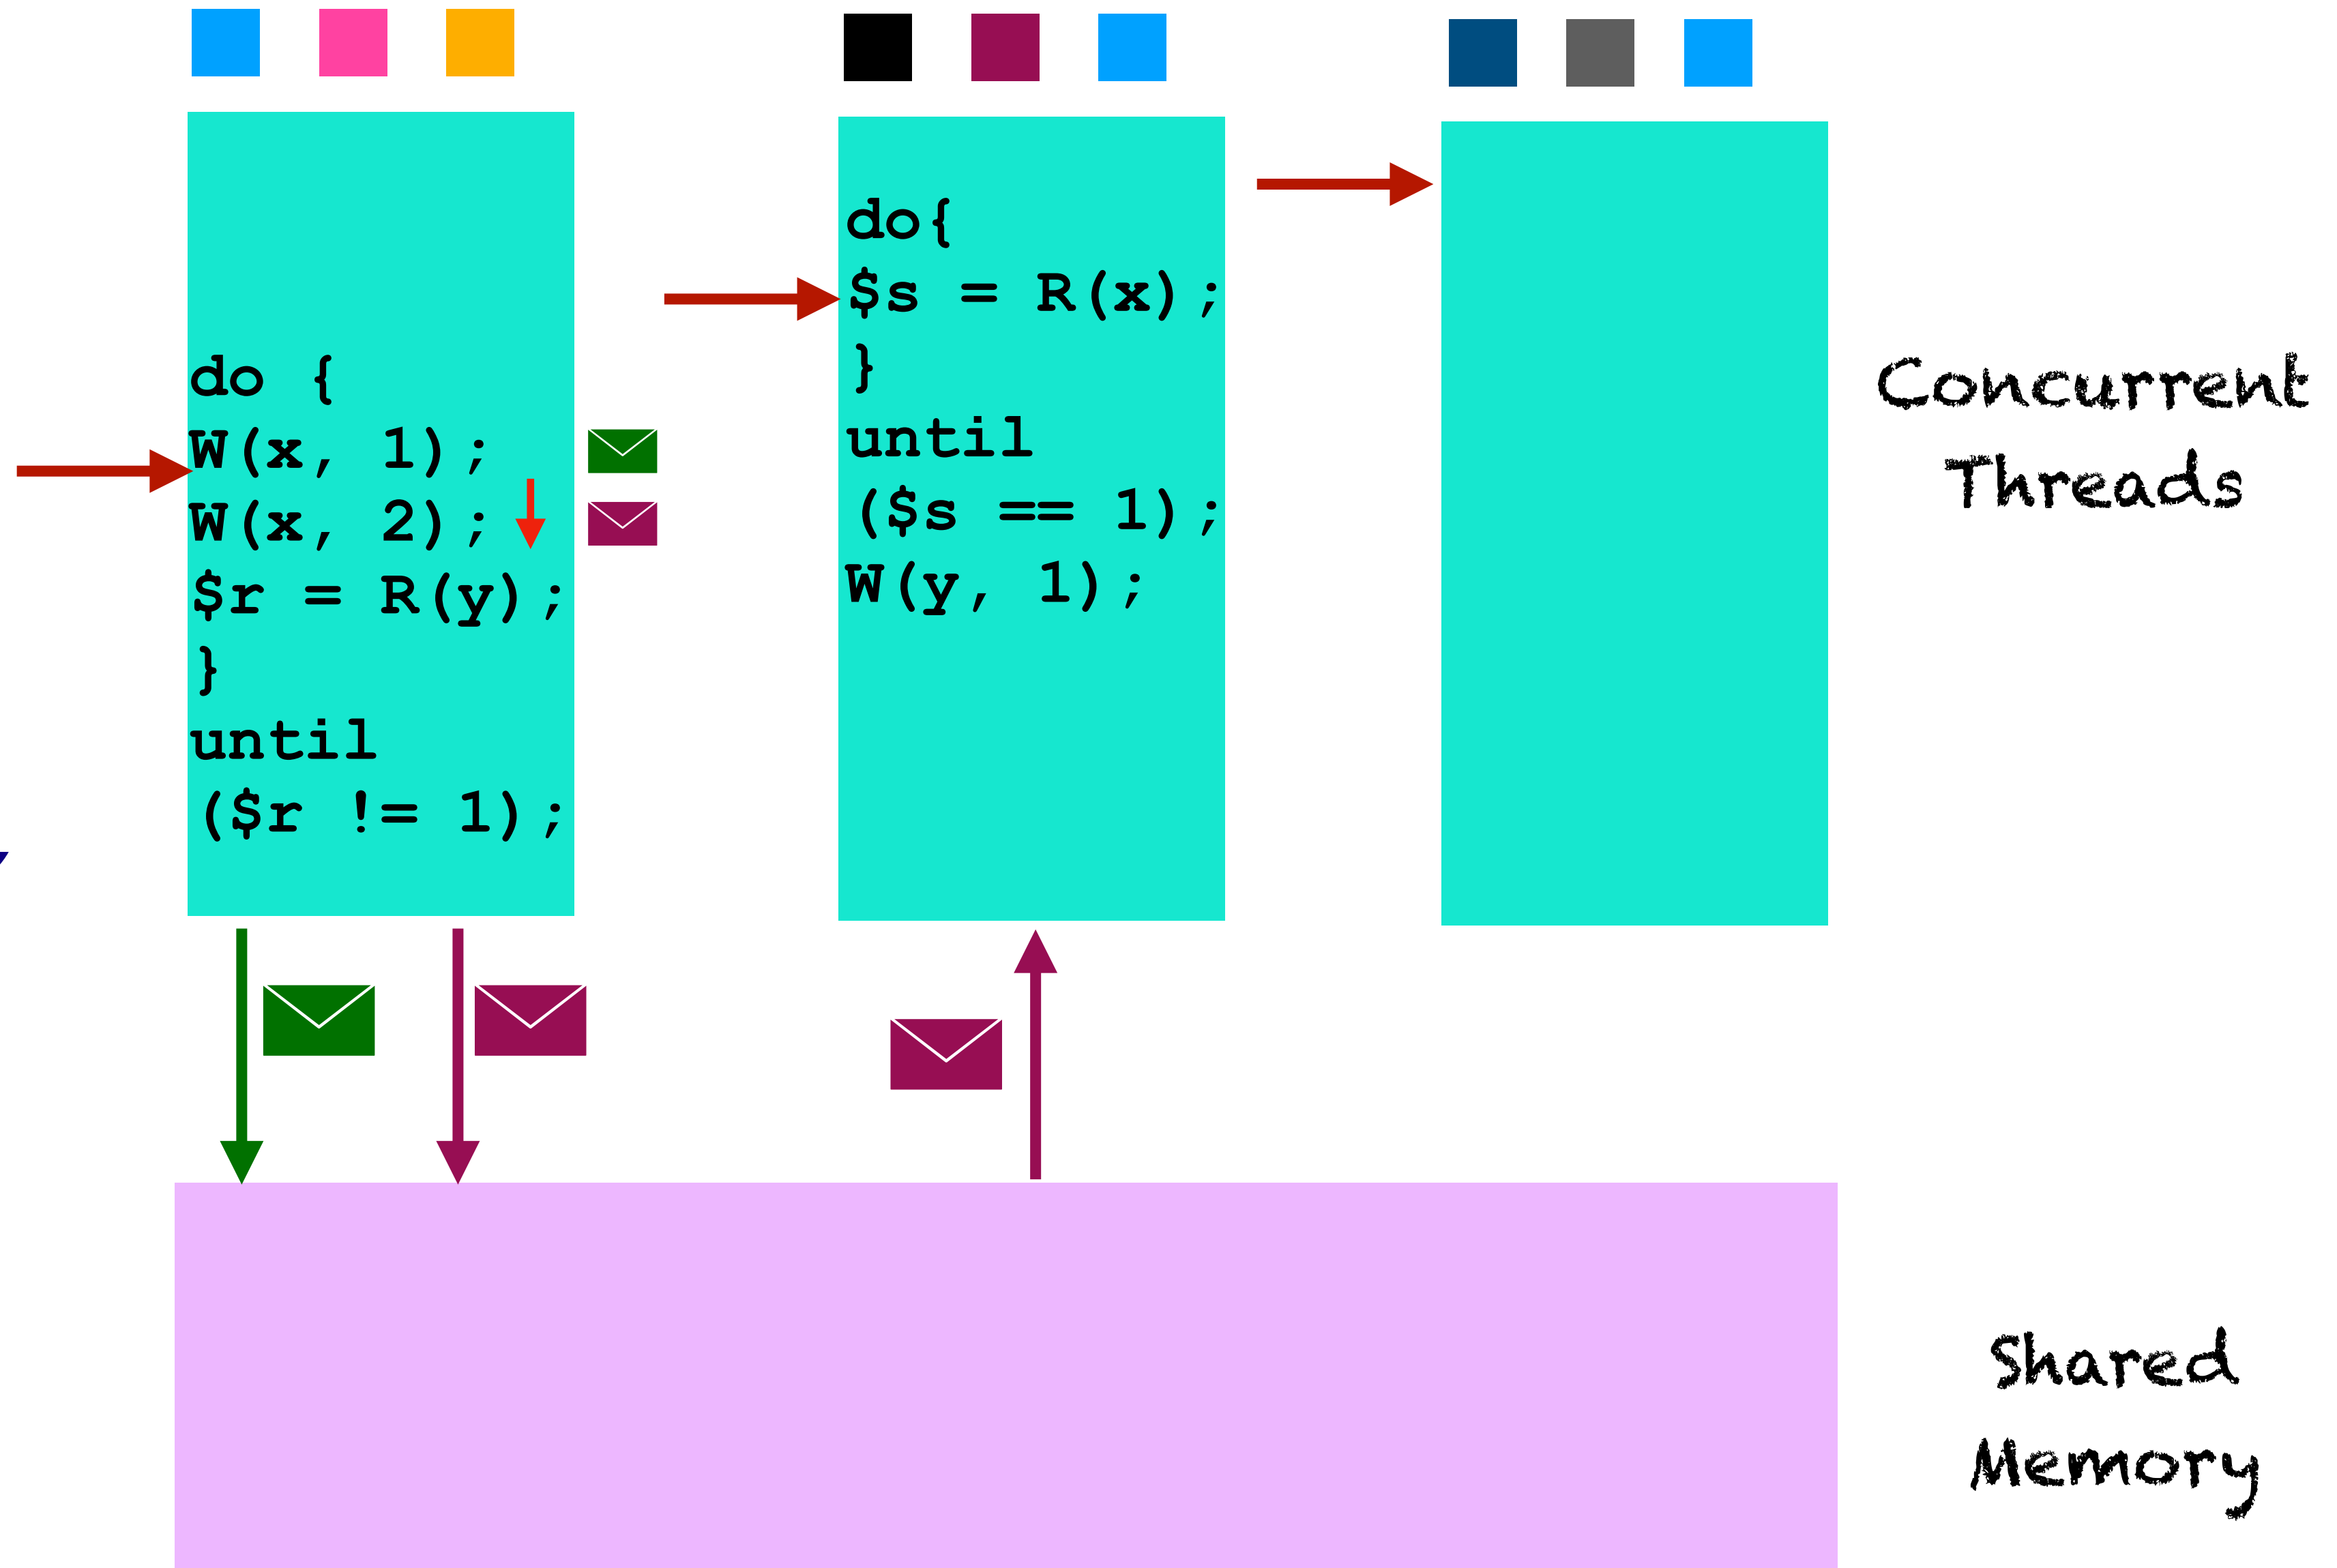


Our connection

# Concurrency and Memory: The Setup



# Need for transition fairness





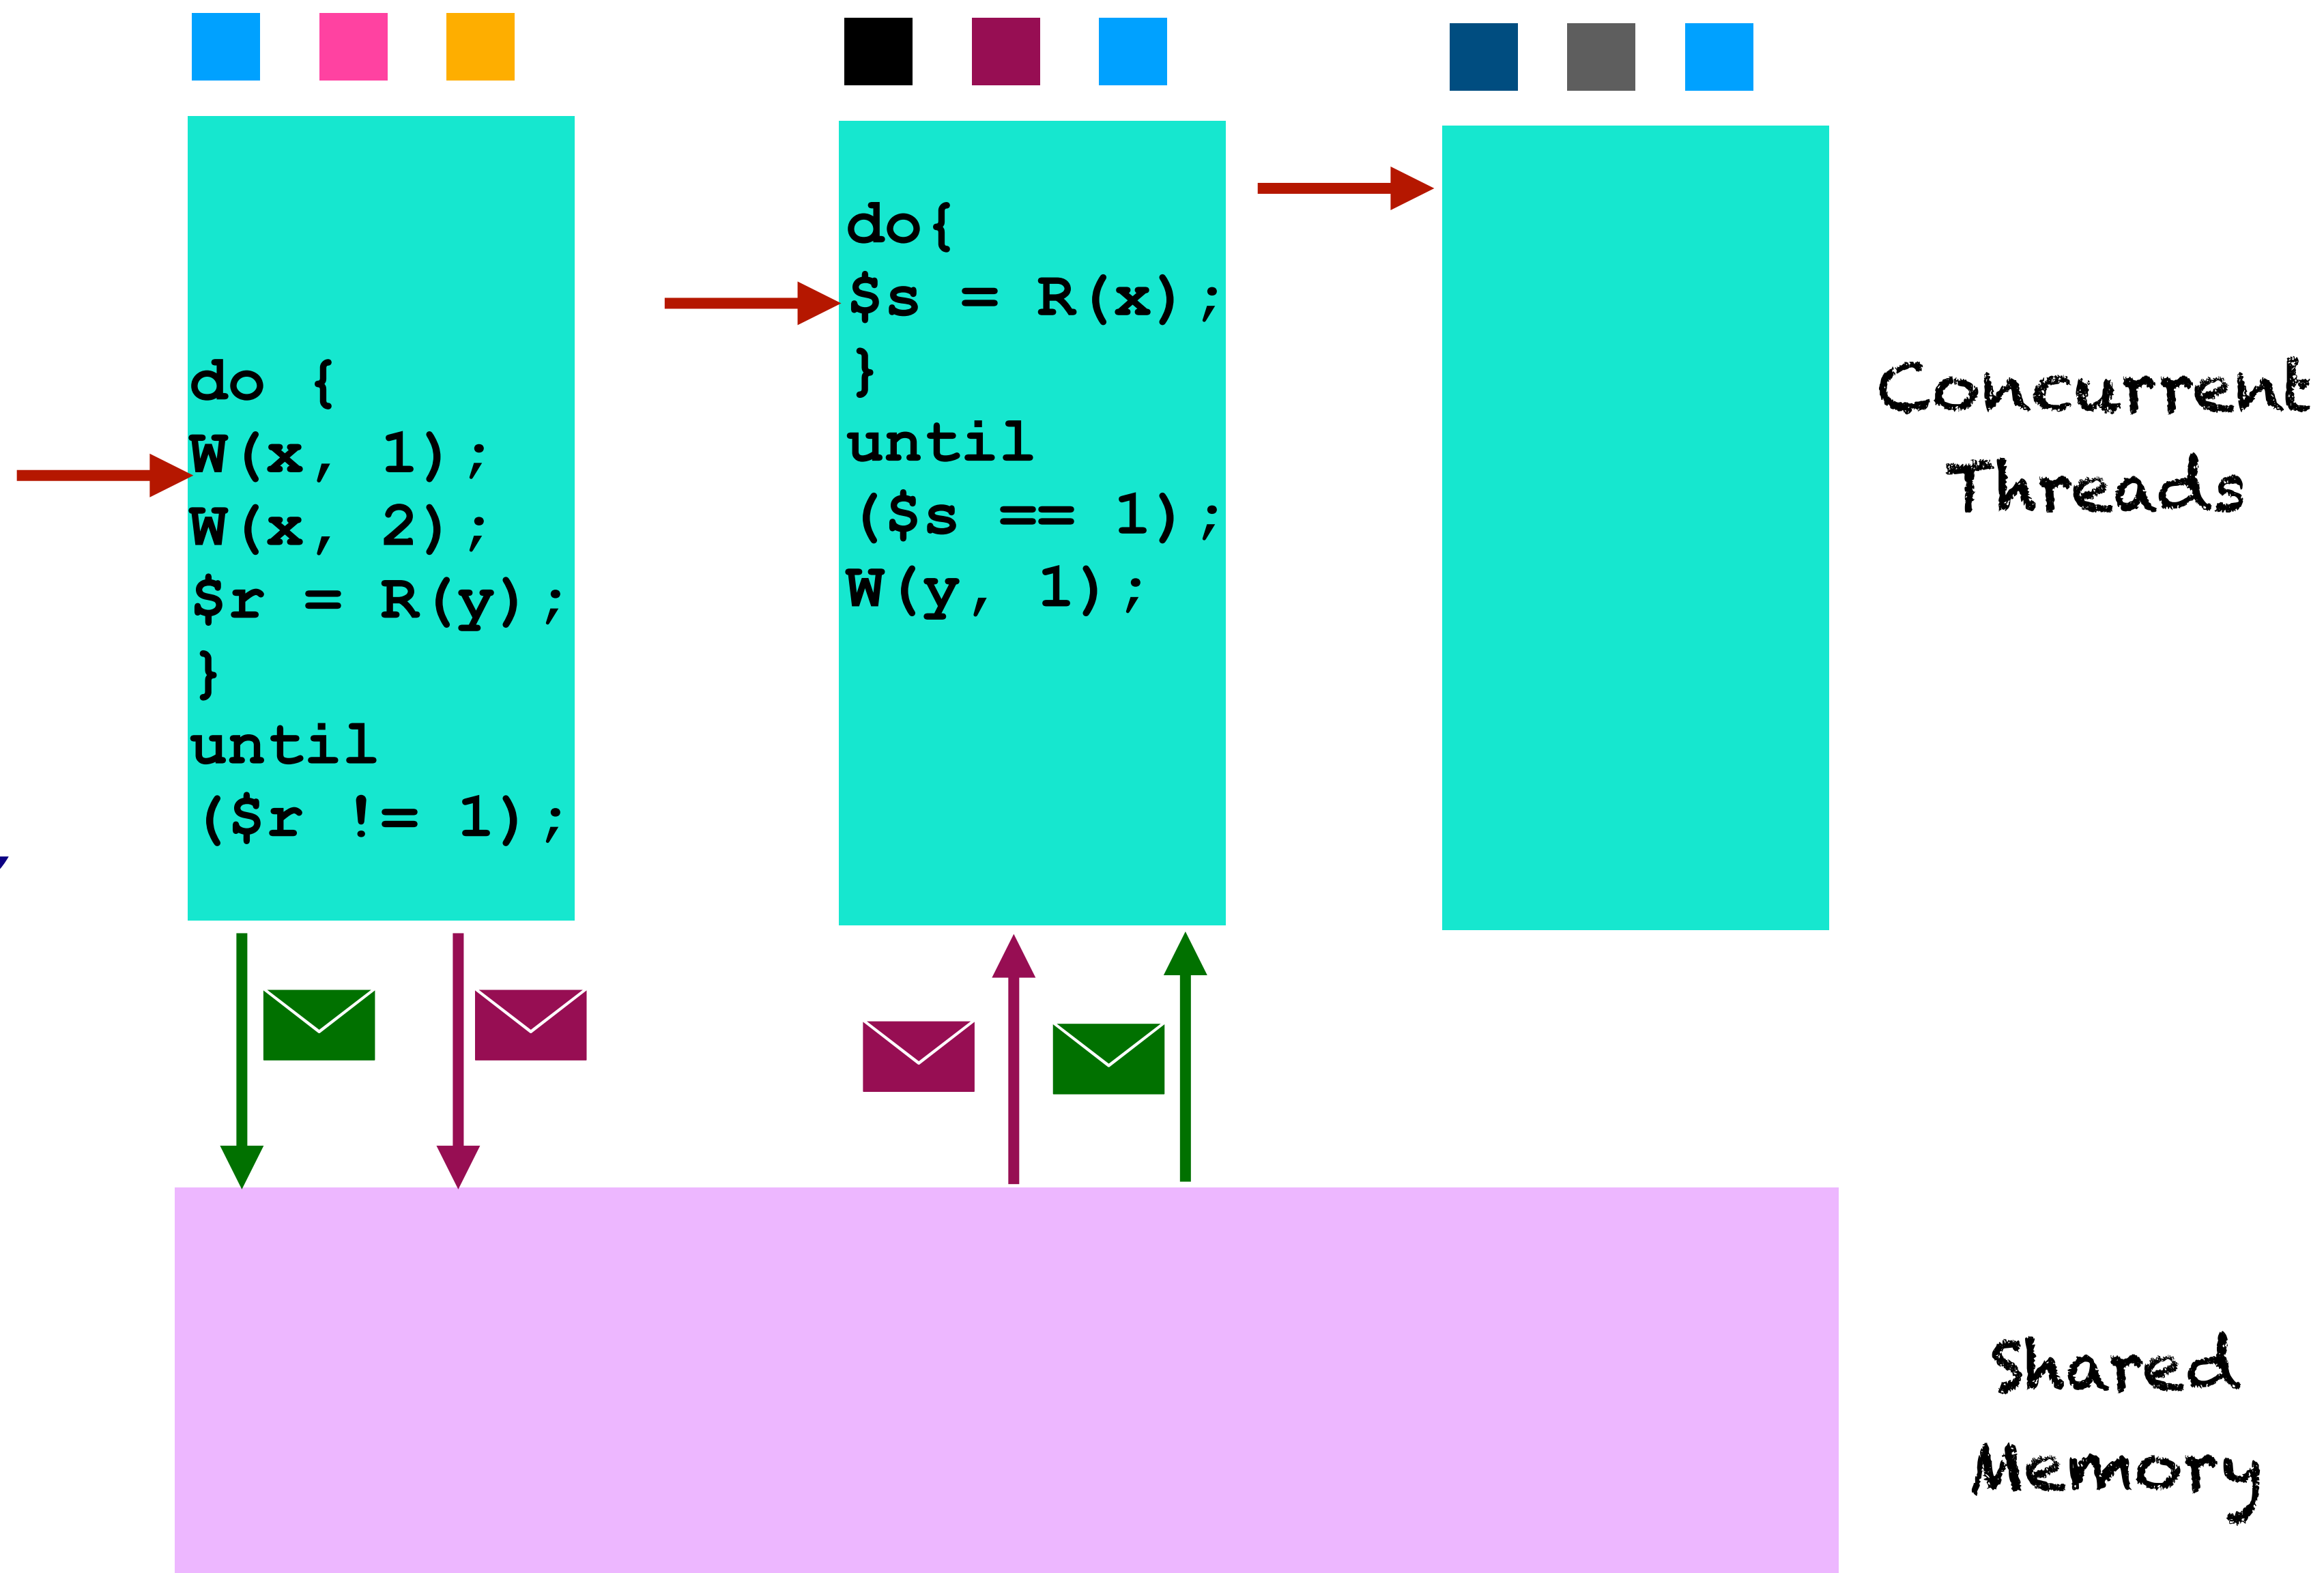
# Transition fairness

If a configuration  $c$  is visited infinitely often, then every transition  $(c, c')$  that is enabled from  $c$  is taken infinitely often.






