THE UNIVERSITY of EDINBURGH
informatics

ΠΑΝΕΠΙΣΤΗΜΙΟ ΚΡΗΤΗΣ
UNIVERSITY OF CRETE

Towards the Synthesis of Coherence/Replication Protocols from Consistency Models via Real-Time Orderings

Vasilis Gavrielatos, Vijay Nagarajan, Panagiota Fatourou

Thanks to:

Shared Memory Systems

Distributed Systems
-- NoSQL Databases
-- Coordination Services
-- DSMs

Computer Architecture
-- Multiprocessors
-- GPUs

Consistency Model

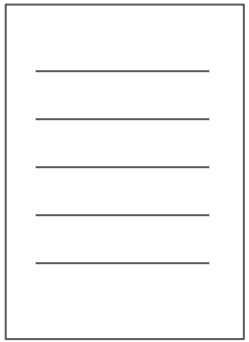

Consistency Model

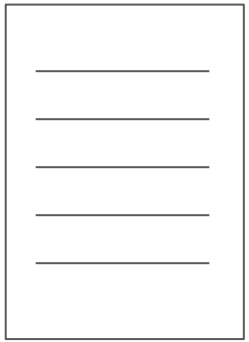

Sequential Consistency
Release Consistency
Causal Consistency
Eventual Consistency
Total Store Order
...

Consistency Model

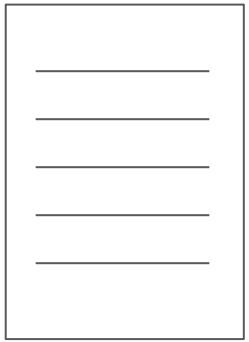

Synchronization Patterns

Consistency Model

Programmers

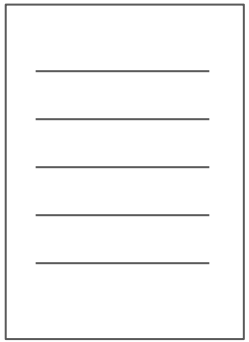

Designers

Consistency Model

Programmers

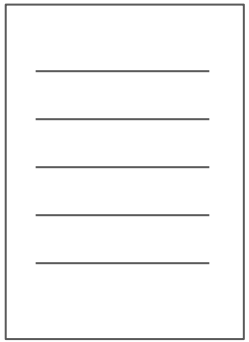

Designers

Consistency Model

Programmers

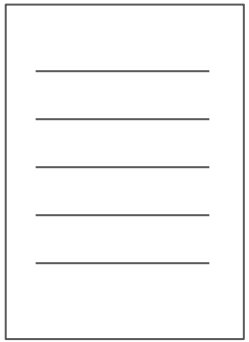

Designers

Replication Protocol:
e.g., Raft, CR, ABD,
ZAB, COPS, Kite

Consistency Model

Programmers

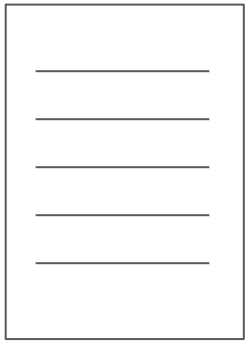

Designers

Replication Protocol:
e.g., Raft, CR, ABD,
ZAB, COPS, Kite

Coherence Protocol:
e.g., MOESI, TSO-CC
HMG

Protocol

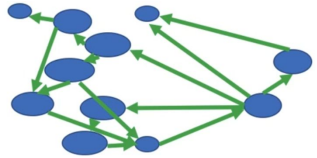

Consistency Model

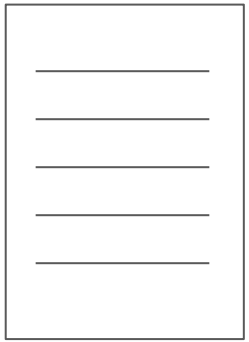

Programmers

Designers

Abstraction

Consistency Model

Protocol High-level

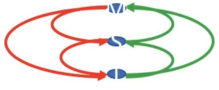

Abstraction

Consistency Model

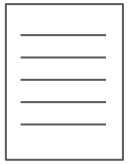

Protocol High-level

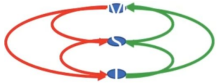

Protocol Low-level

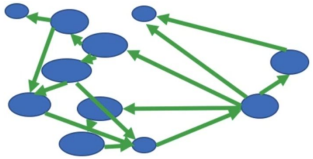

Verification

Simulation

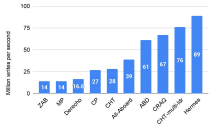

Implementation

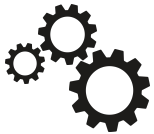

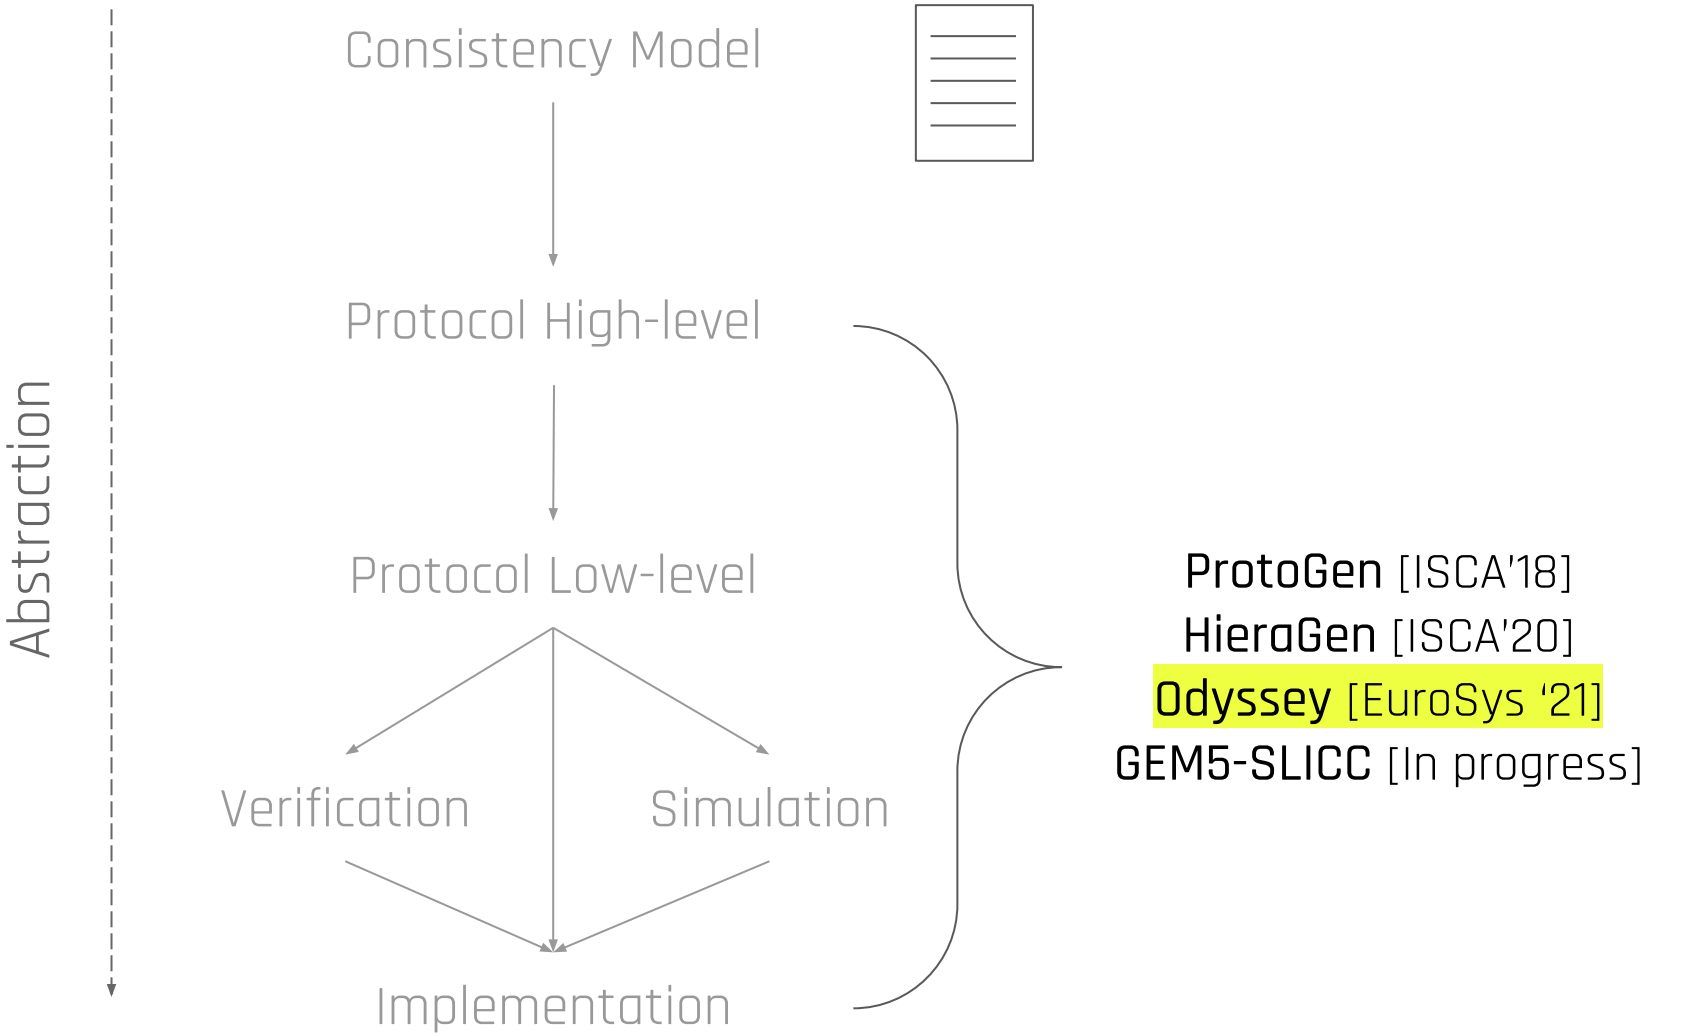
Abstraction
Consistency Model
Protocol High-level
Protocol Low-level
Verification
Simulation
Implementation
ProtoGen [ISCA'18]
HieraGen [ISCA'20]
Odyssey [EuroSys '21]
GEM5-SLICC [In progress]

Abstraction

Consistency Model

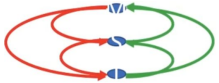

Protocol High-level

Protocol Low-level

Verification

Simulation

Implementation

ProtoGen [ISCA'18]
HieraGen [ISCA'20]
Odyssey [EuroSys '21]
GEM5-SLICC [In progress]

Consistency Model

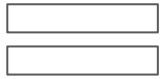

Synchronization Patterns

Alglave et. al [Herding Cats TOPLAS '14]

Consistency Model 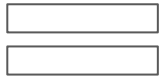 Synchronization Patterns

Protocol 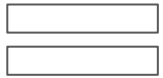 Real time orderings

Consistency Model

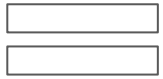

Synchronization Patterns

Mapping

Protocol

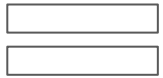

Real time orderings

Consistency Model

Synchronization Patterns

Mapping

Protocol

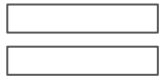

Real time orderings

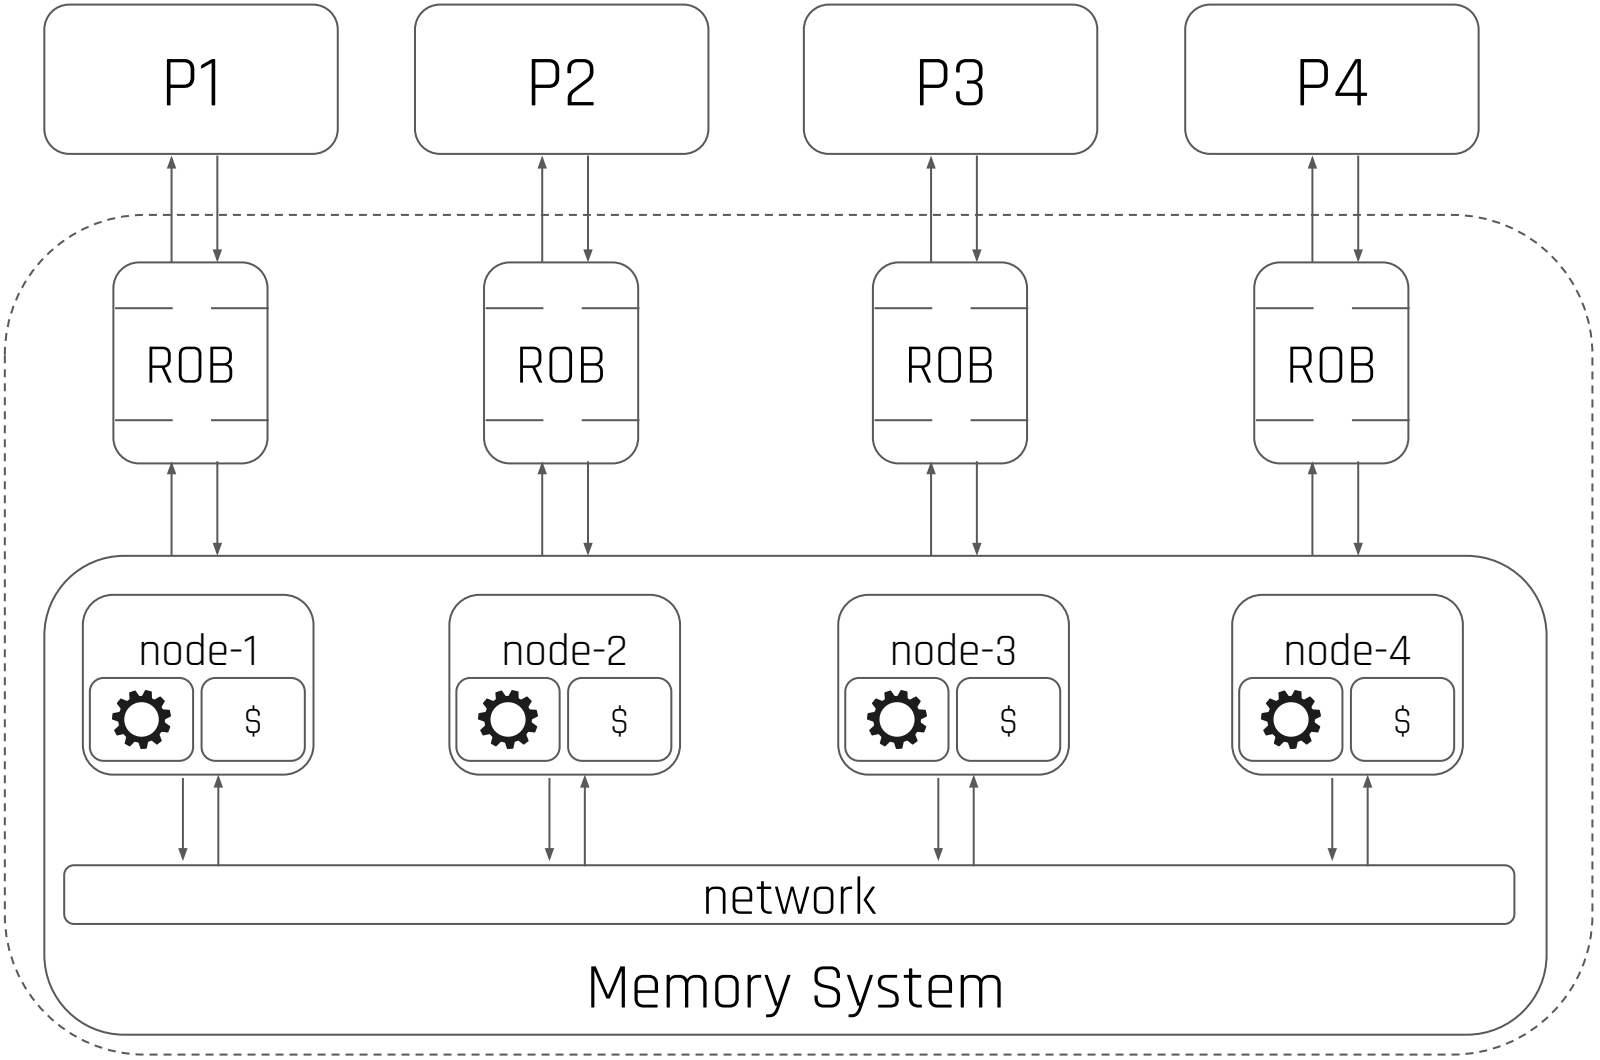
P1
P2
P3
P4
ROB
ROB
ROB
ROB
node-1
$
node-2
$
node-3
$
node-4
$
network
Memory System

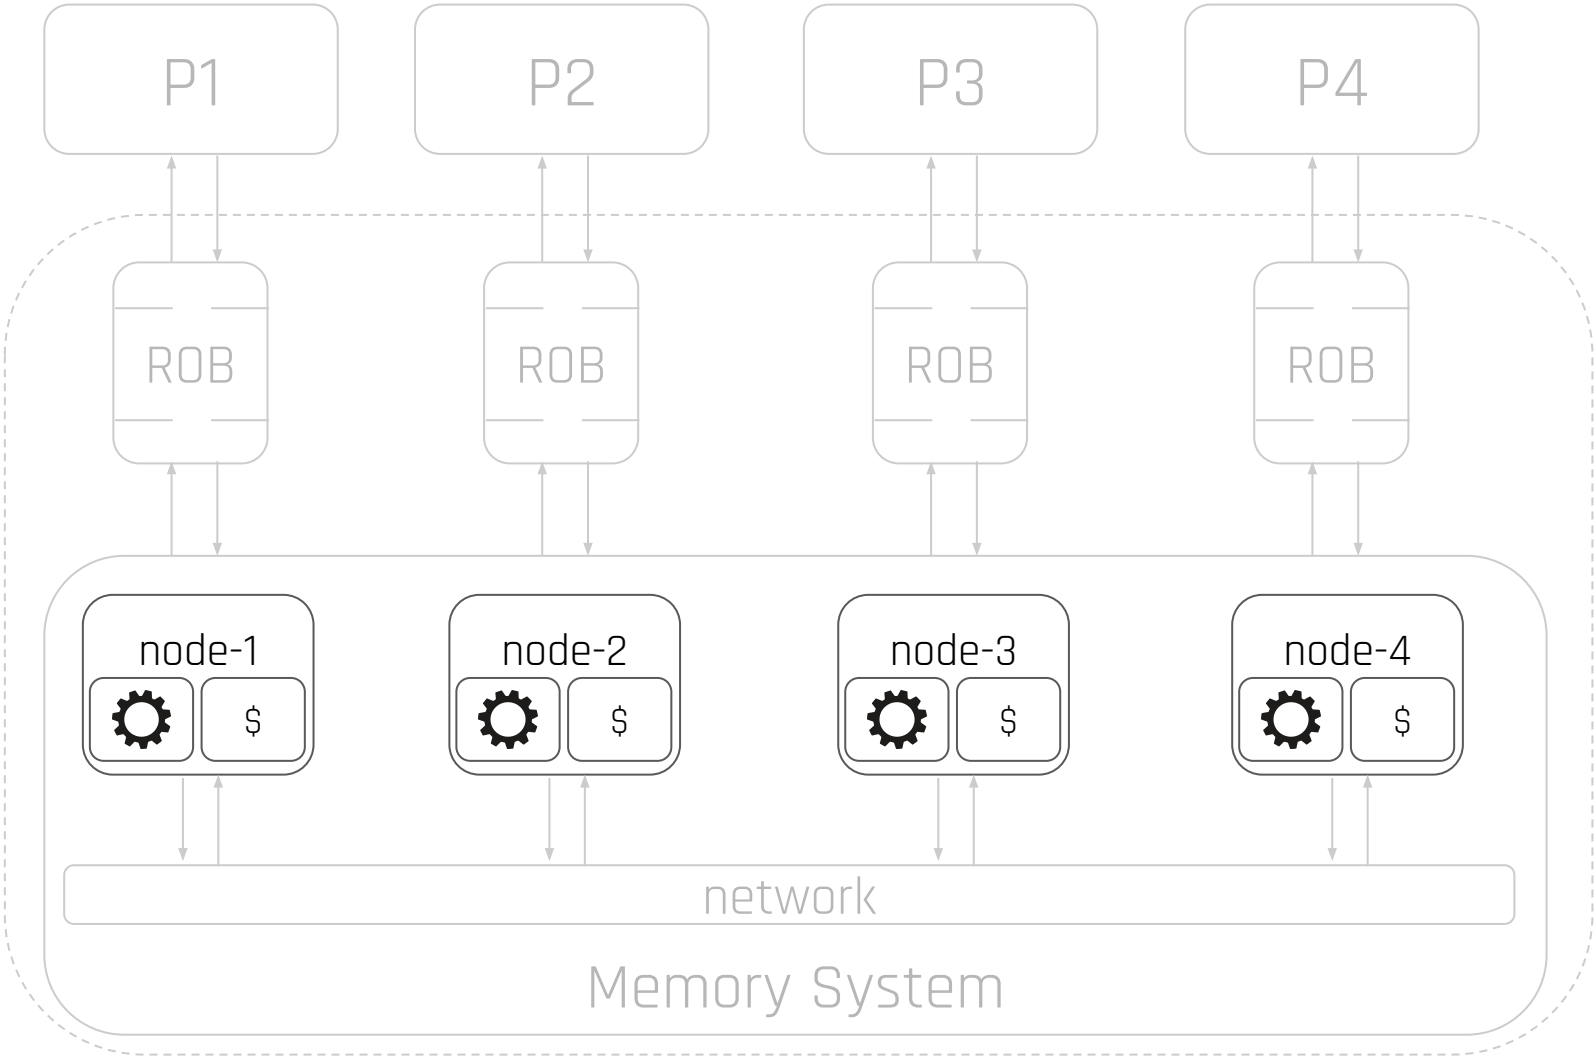
P1
P2
P3
P4
ROB
ROB
ROB
ROB
node-1
$
node-2
$
node-3
$
node-4
$
network
Memory System

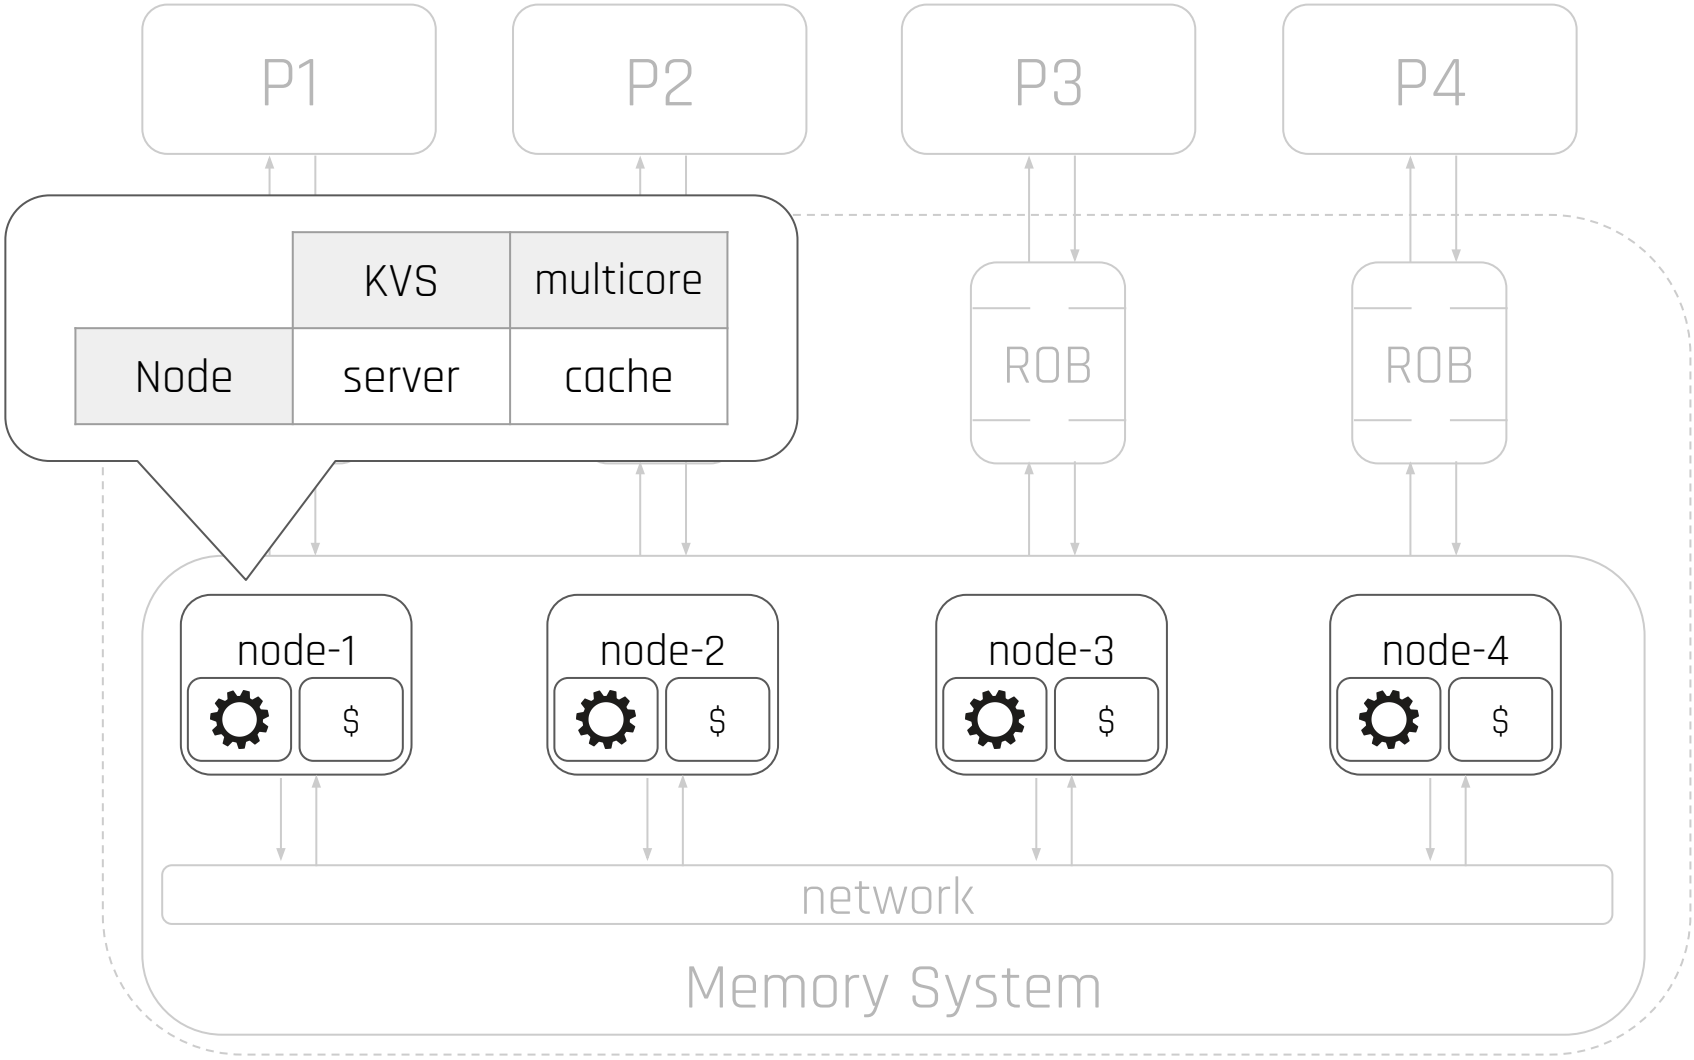
P1
P2
P3
P4
KVS
multicore
Node
server
cache
ROB
ROB
node-1
$
node-2
$
node-3
$
node-4
$
network
Memory System