# The resolution



# Reality is more complex: example




# Transition fairness falls short


**If** a configuration  $c$  is visited infinitely often,  
**then** every transition  $(c, c')$  that is enabled from  $c$   
is taken infinitely often.

But what if there are infinitely many configurations?  
An infinite run need not visit any configuration repeatedly!


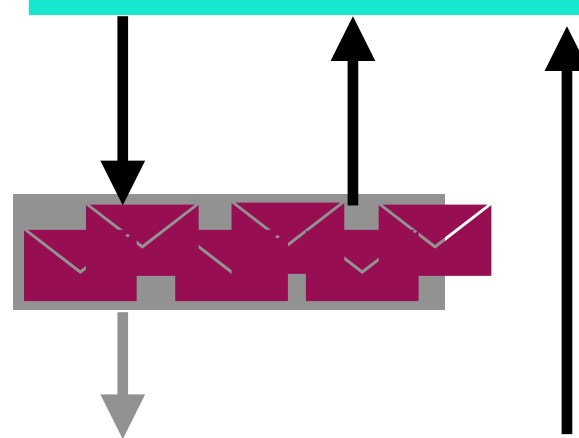
# Need for memory fairness



```
do
{
W(x, 1);
$r1 = R(x);
$r2 = R(y);
}
until
($r1 == 2
or $r2 == 1)
;
W(y, 1);
```



```
do
{
W(x, 2);
$s1 = R(x);
$s2 = R(y);
}
until
($s1 == 2
or $s2 == 1)
;
W(y, 1);
```