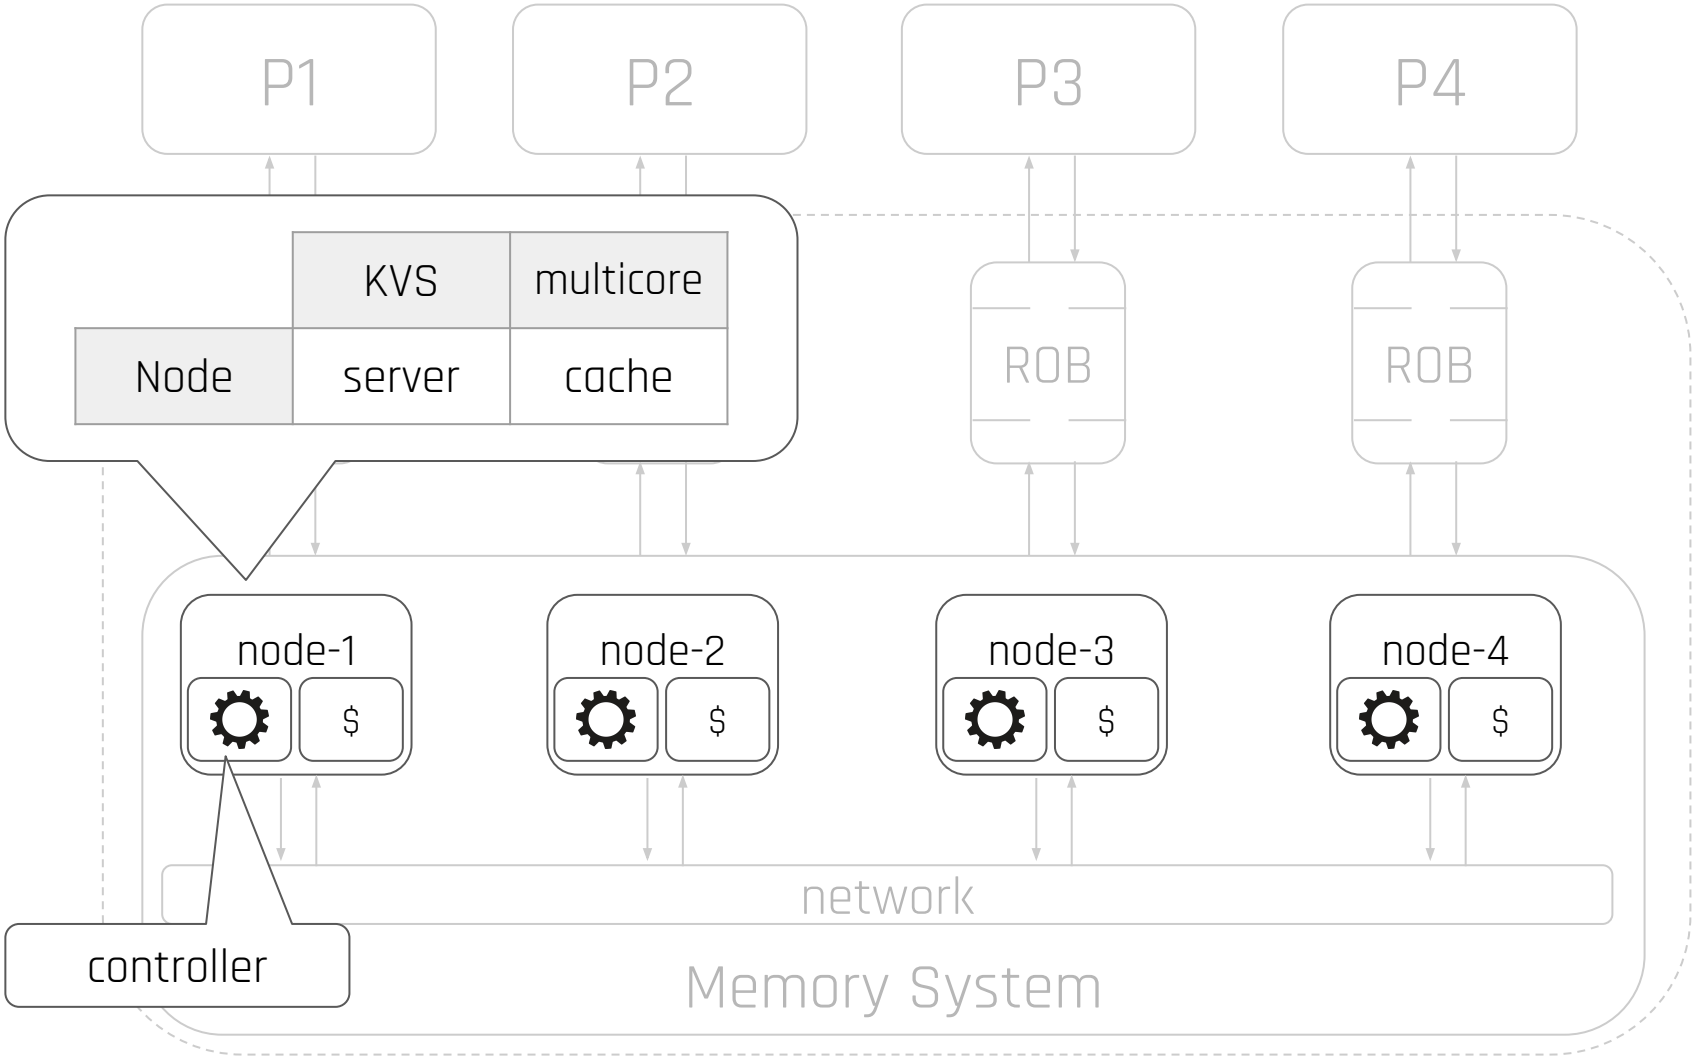
P1
P2
P3
P4
KVS
multicore
Node
server
cache
ROB
ROB
node-1
node-2
node-3
node-4
$
$
$
$
network
controller
Memory System

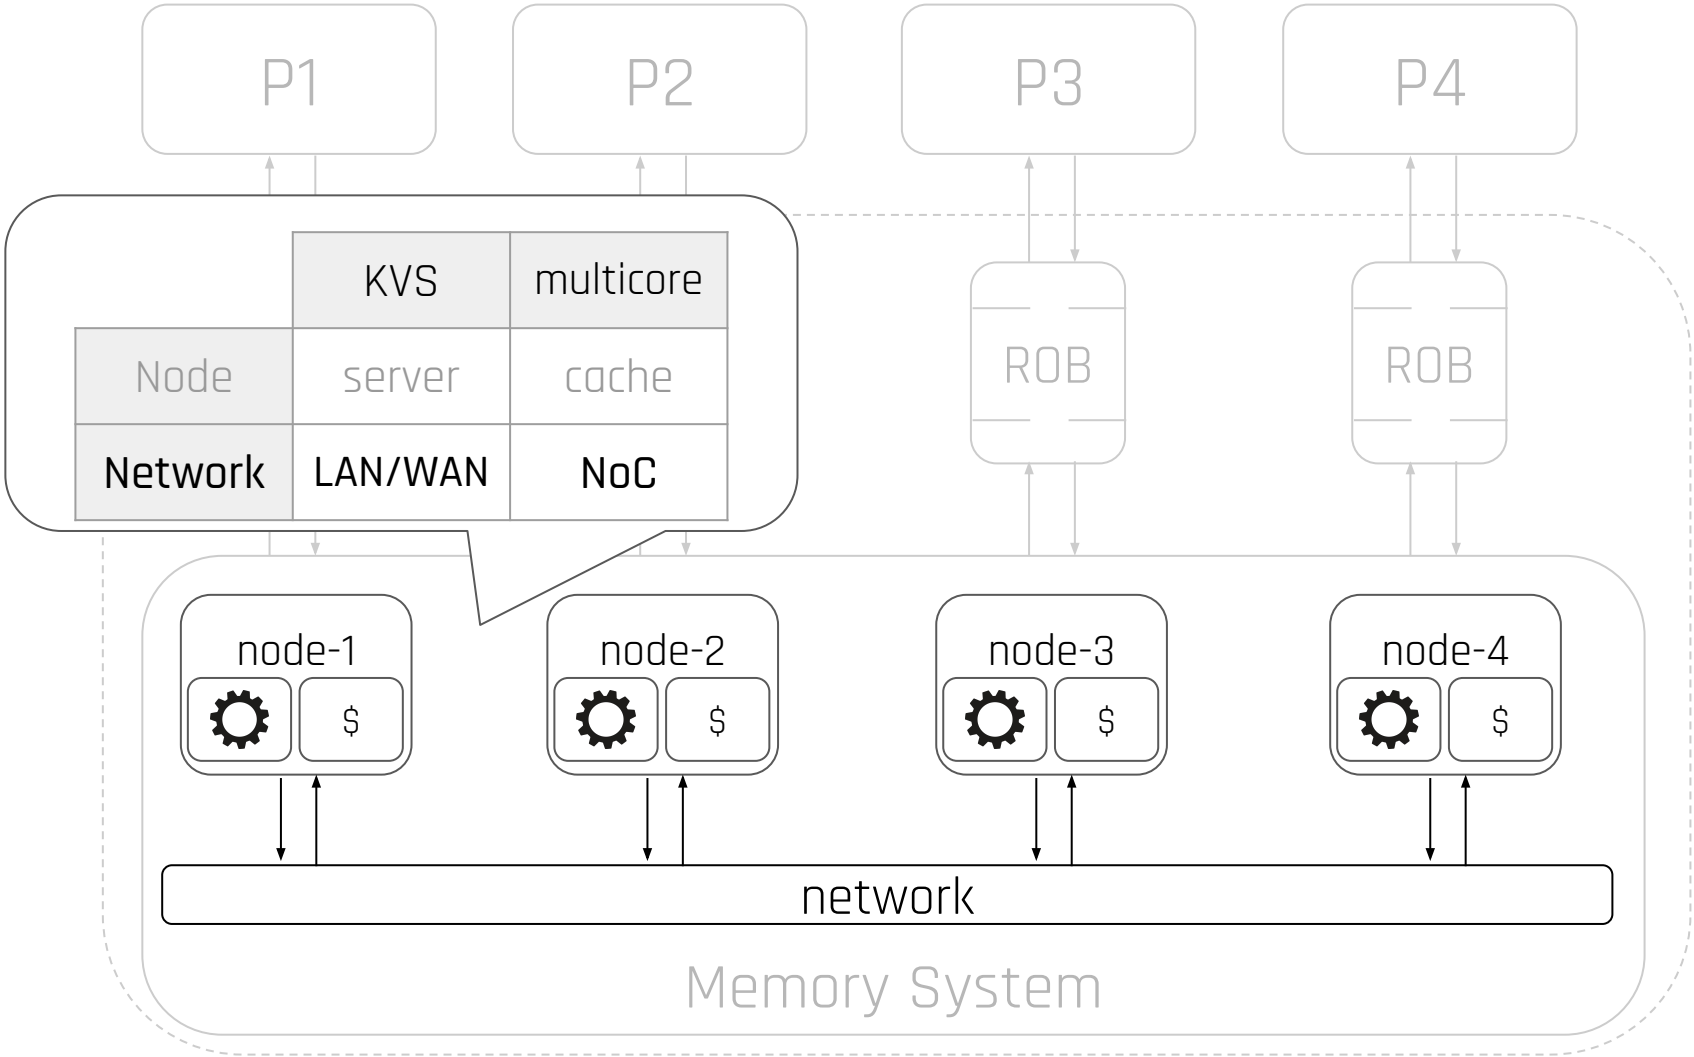
P1
P2
P3
P4
KVS
multicore
Node
server
cache
Network
LAN/WAN
NoC
ROB
ROB
node-1
$
node-2
$
node-3
$
node-4
$
network
Memory System

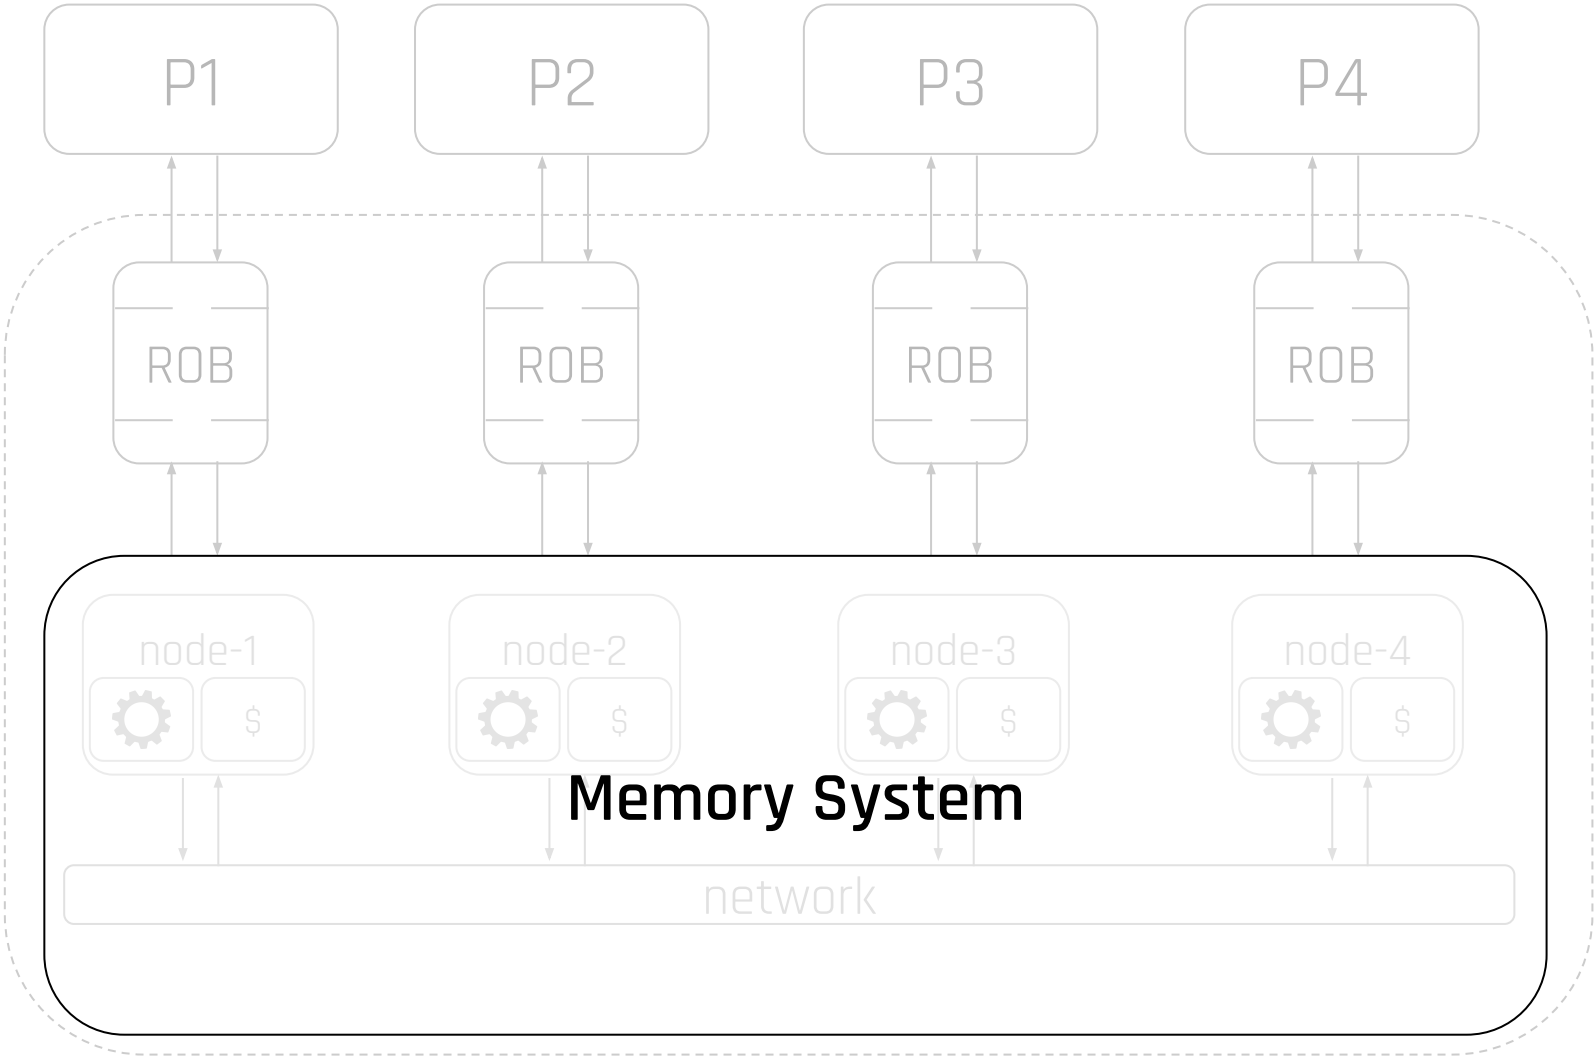
P1
P2
P3
P4
ROB
ROB
ROB
ROB
node-1
$
node-2
$
node-3
$
node-4
$
Memory System
network

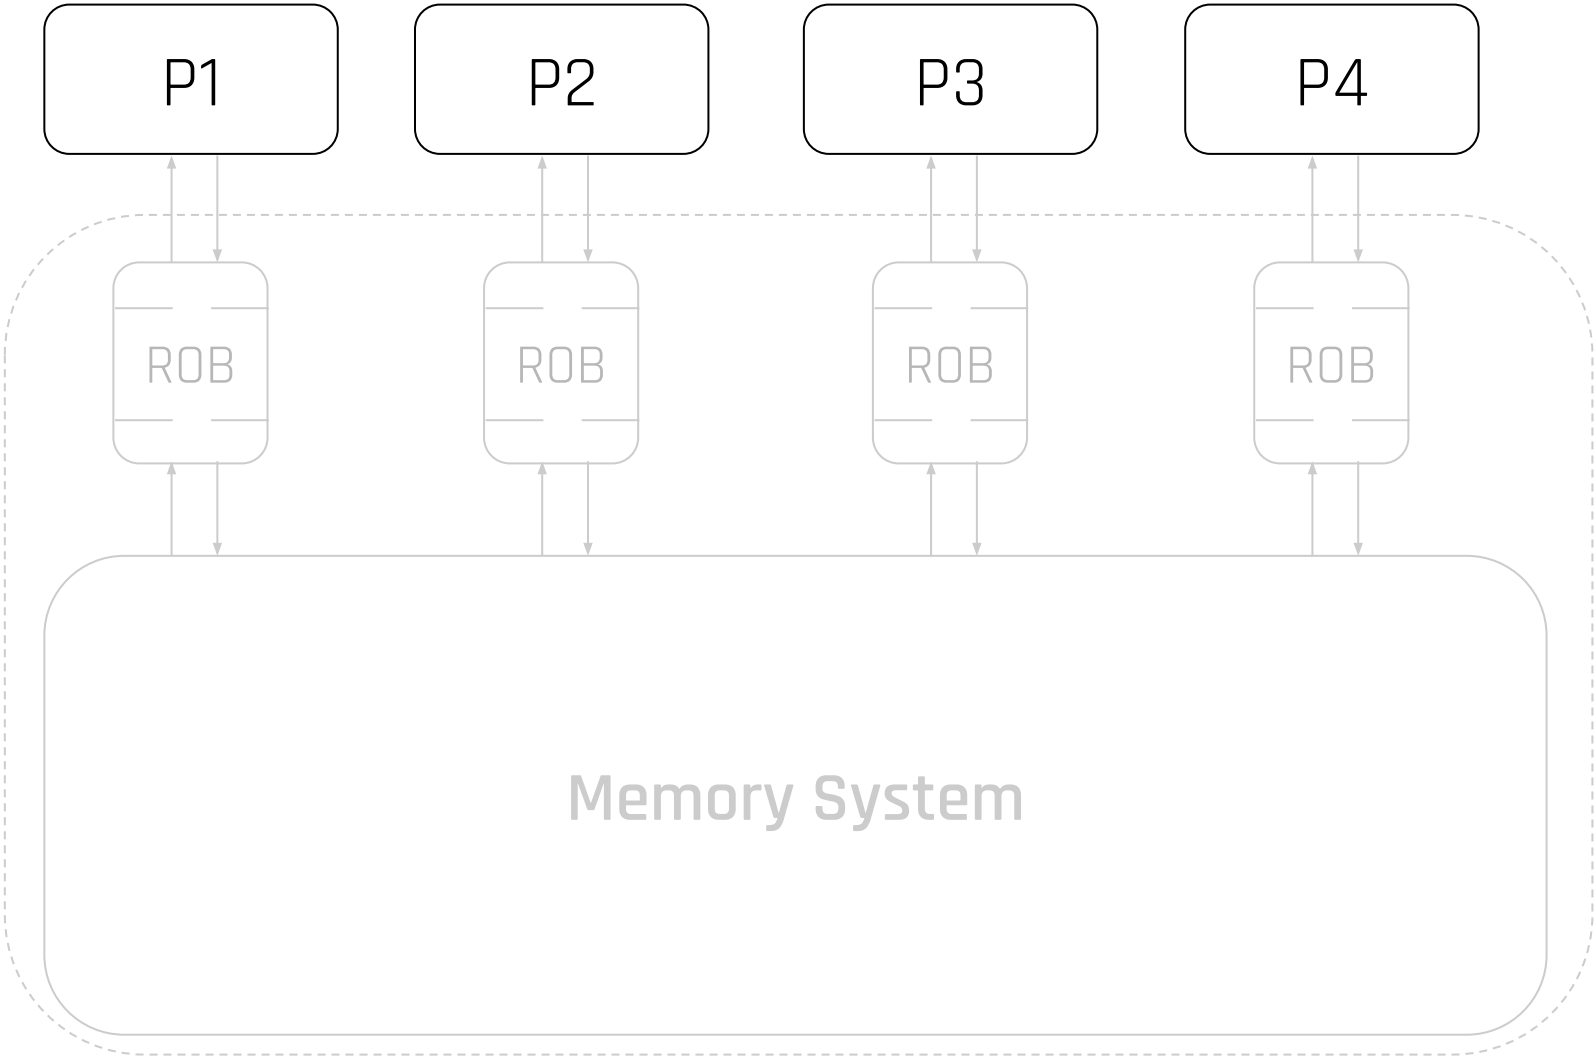
P1
P2
P3
P4
ROB
ROB
ROB
ROB
Memory System