Concurrent  
Threads

Intermediate  
Buffers

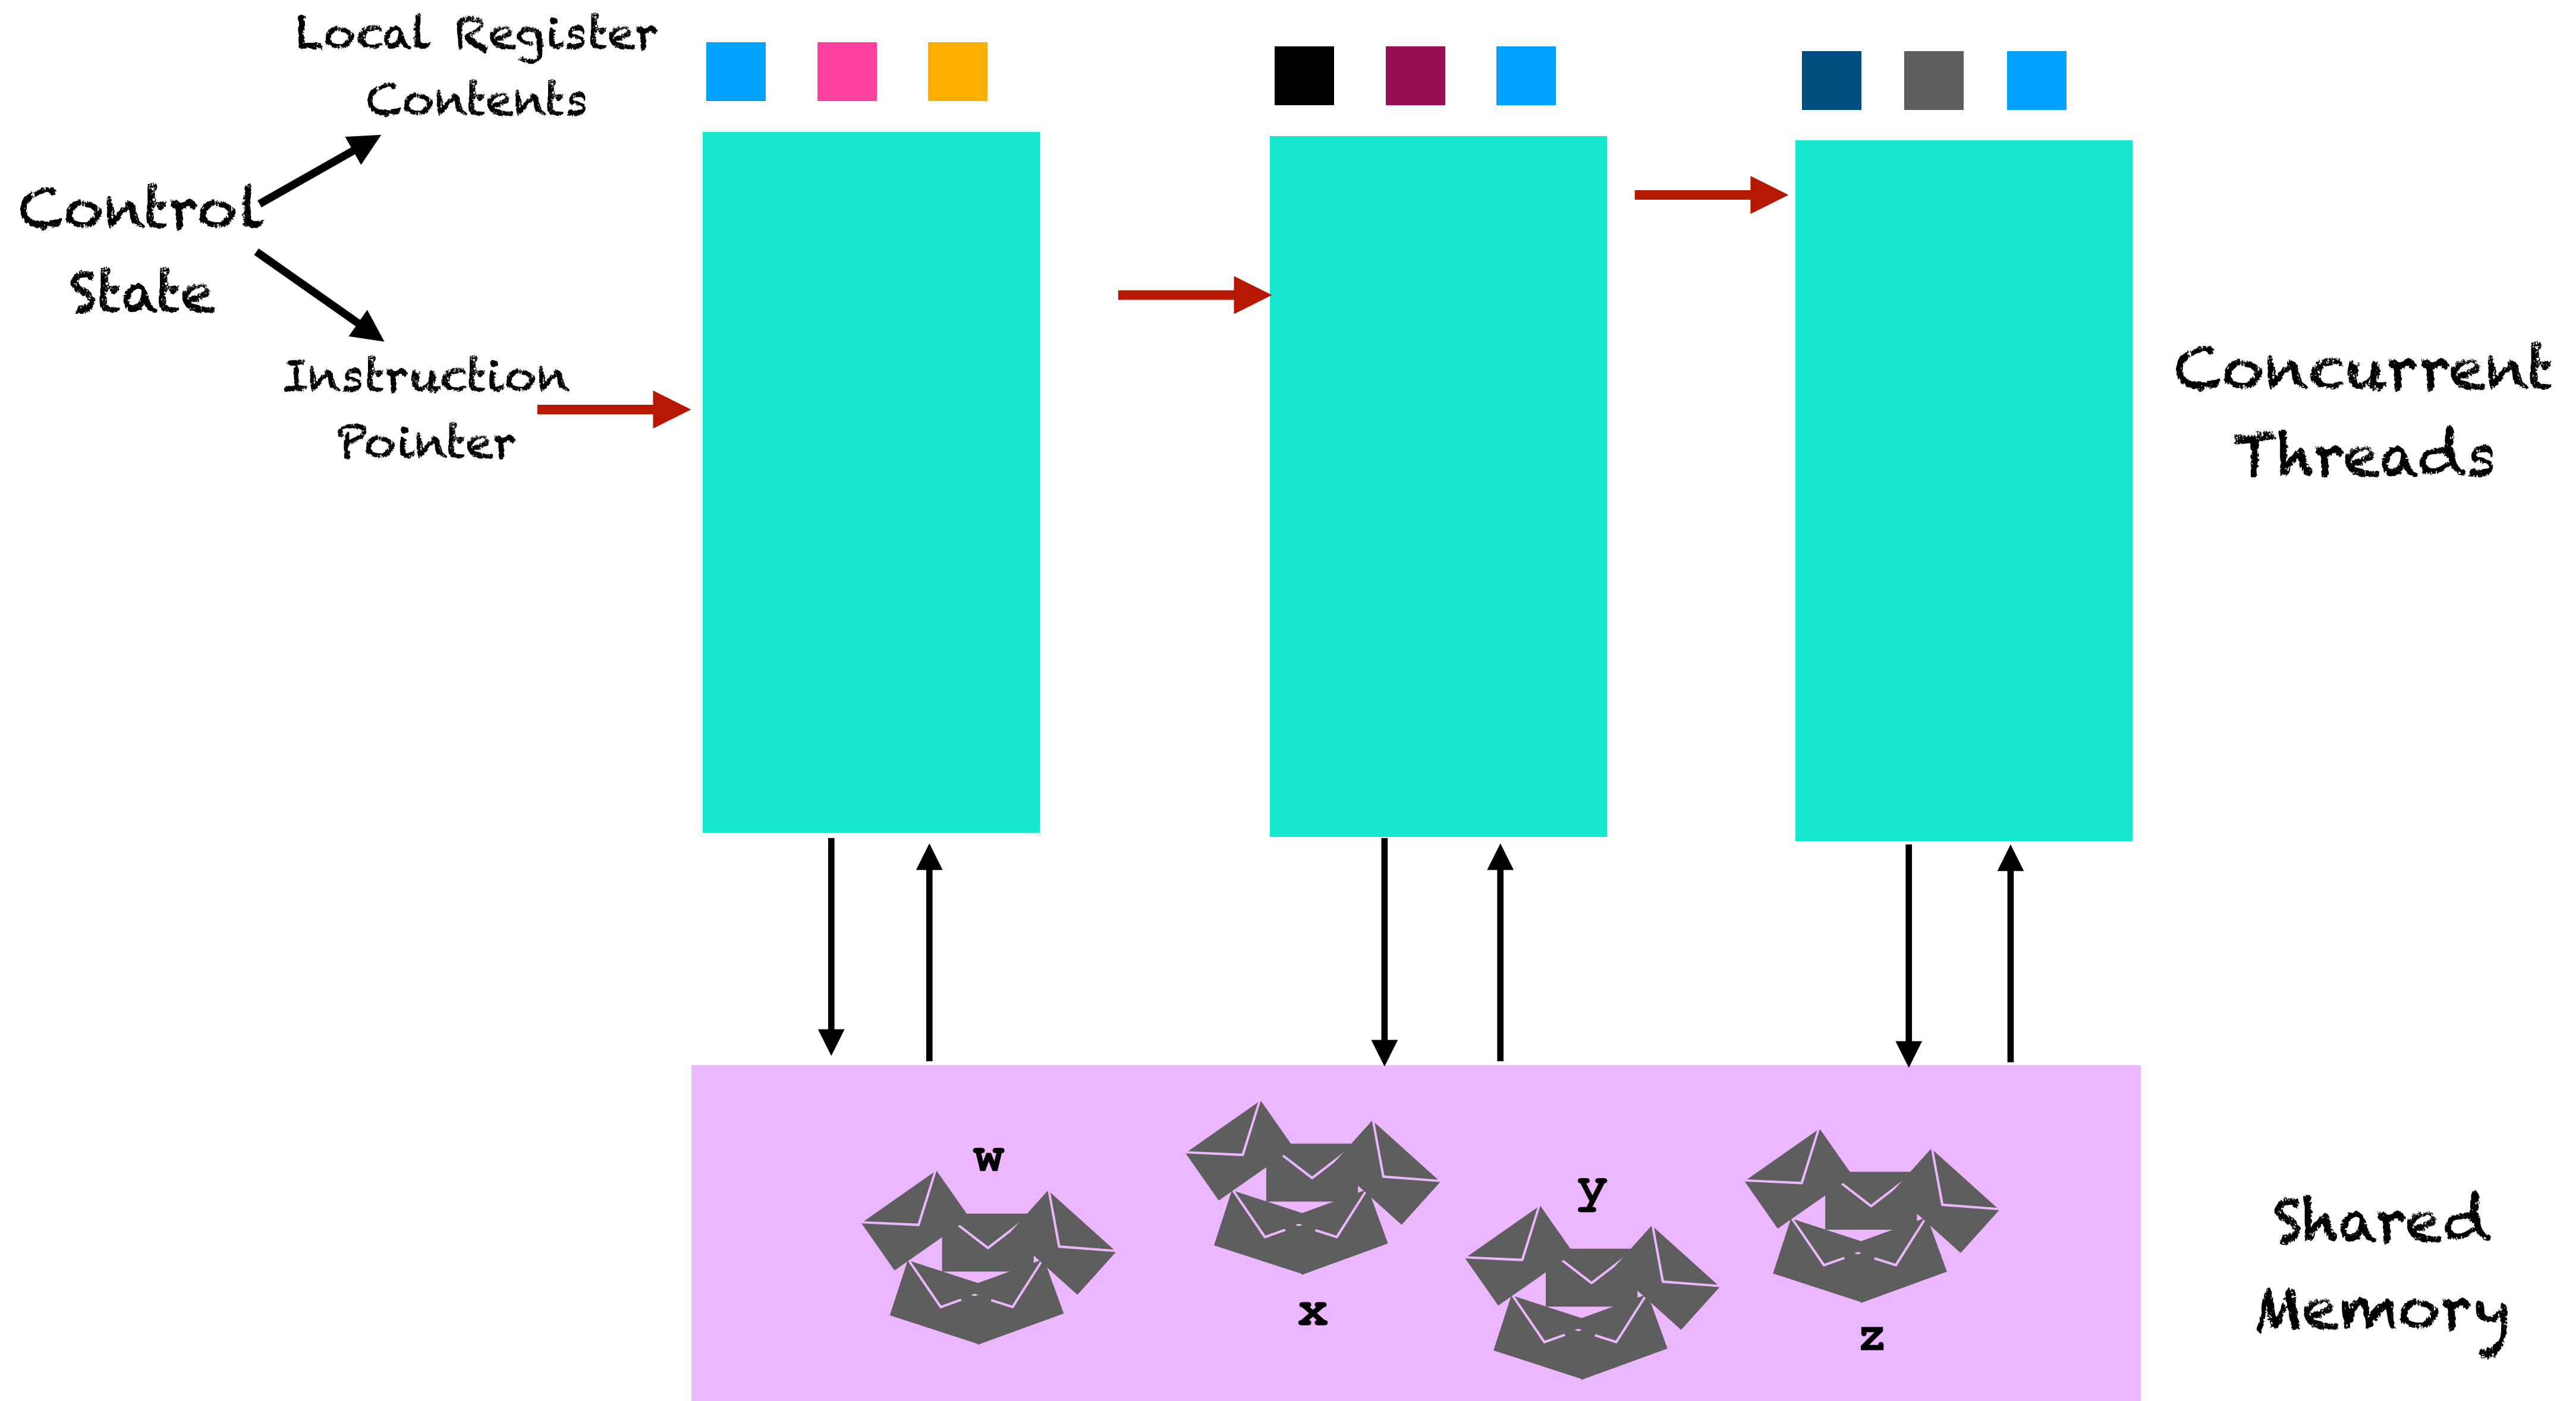
Shared  
Memory

# Memory fairness, informally



The “buffers” are flushed “regularly”.

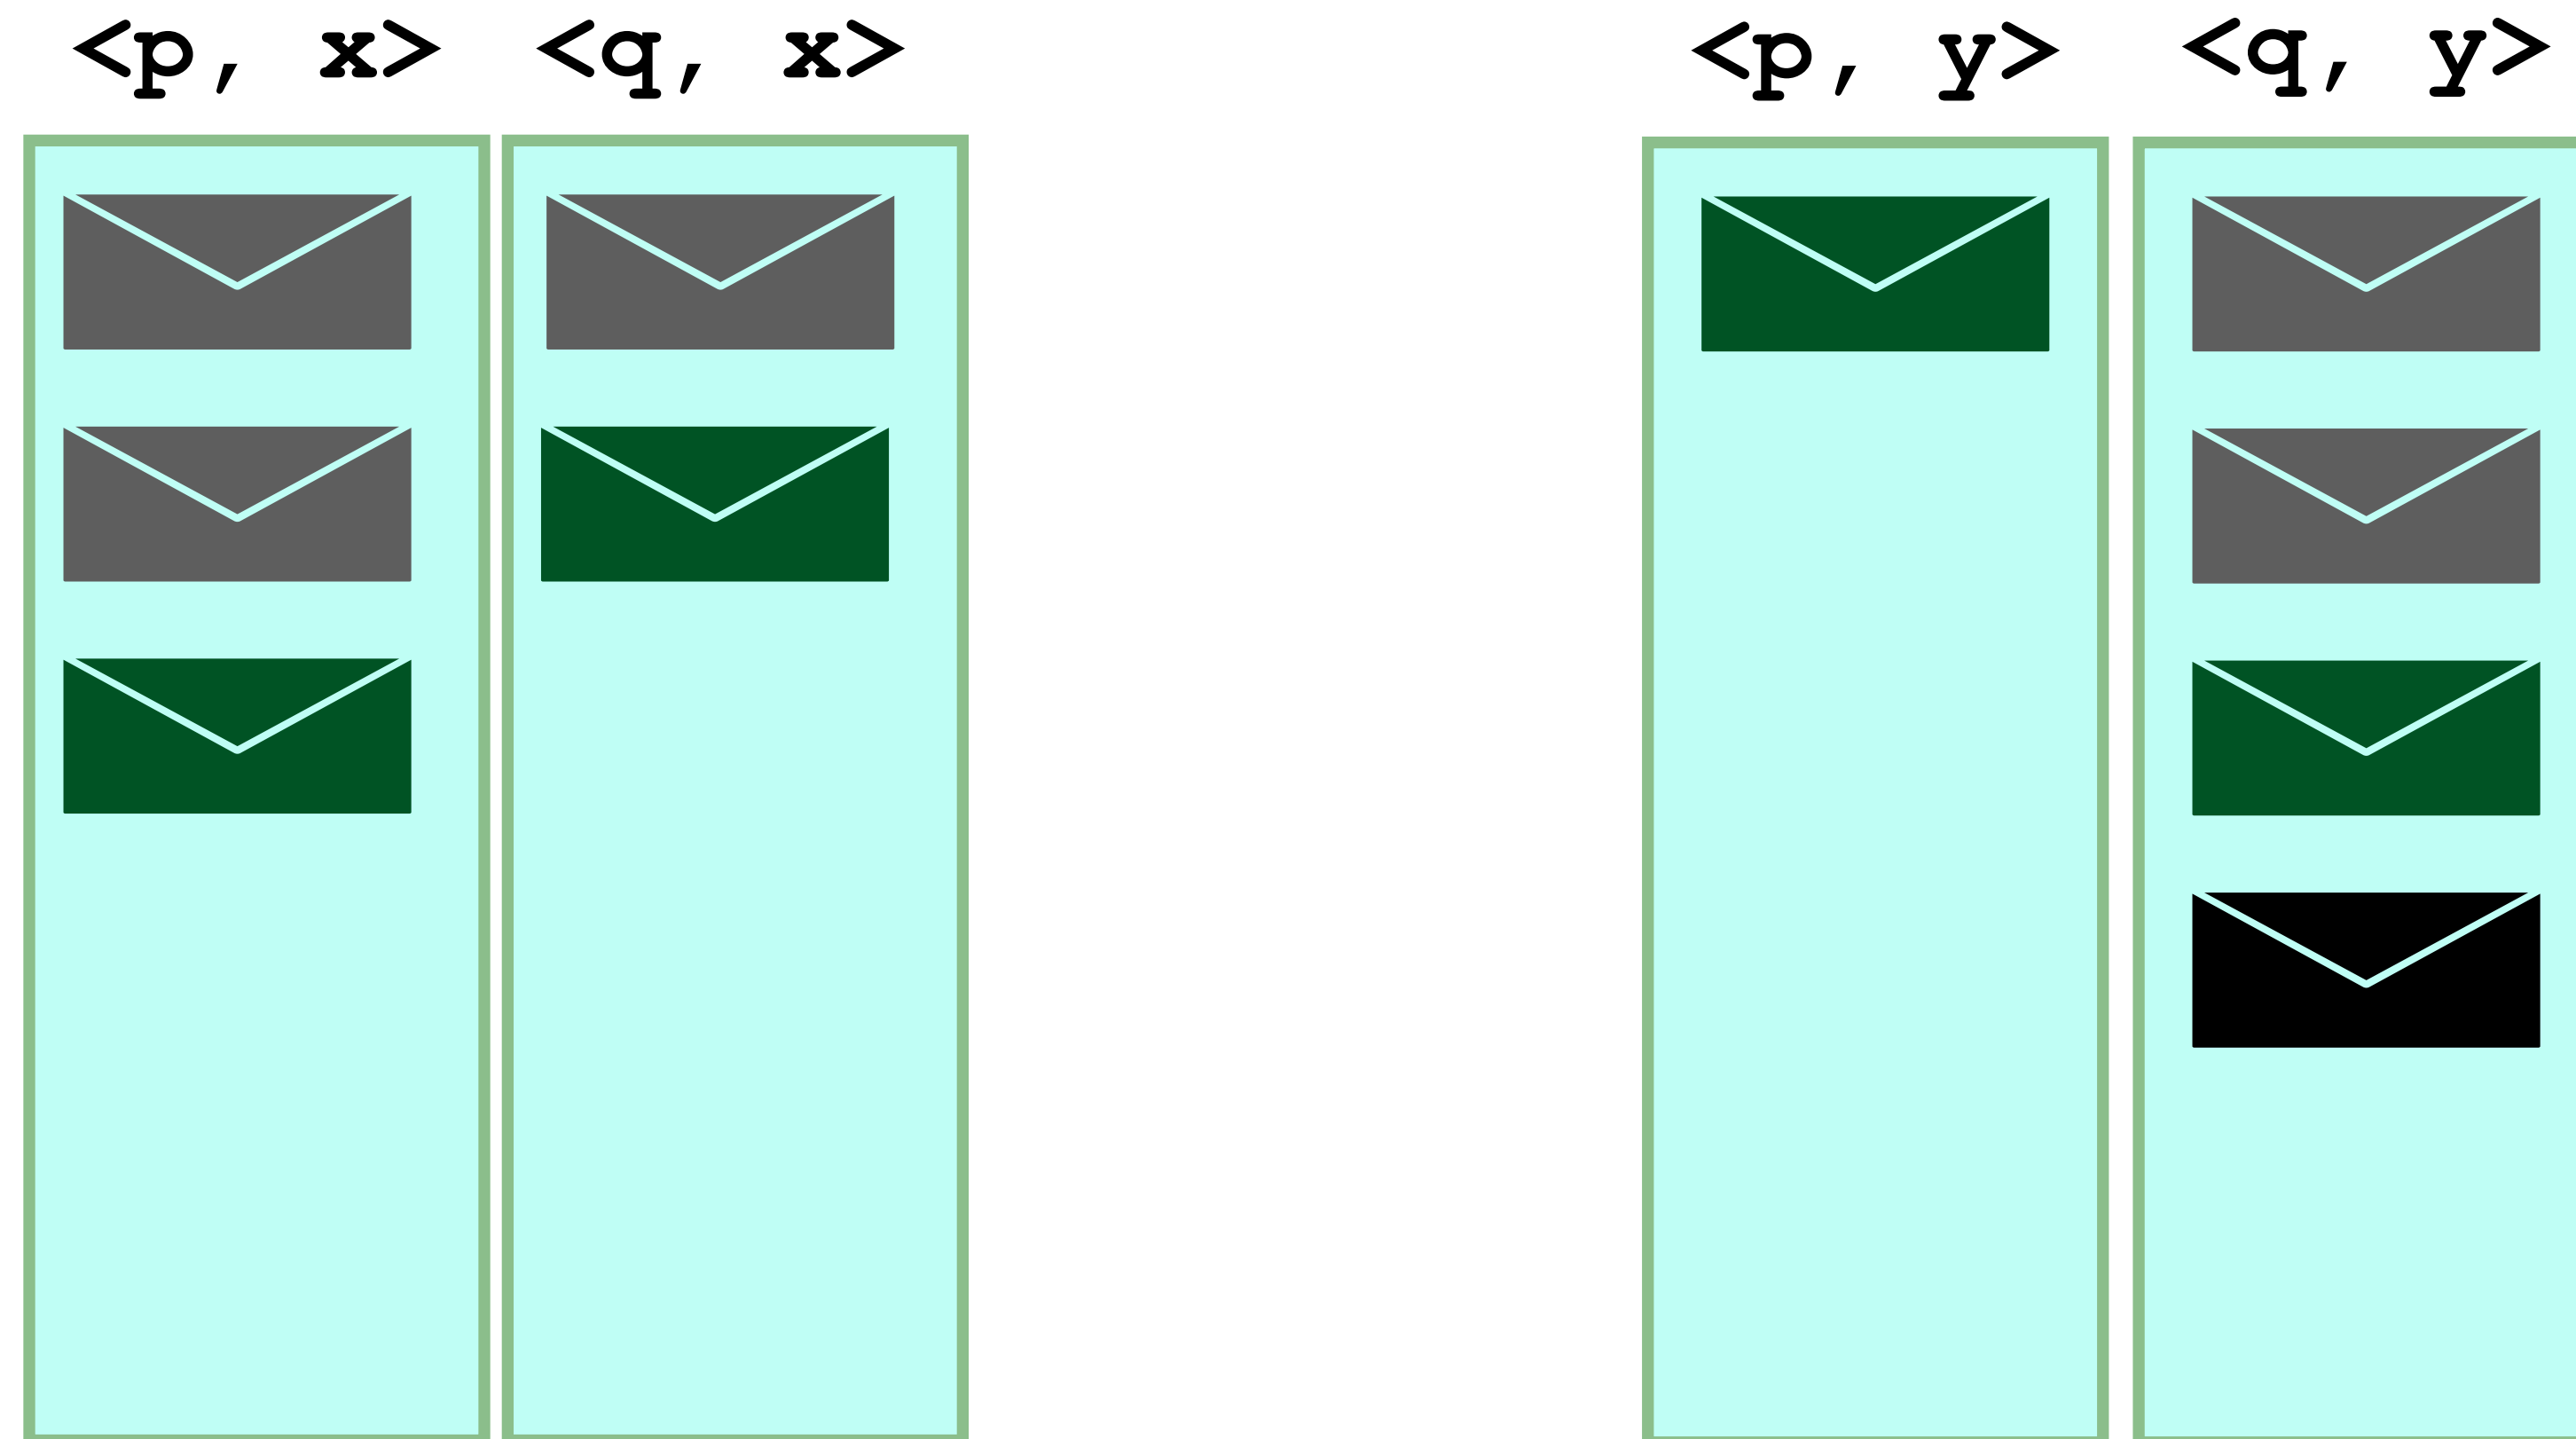
# Weak memory: an abstract idea



# How does weak memory propagate messages?


Consider writes by thread p to variable  $x$ .

They will always be observed in the same order in which they were made!





**Relative Strength of Memory Models: An arrow from A to B denotes that all behaviours of B are allowed by A.**



**Blue** denotes that the underlying reachability is **decidable**, **purple** denotes it is **undecidable**.

**Turquoise** arrows indicate that relative strength follows from **design**.

The **orange** arrow indicates the enforcement of **acquire semantics** on reads.

**Brown** arrows indicate the enforcement of **multi copy atomicity** on the memory model.

# Configuration Size


Constraints imposed by the memory model make messages redundant as the run progresses

We only keep track of messages that are not redundant!


The number of messages stored in a memory configuration is called its size.