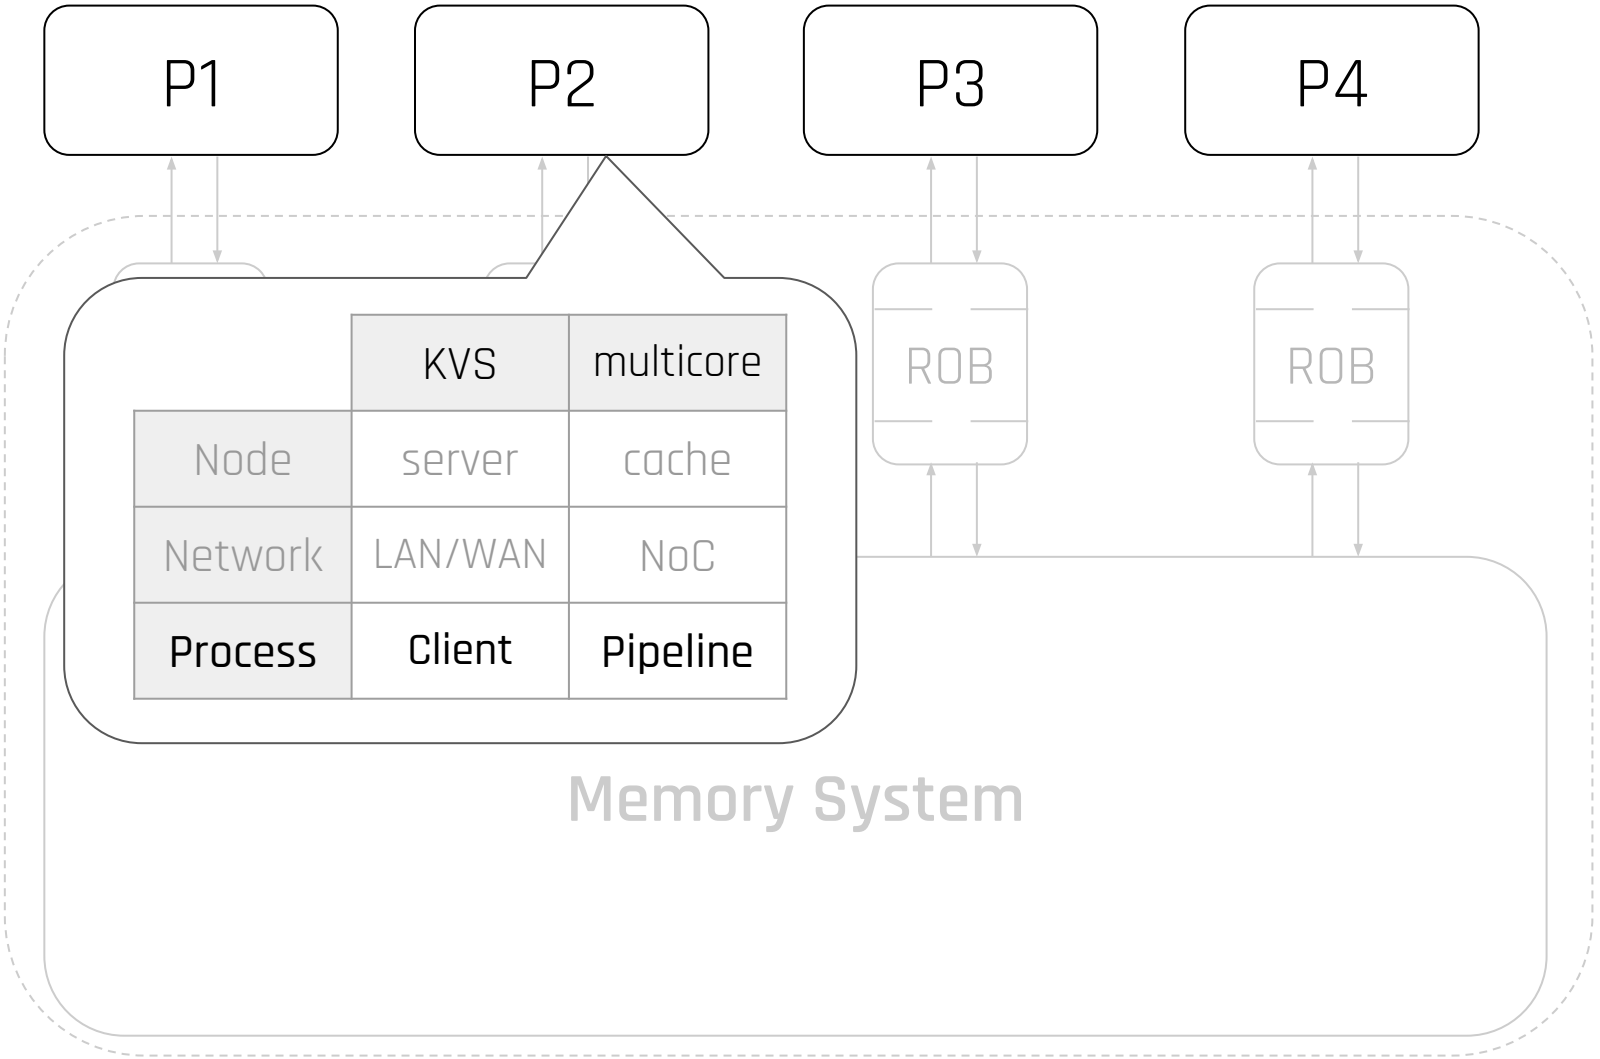
P1
P2
P3
P4
ROB
ROB
KVS
multicore
Node
server
cache
Network
LAN/WAN
NoC
Process
Client
Pipeline
Memory System

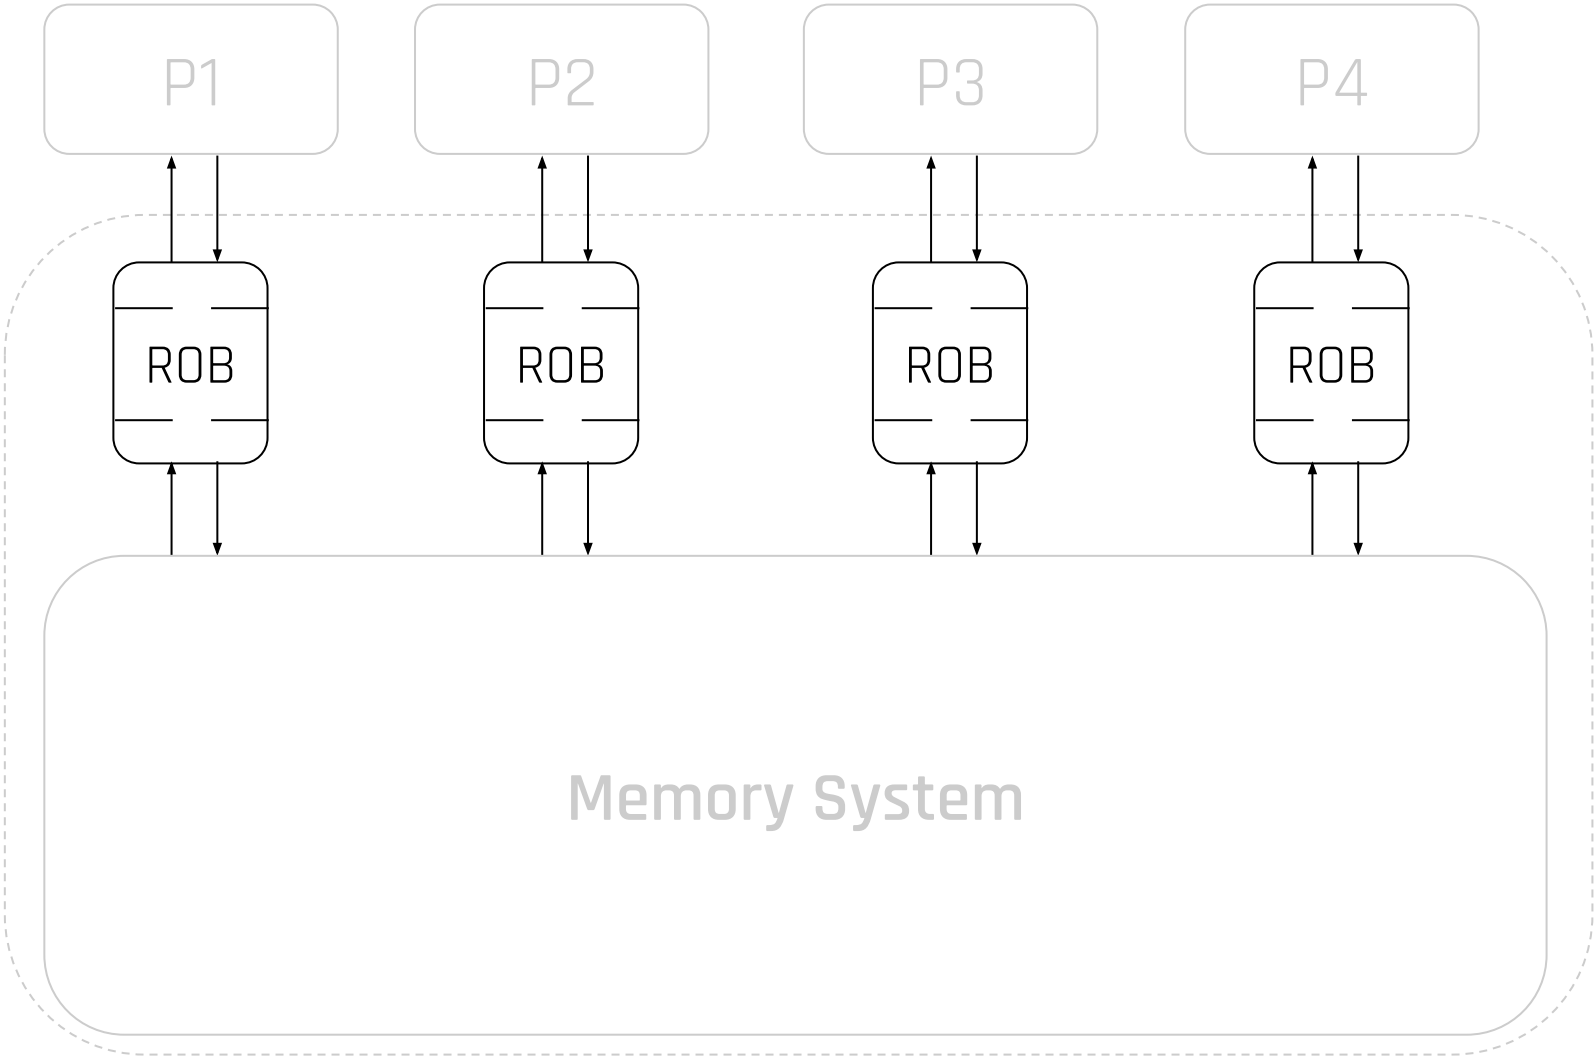
P1
P2
P3
P4
ROB
ROB
ROB
ROB
Memory System