# Size Bounded Executions

Memory Fairness



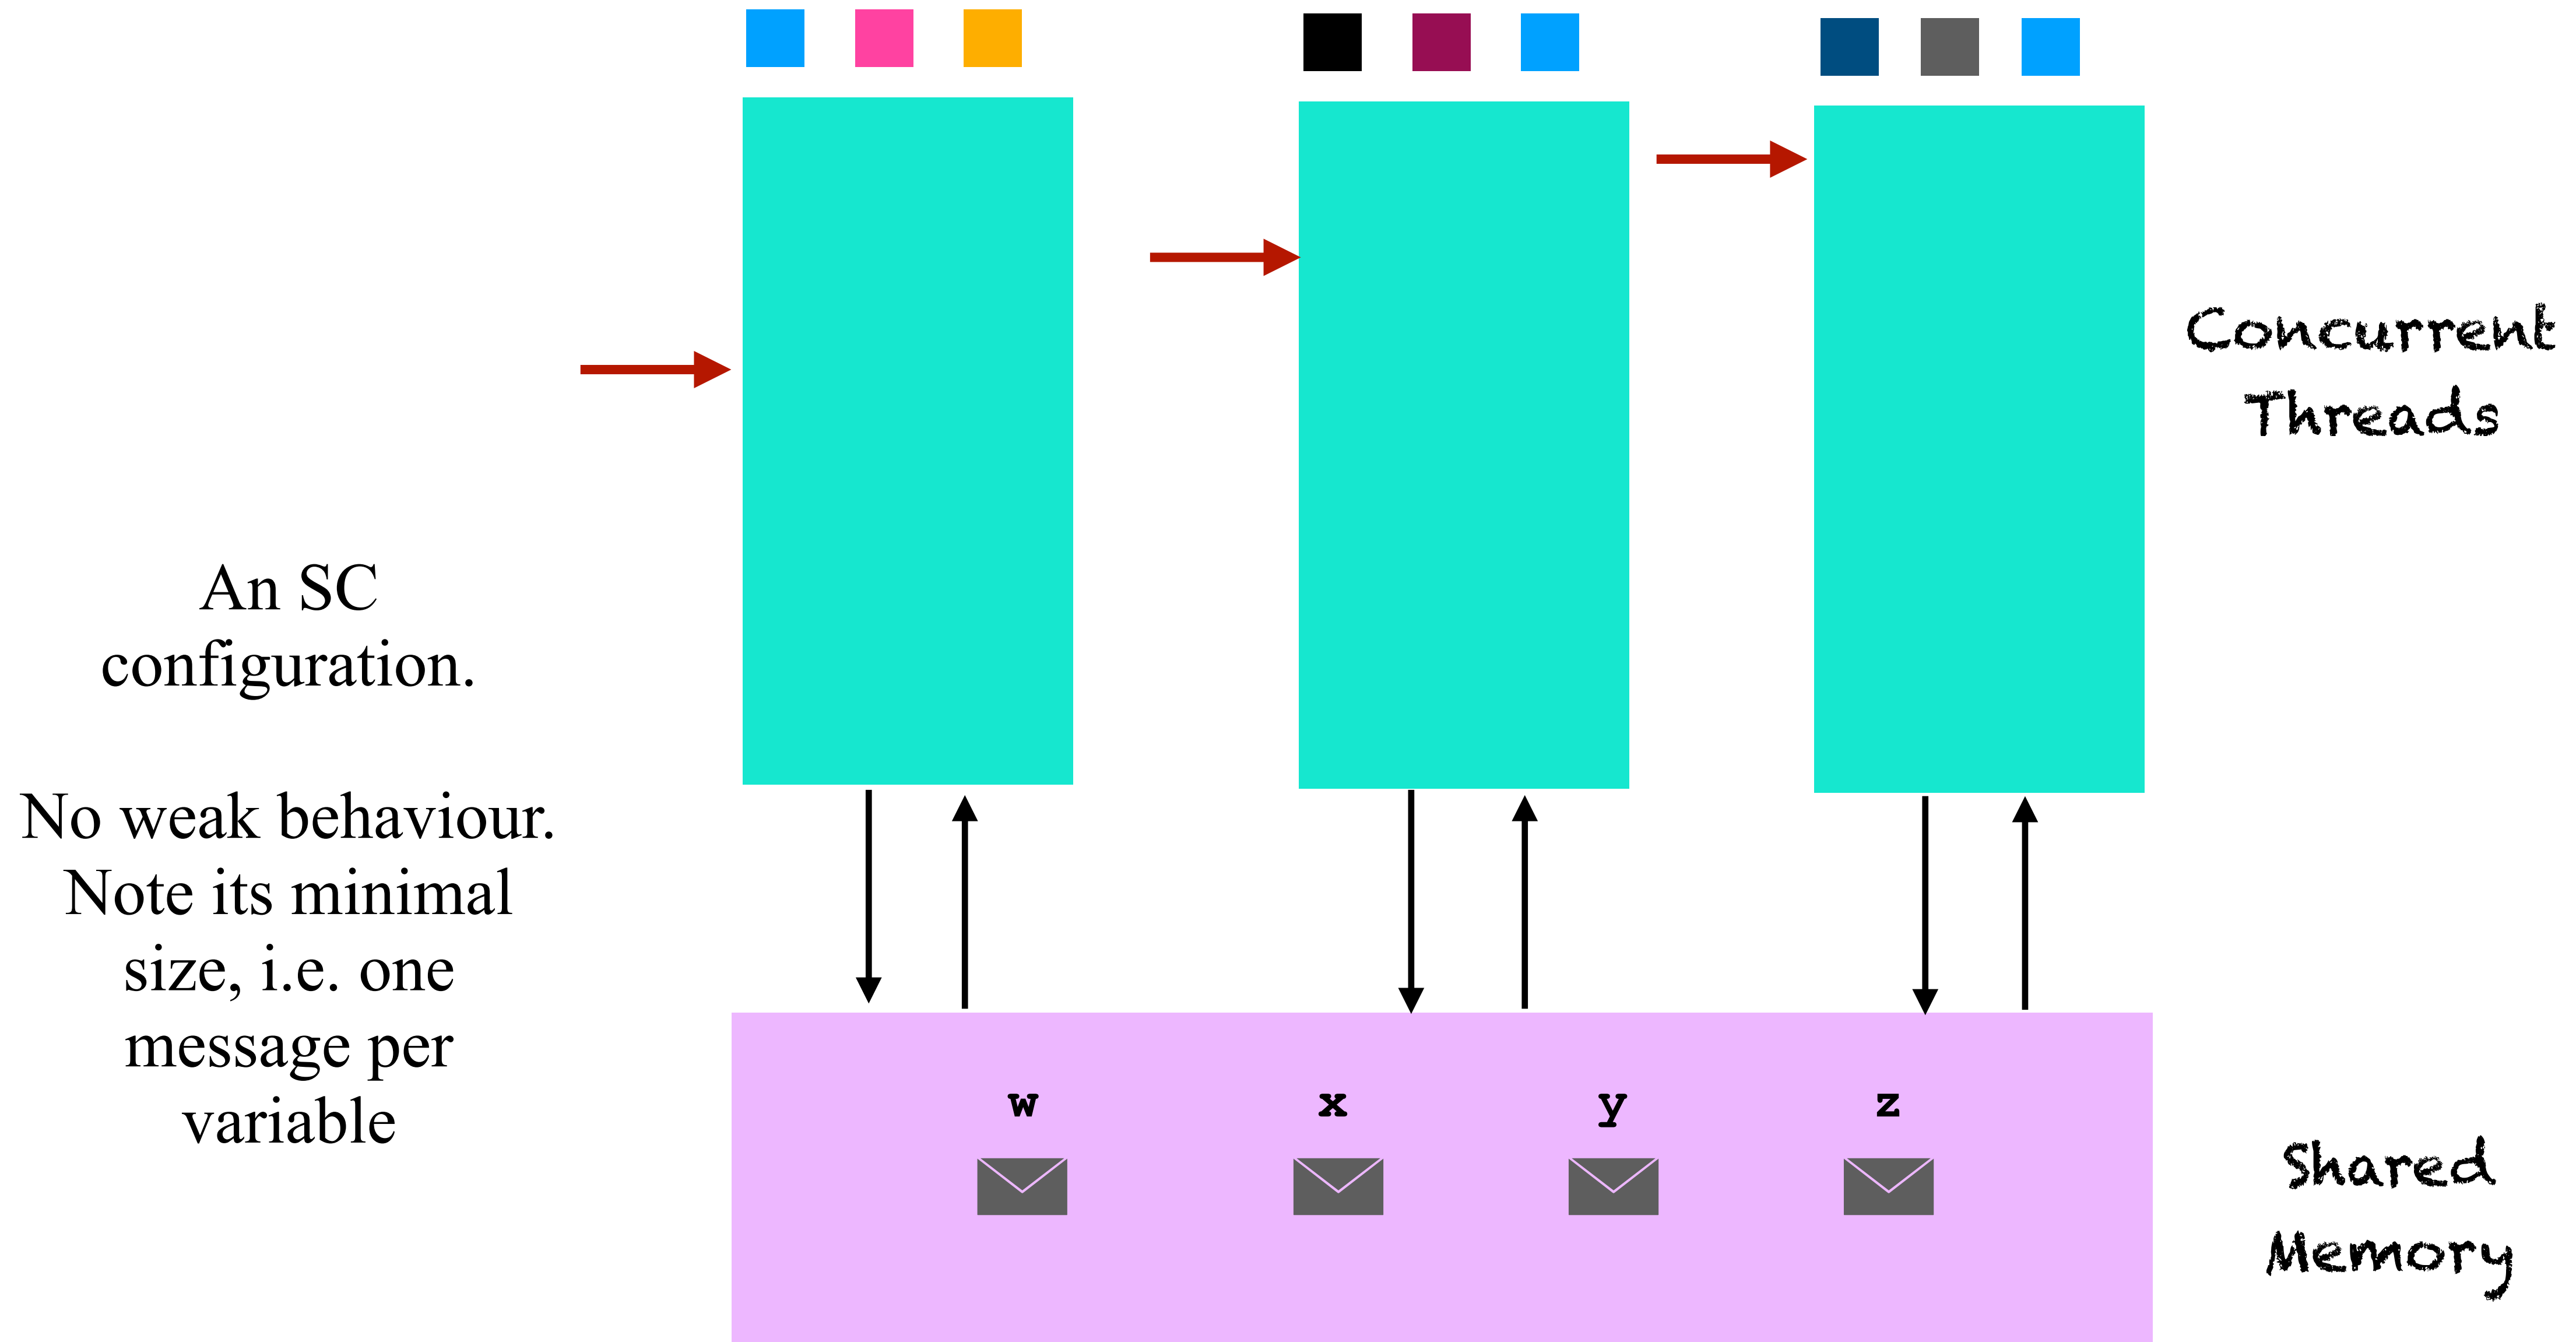
Unified framework



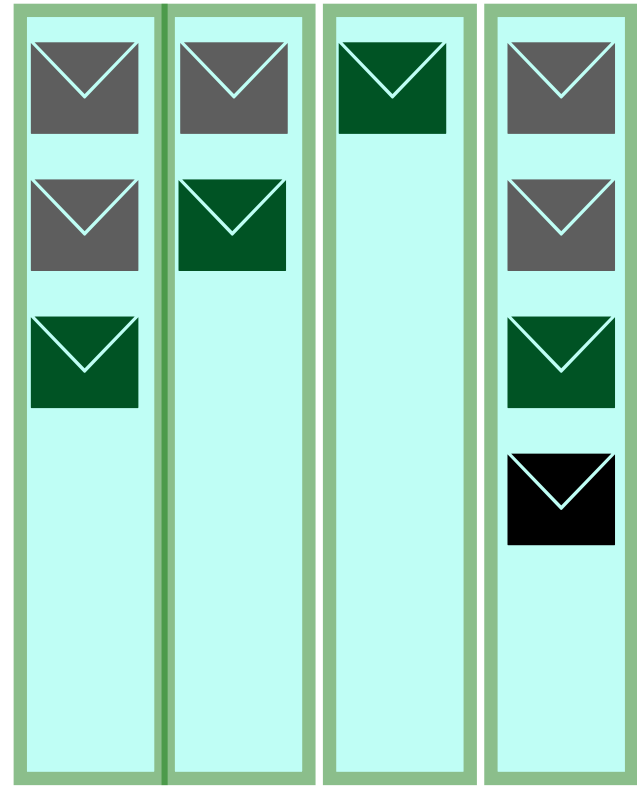
An execution is called size bounded if there exists an  $N$  such that each configuration is of size at most  $N$ .

If  $N$  is specified, we refer to the execution as  $N$ -bounded.

# Configuration size and weakness



# Plain Configurations




**Configurations with exactly one message per variable are called plain.**

**There are finitely many plain configurations**


# Repeatedly Plain Executions

Memory Fairness





An infinite execution is repeatedly plain if plain configurations occur infinitely often.