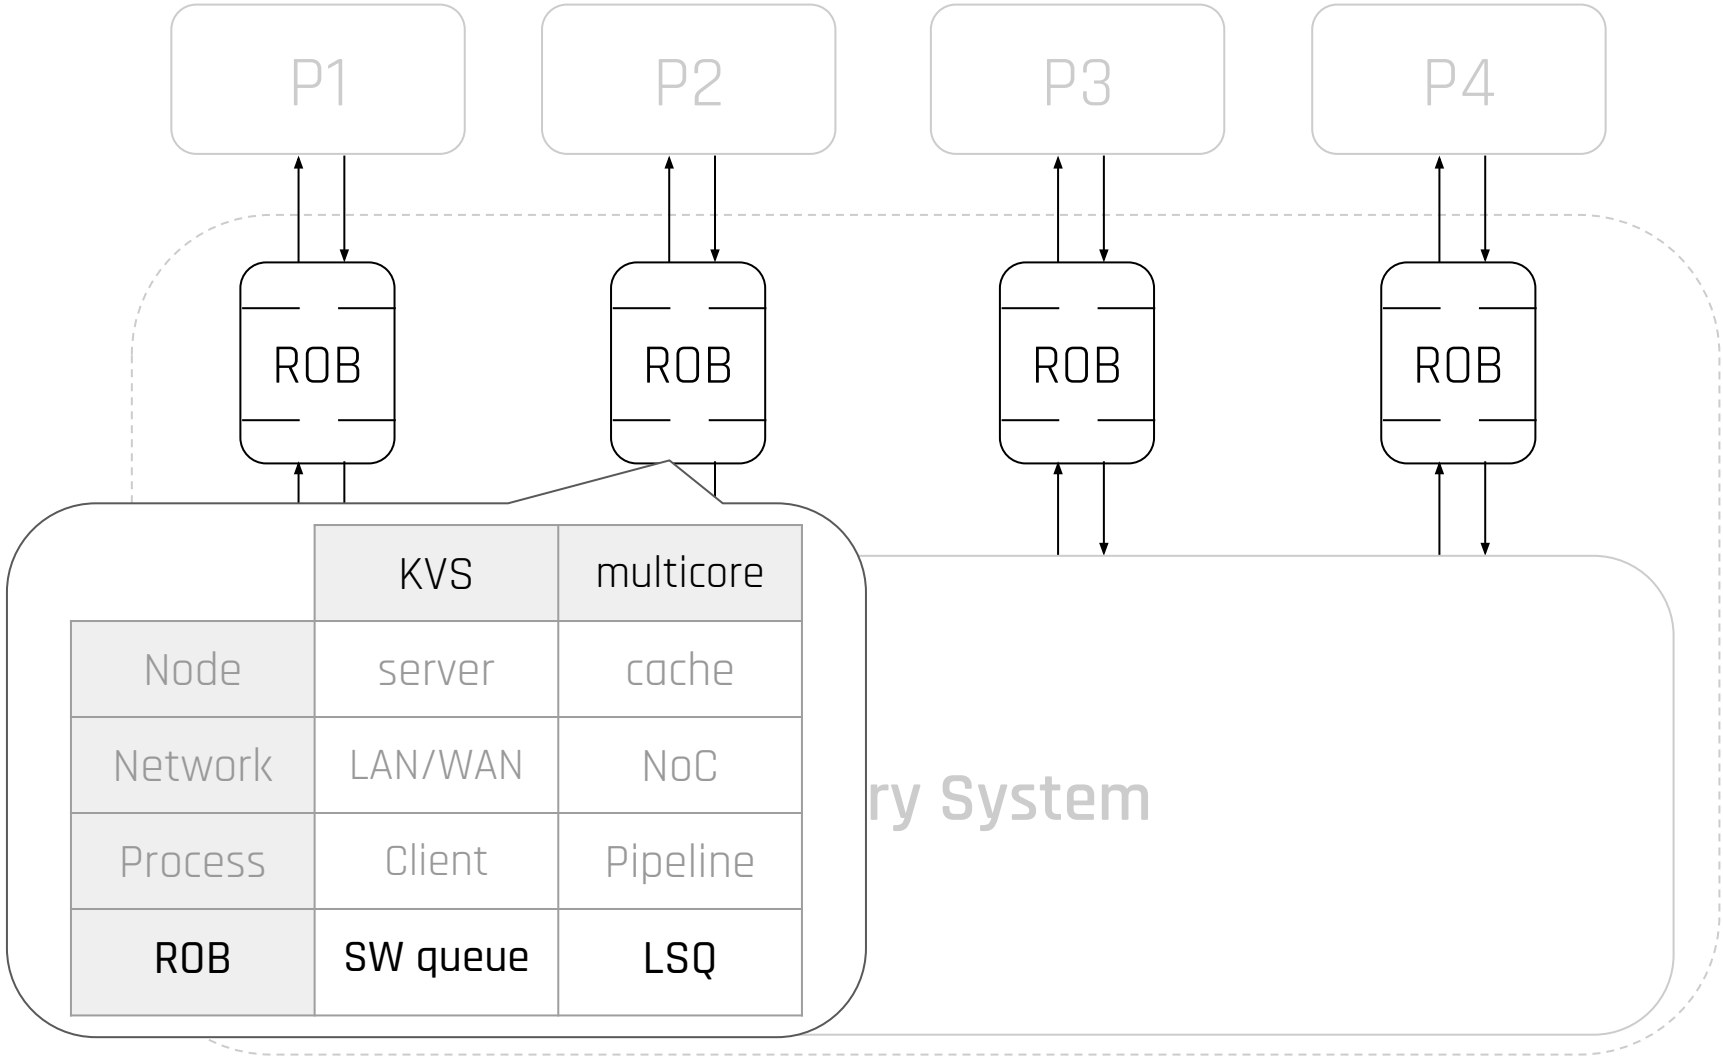

	KVS	multicore
Node	server	cache
Network	LAN/WAN	NoC
Process	Client	Pipeline
ROB	**SW queue**	**LSQ**

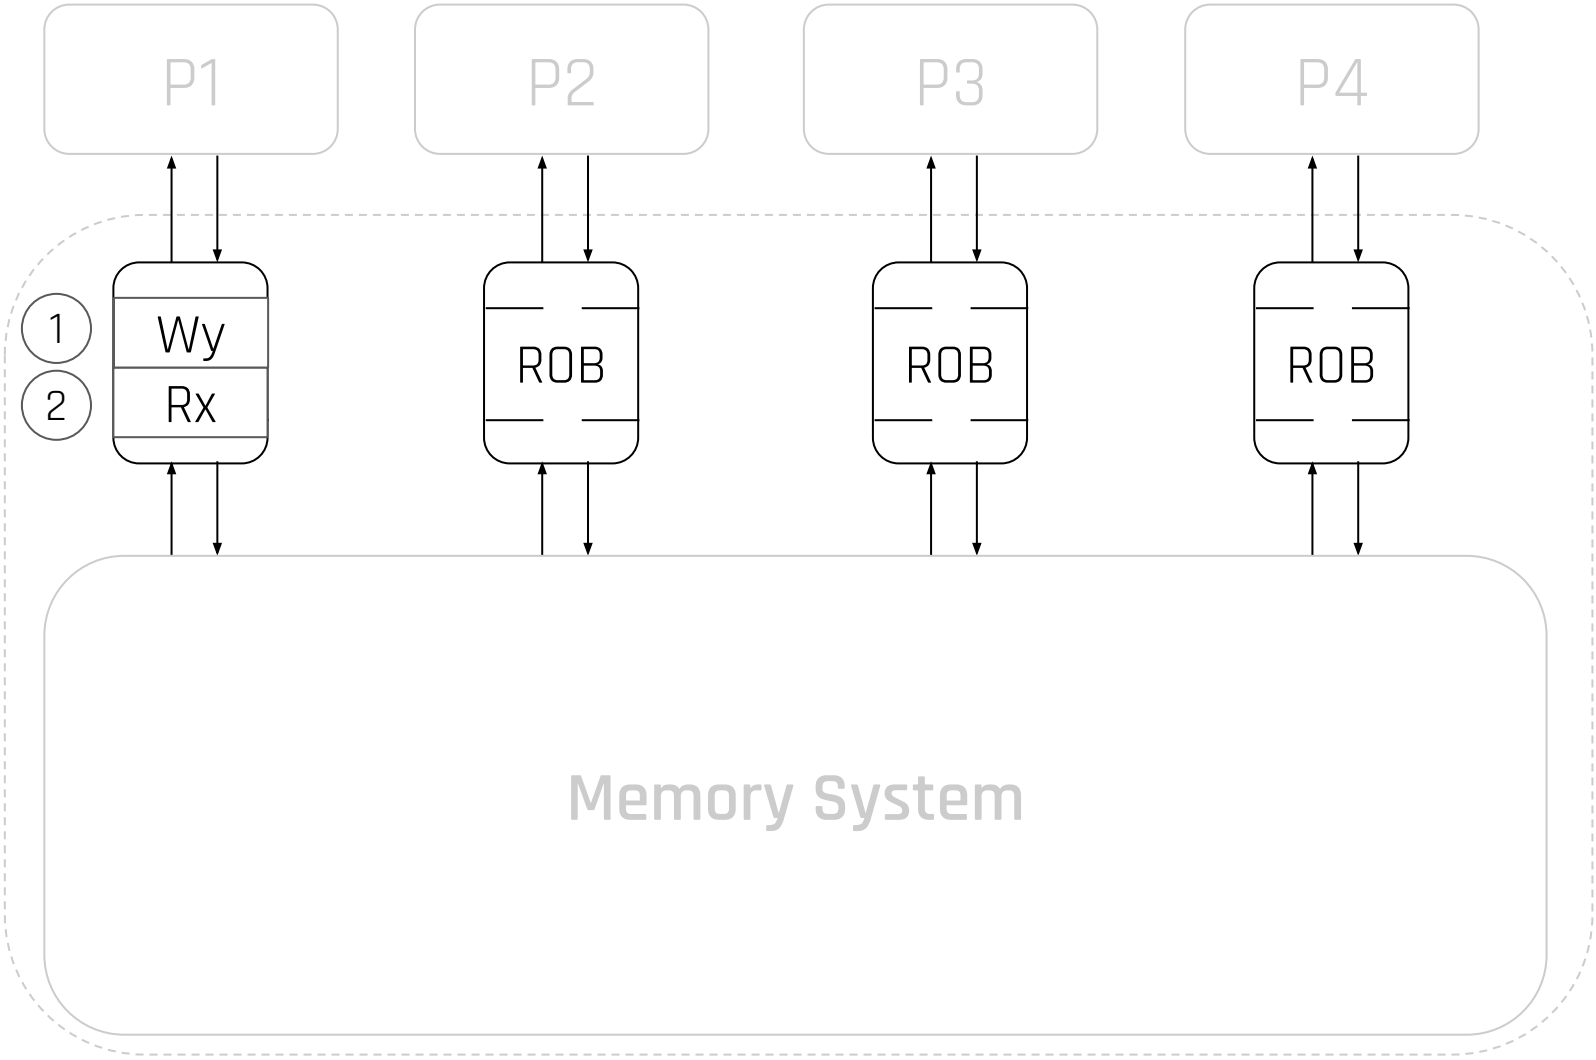
P1
P2
P3
P4
1
2
Wy
Rx
ROB
ROB
ROB
Memory System

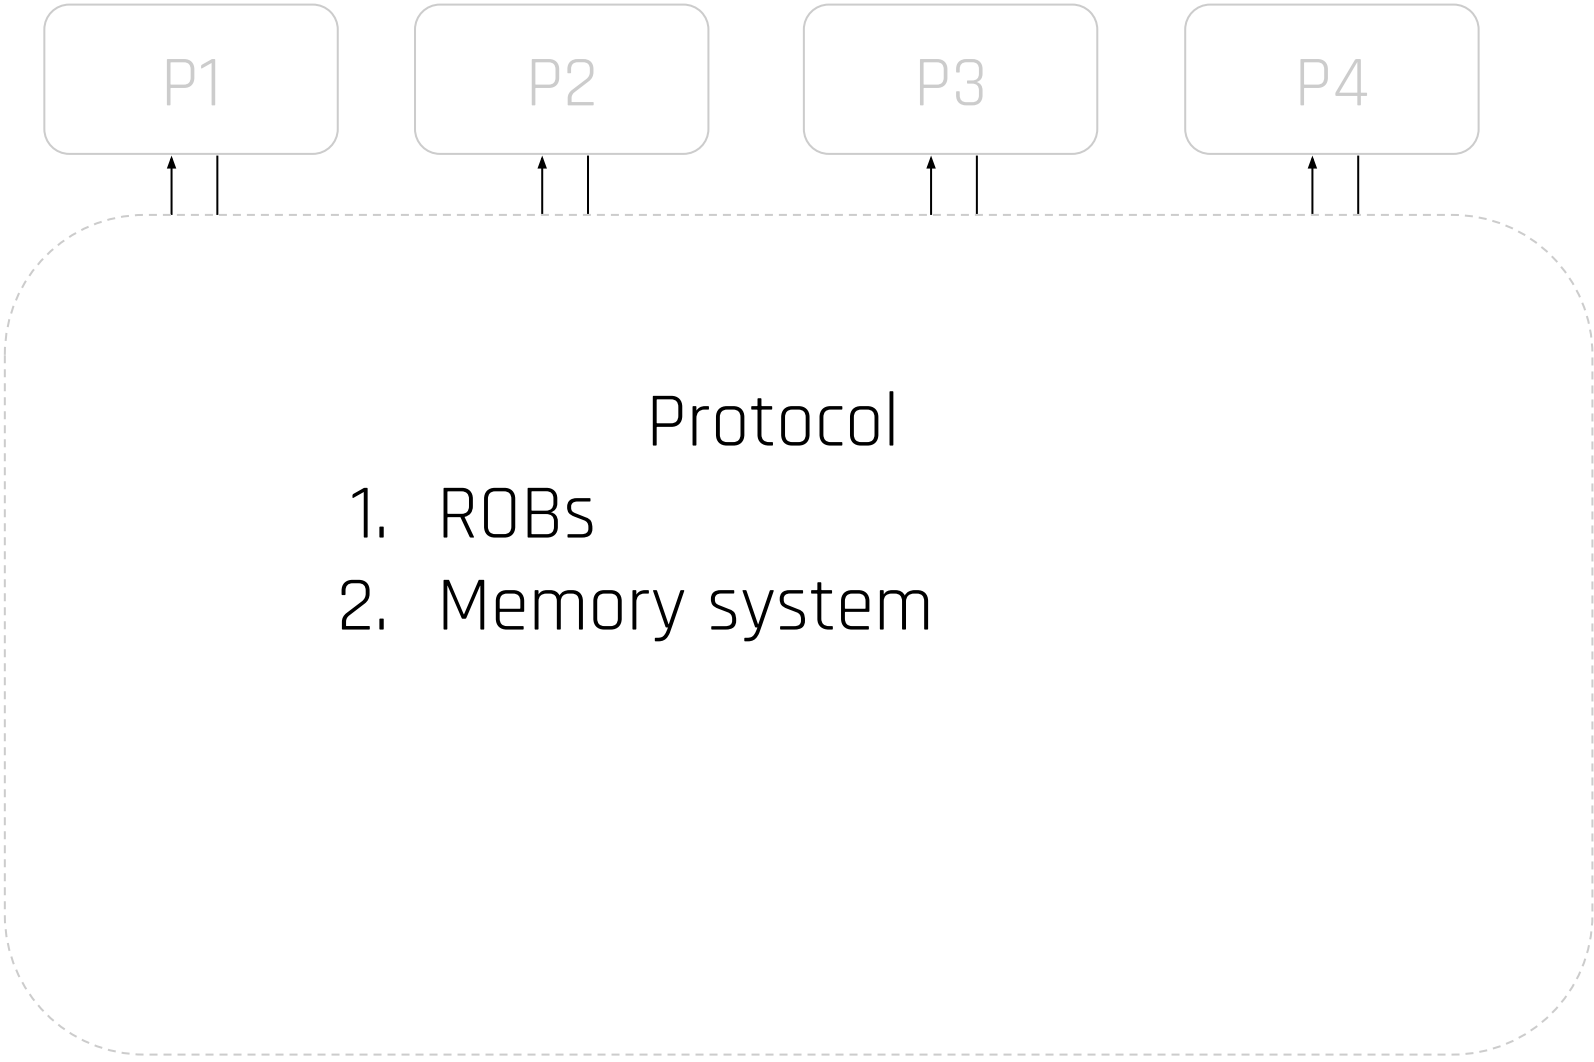
P1
P2
P3
P4
Protocol
1. ROBs
2. Memory system