Unified framework




# Transition + Memory Fairness, Formally

Memory Fairness



Unified framework




We have the following fairness conditions on infinite executions

- N-bounded transition fairness
- Repeatedly plain transition fairness




# A Probabilistic Analog


Memory Fairness



A Markov chain induced by the system satisfies Probabilistic Memory Fairness if the set of plain configurations is visited infinitely often with probability 1.



Unified framework



Such Markov Chains are "decisive" by dint of having the set of plain configurations as a "finite attractor"




Decisive Markov Chains are well studied

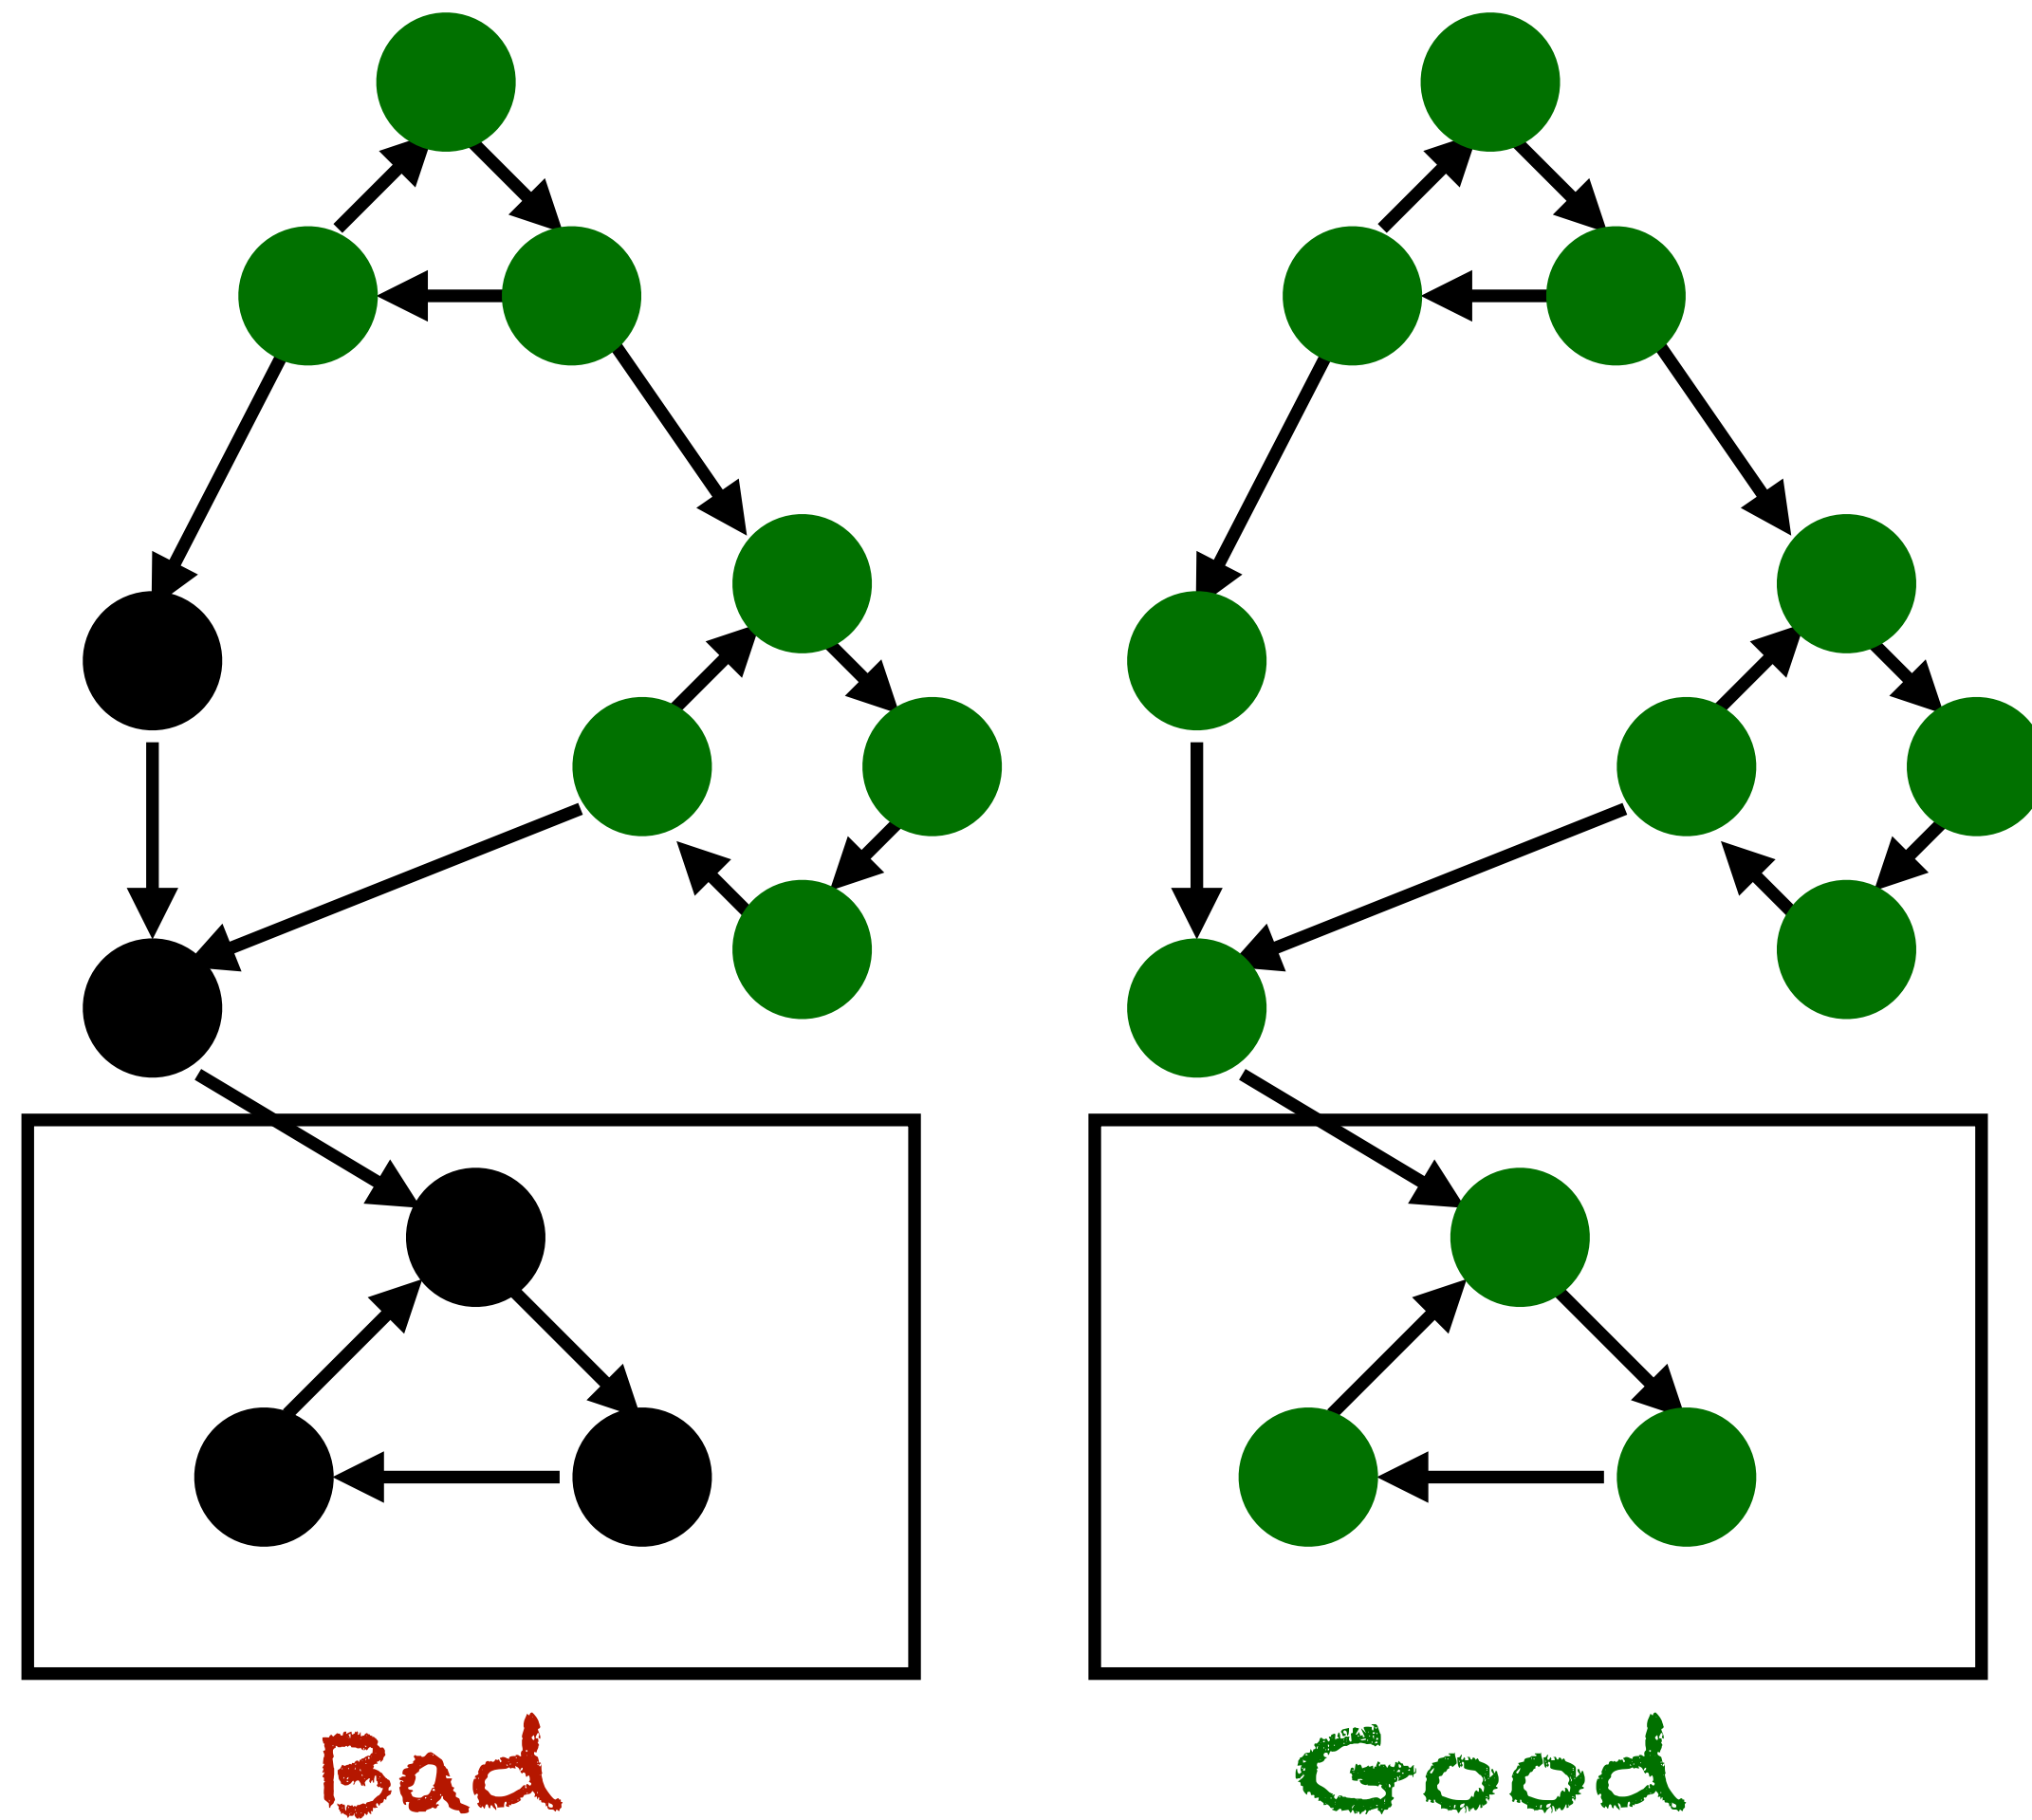
# The connection

The following fairness conditions are equivalent for termination and repeated control state reachability

- Probabilistic Memory Fairness
- $N$ -bounded transition fairness for sufficiently large  $N$
- Repeatedly plain transition fairness




# Proof Sketch



1) For each  $N$ , construct a graph  $G(N)$  with plain configurations as vertices

2) Draw an edge  $(\gamma, \gamma')$  if  $\gamma'$  is reachable from  $\gamma$  via configurations of size at most  $N$

3) Paint a node green if the control state of interest is reachable via configurations of size at most  $N$




# Proof Sketch


Notice, edges can only be added, and nodes can only go from black to green!

The finite graph saturates; let it be  $G$  for all sufficiently large  $N$


For all our fairness notions, liveness holds if and only if **all bottom scc's of  $G$  are green**



Bad




Good




# Liveness, Verified

All that remains is to construct  $G$  using reachability queries

This can be done by translating our framework into those used for verifying safety



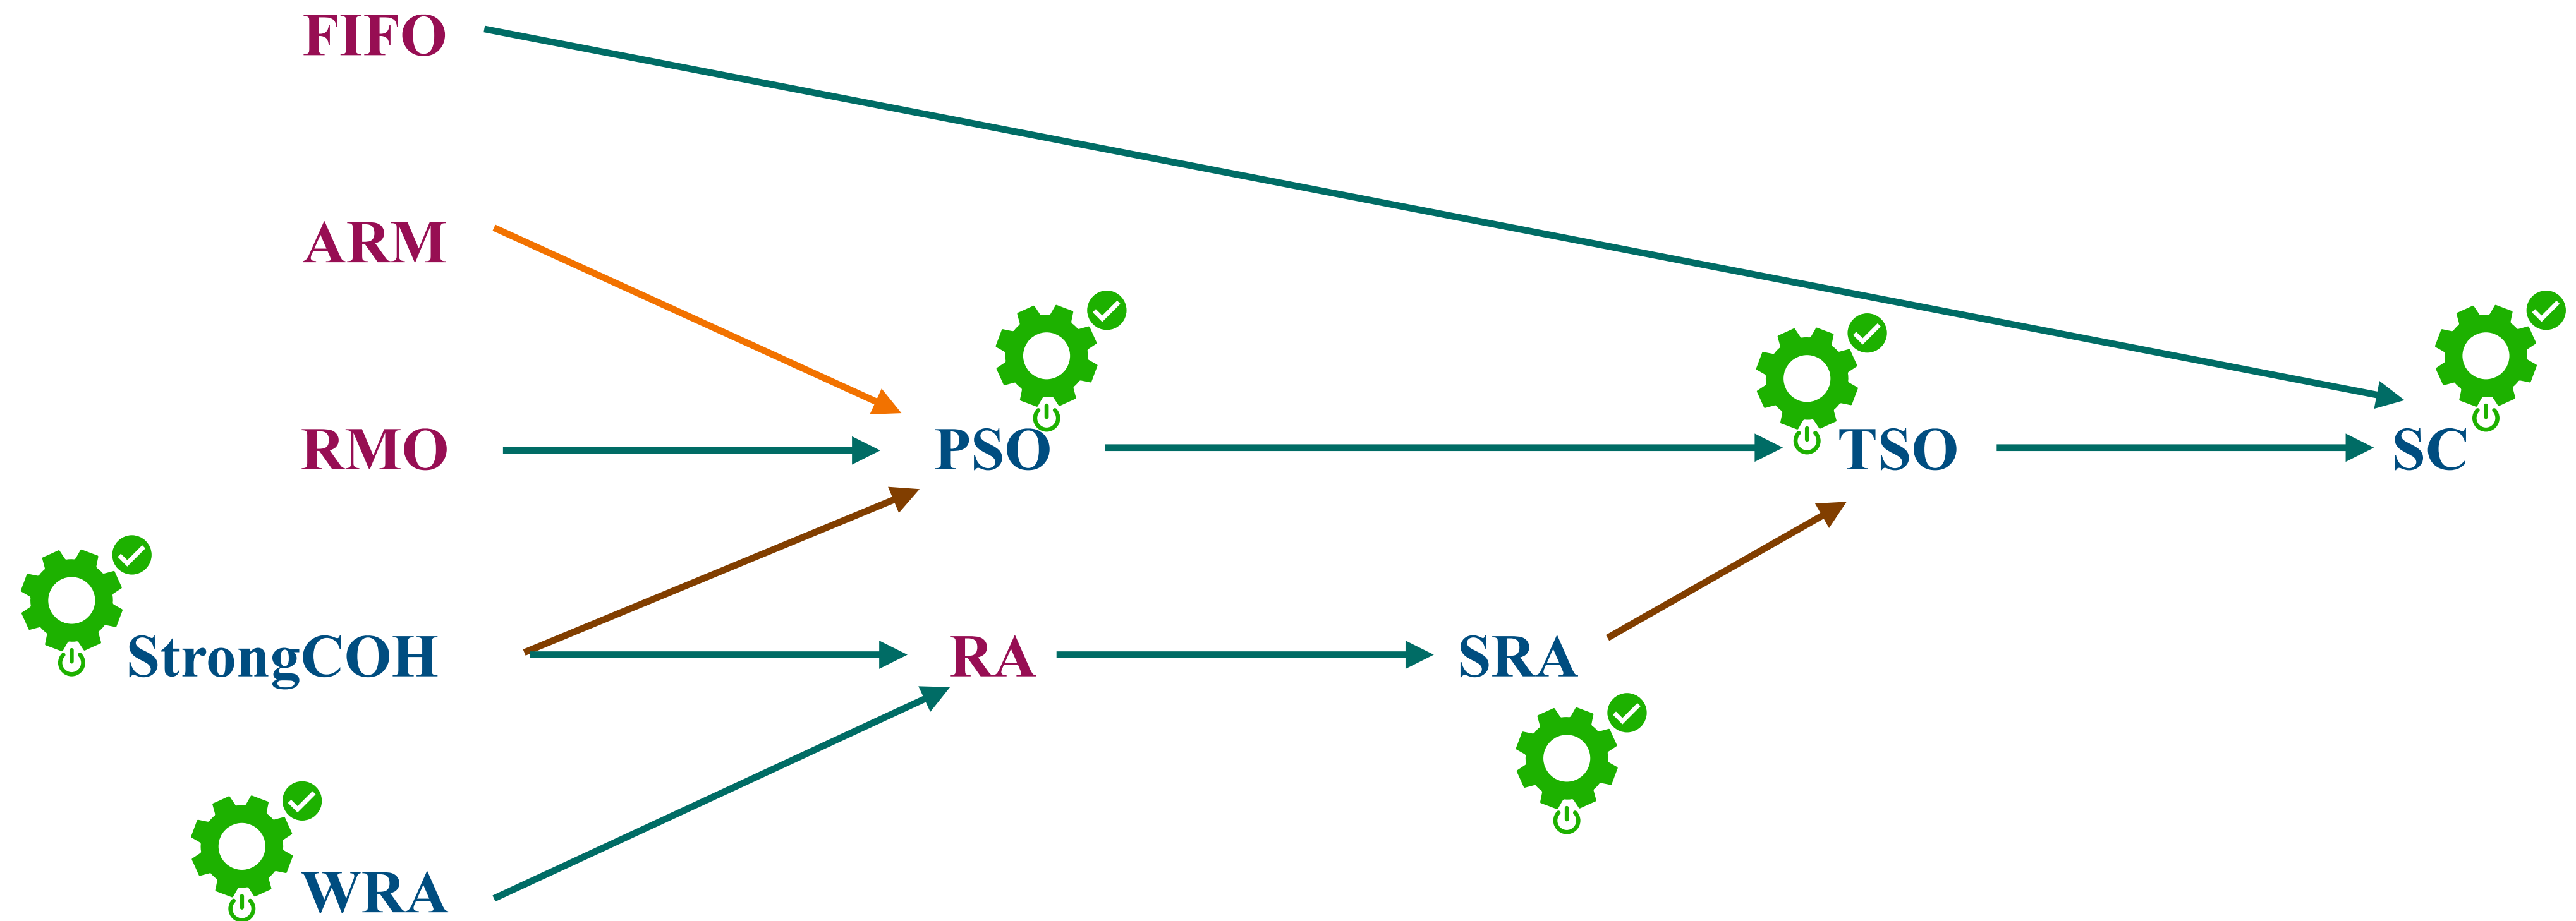
Bad



Good




# Verifying concrete models




Reachability queries result in liveness decision procedures


# Thank You!



The Setup



The Model



The Procedure