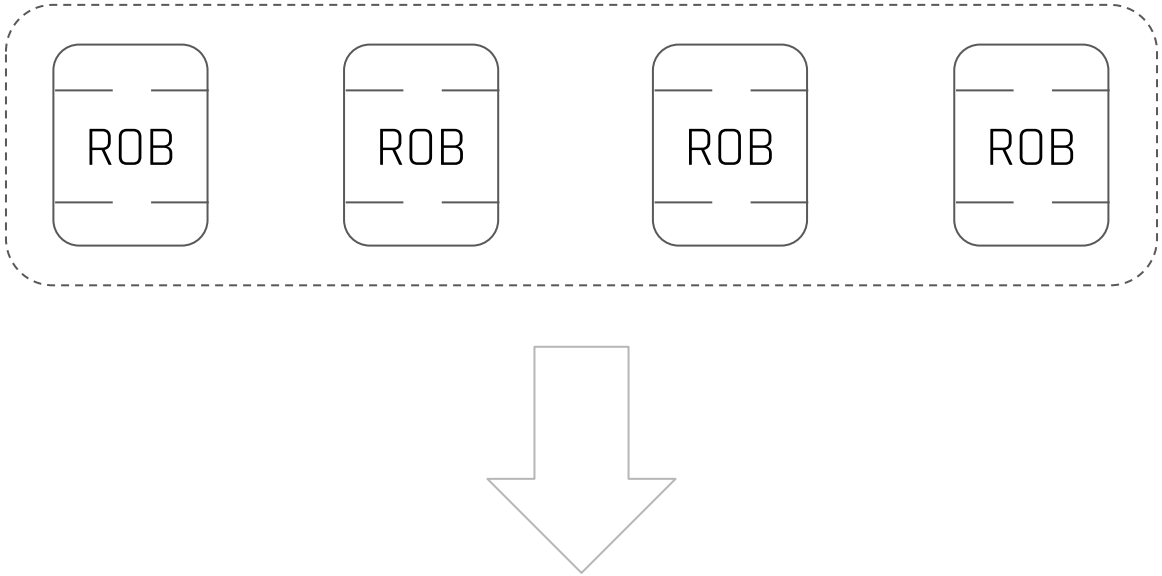

Program-order **R**eal **T**ime orderings

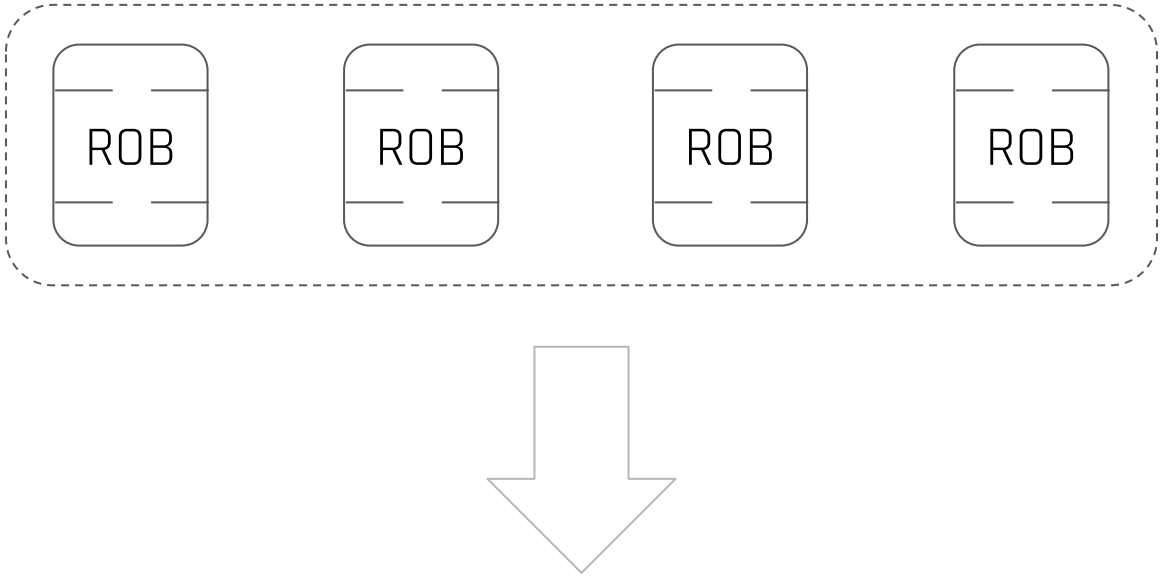
ROB
ROB
ROB
ROB
Prt-orderings

1. prt_{WR} (W⟶ R)

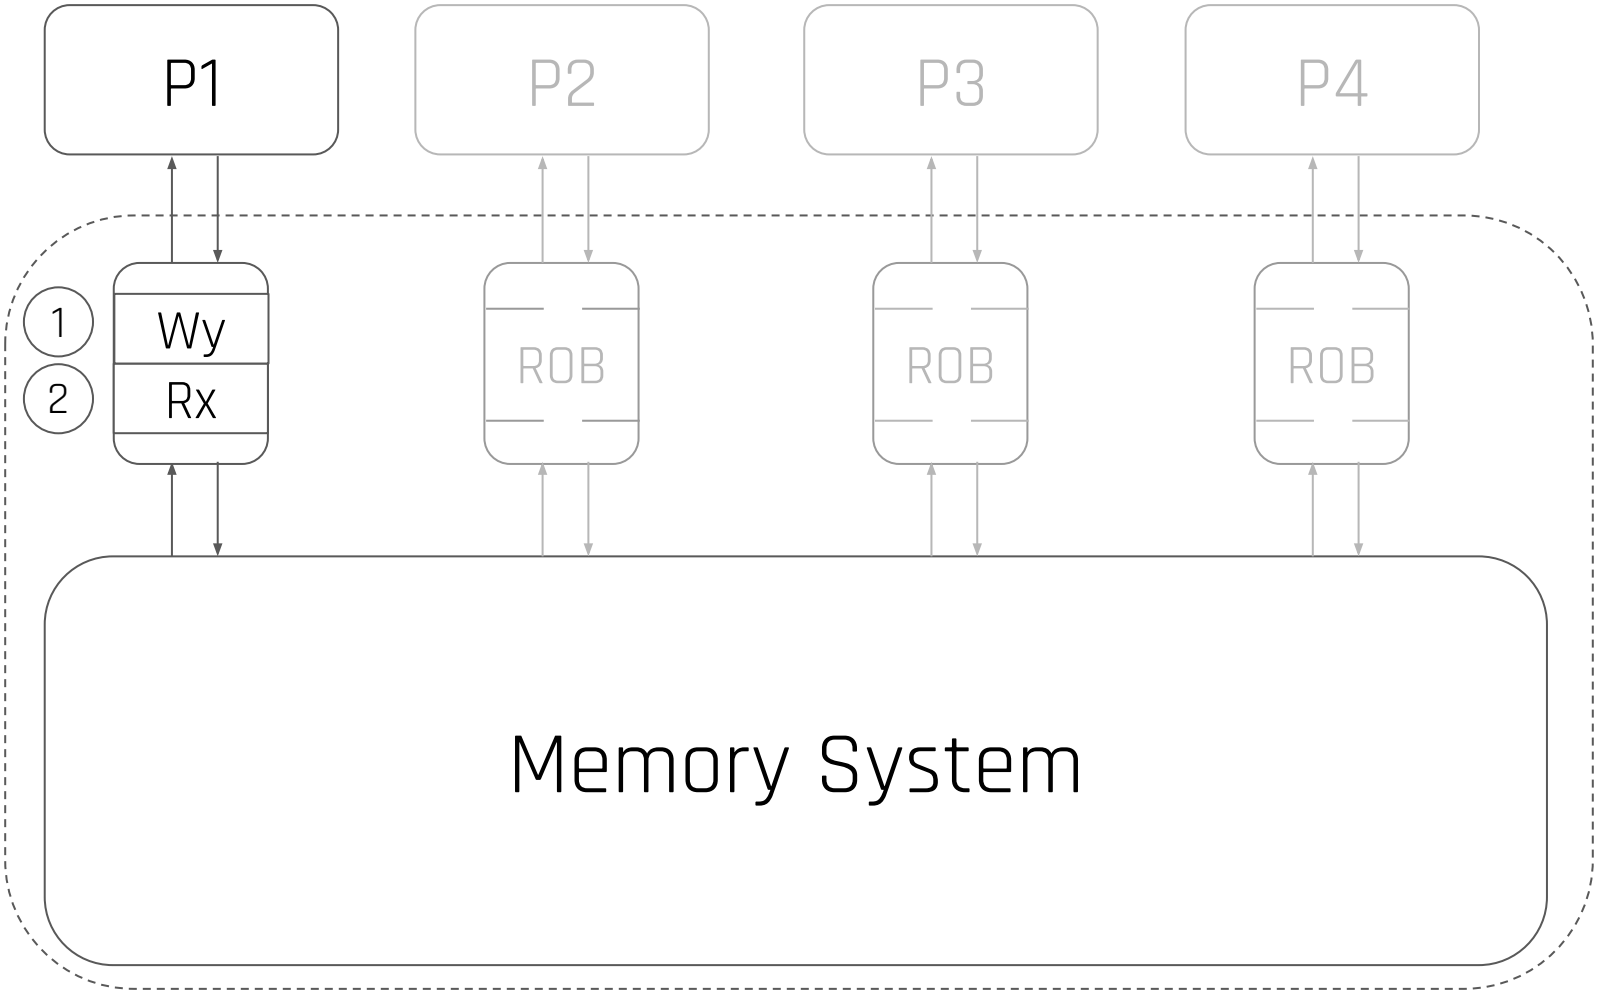
P1
P2
P3
P4
1
2
Wy
Rx
ROB
ROB
ROB
Memory System

1. prt_{WR} (W⟶ R)
2. prt_{WW} (W⟶ W)
3. prt_{RR} (R⟶ R)
4. prt_{RW} (R⟶ W)

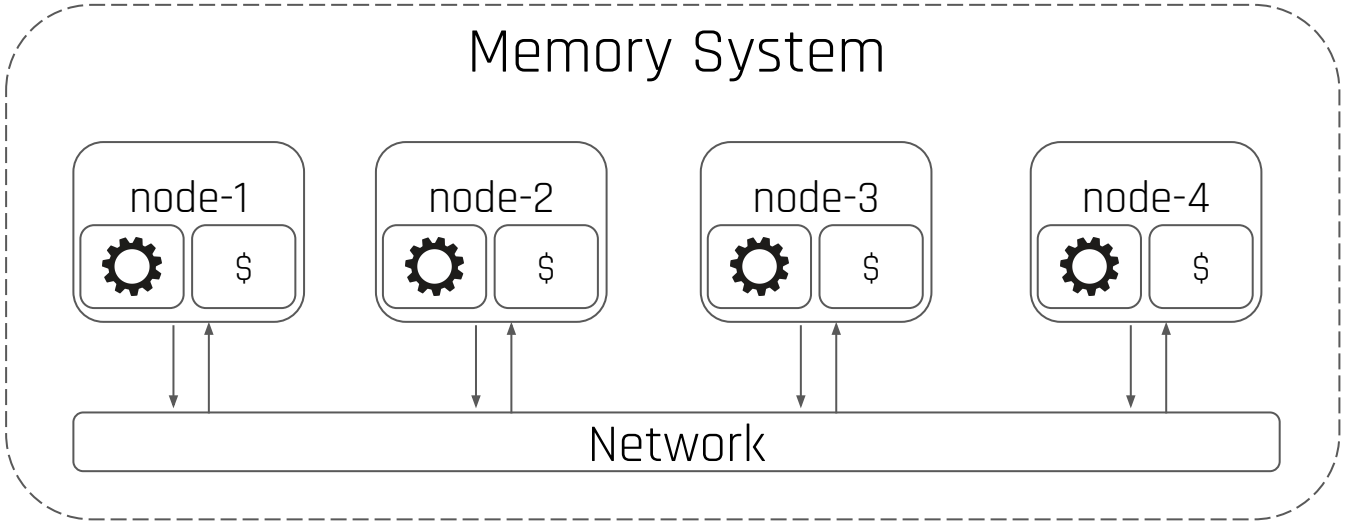

Synchronization **R**eal **T**ime orderings

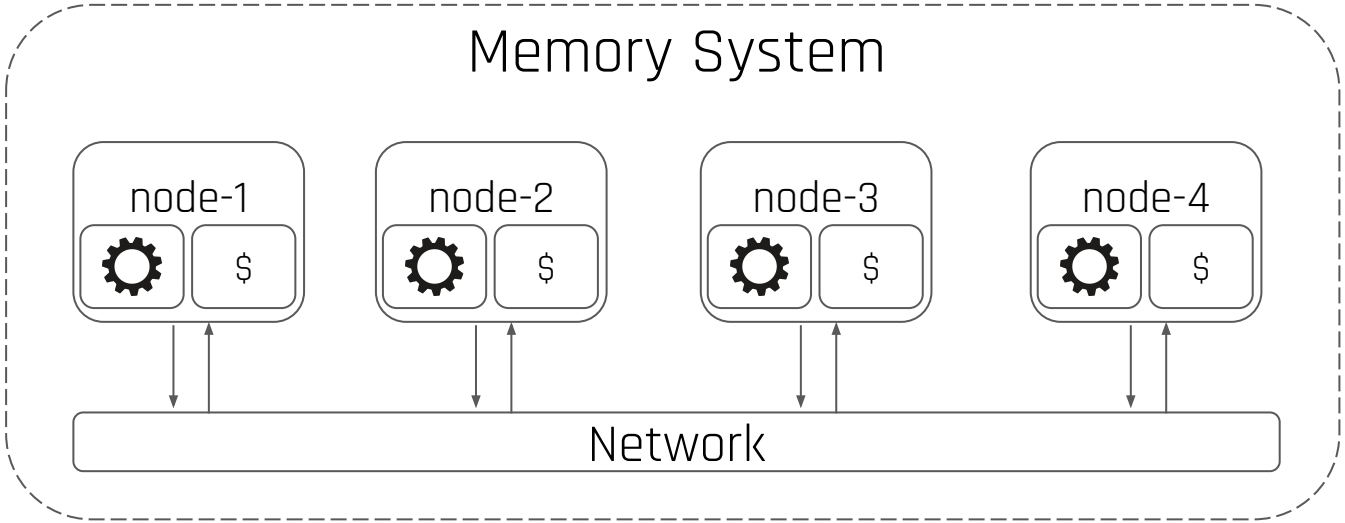

Srt-orderings

1. srt_{WR} (W⟶ R)

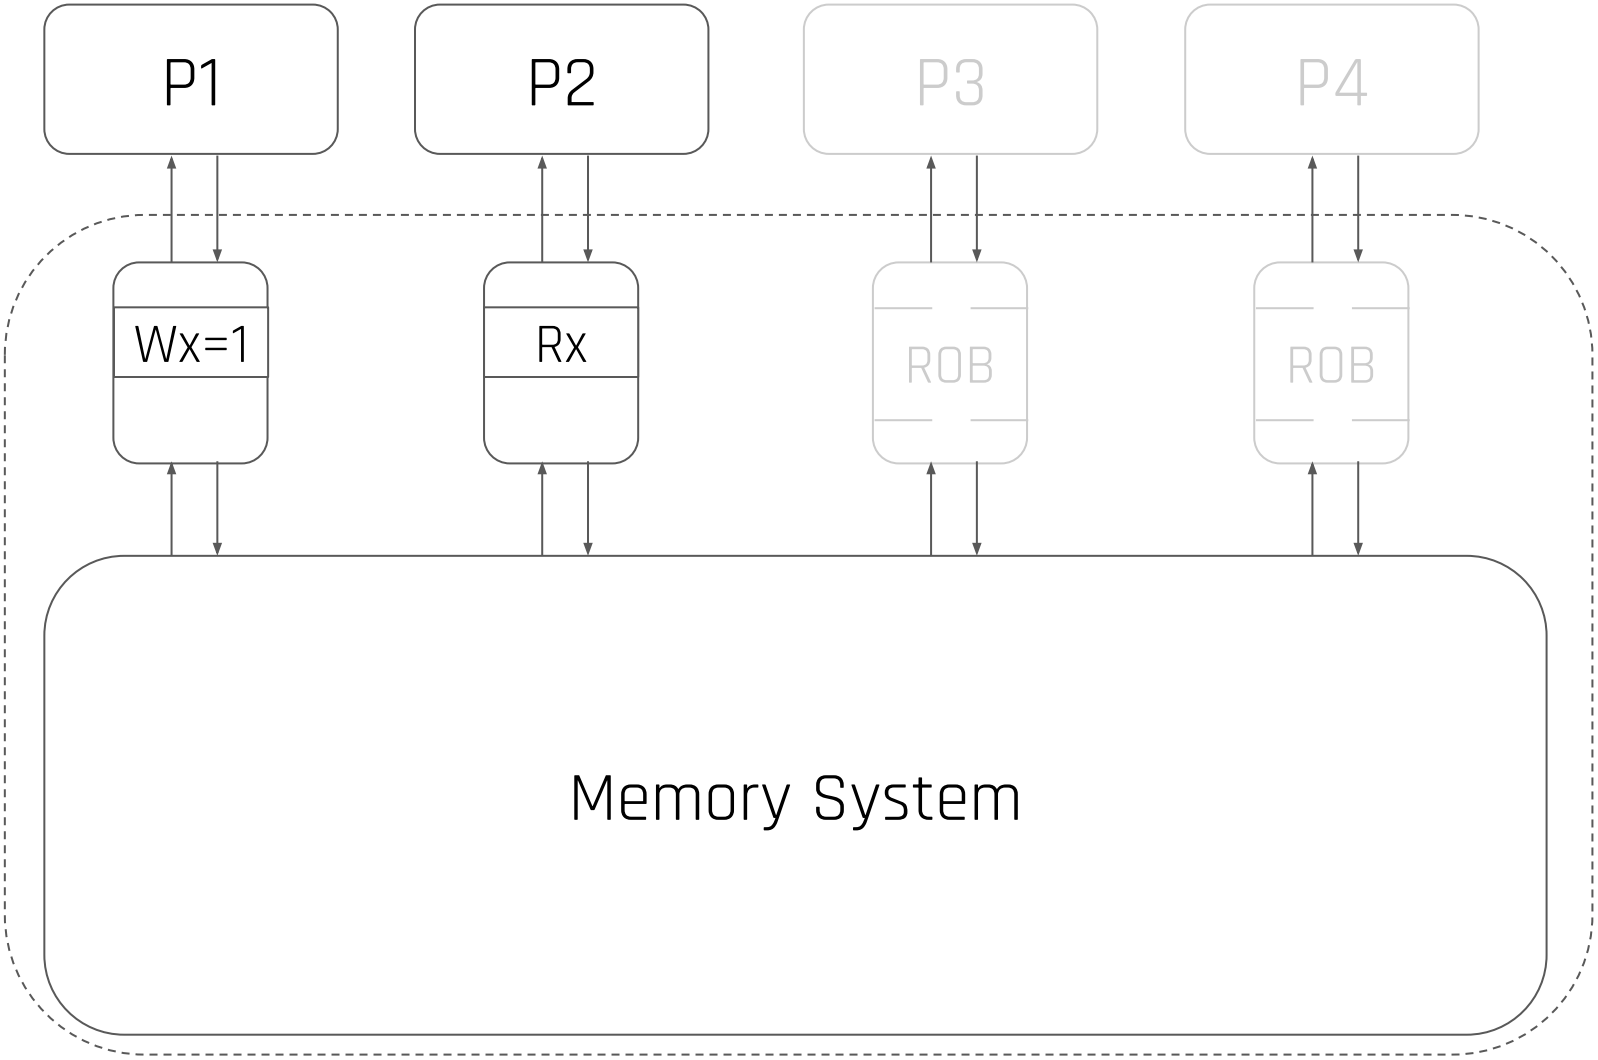
P1
P2
P3
P4
Wx=1
Rx
ROB
ROB
Memory System

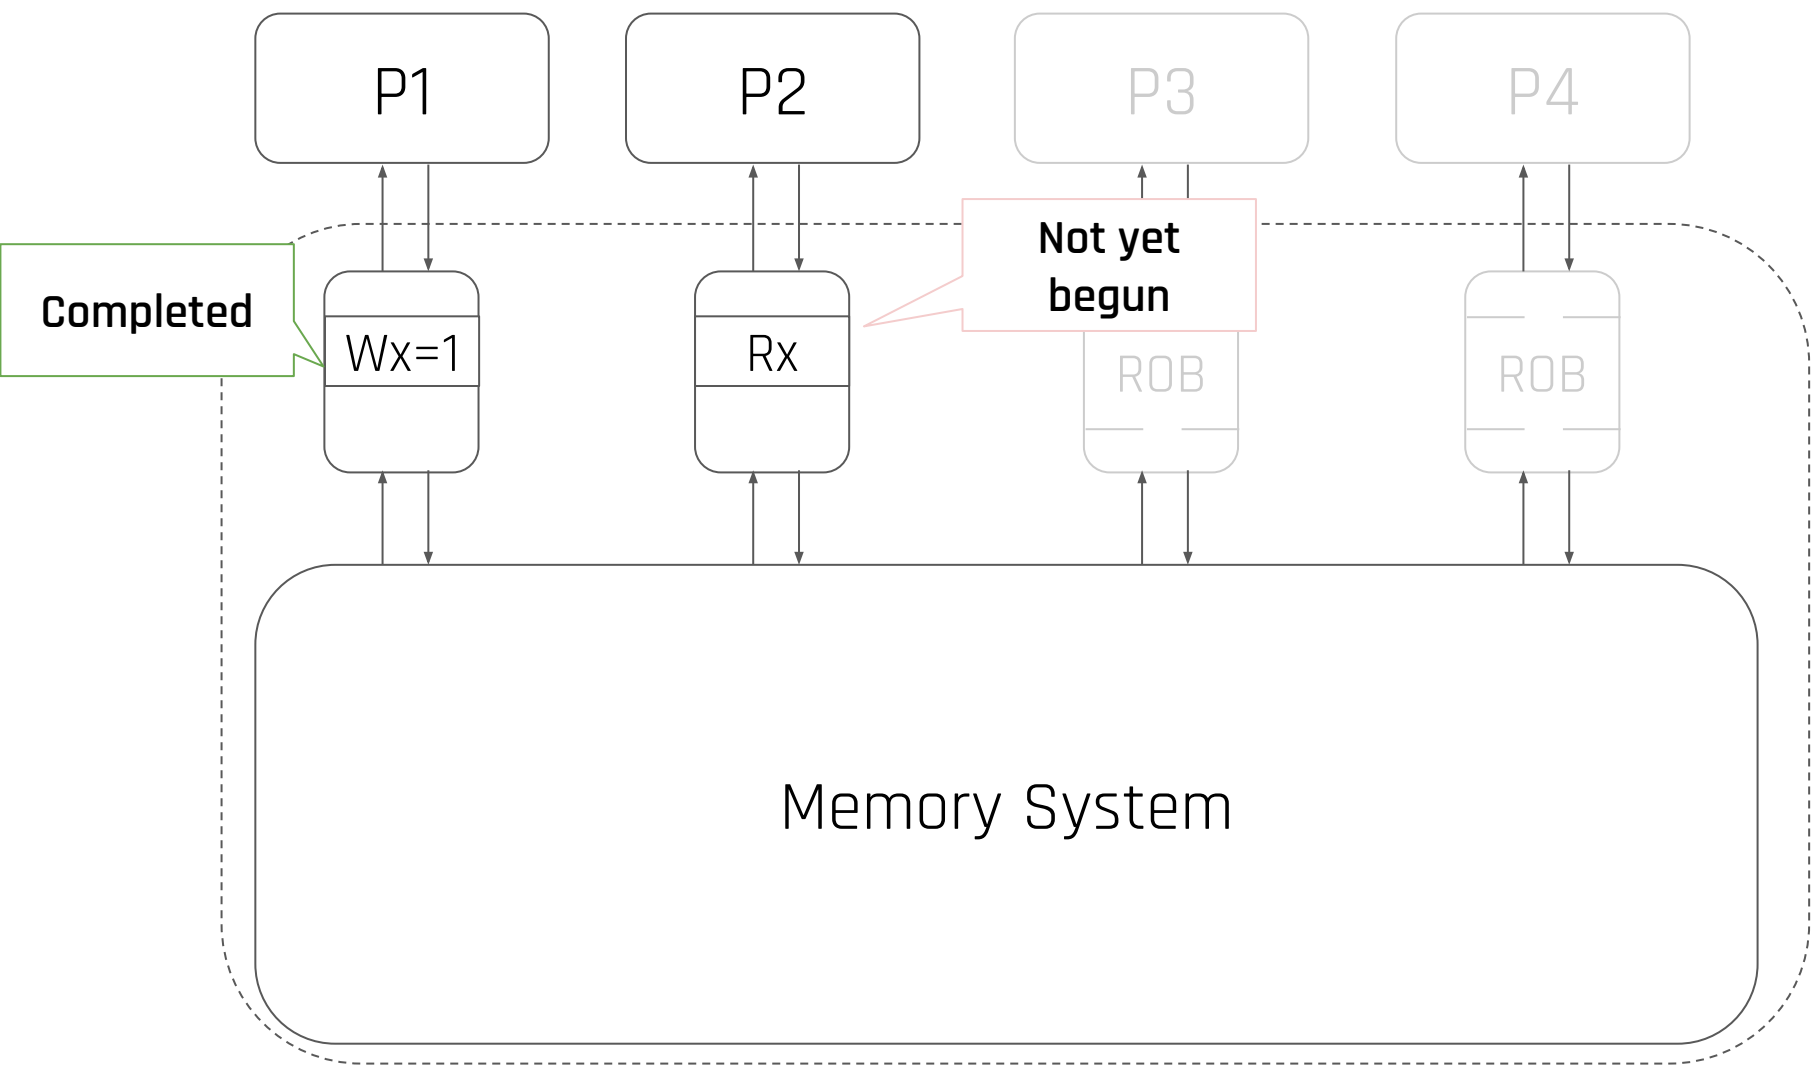
P1
P2
P3
P4
Completed
Wx=1
Rx
Not yet begun
ROB
ROB
Memory System

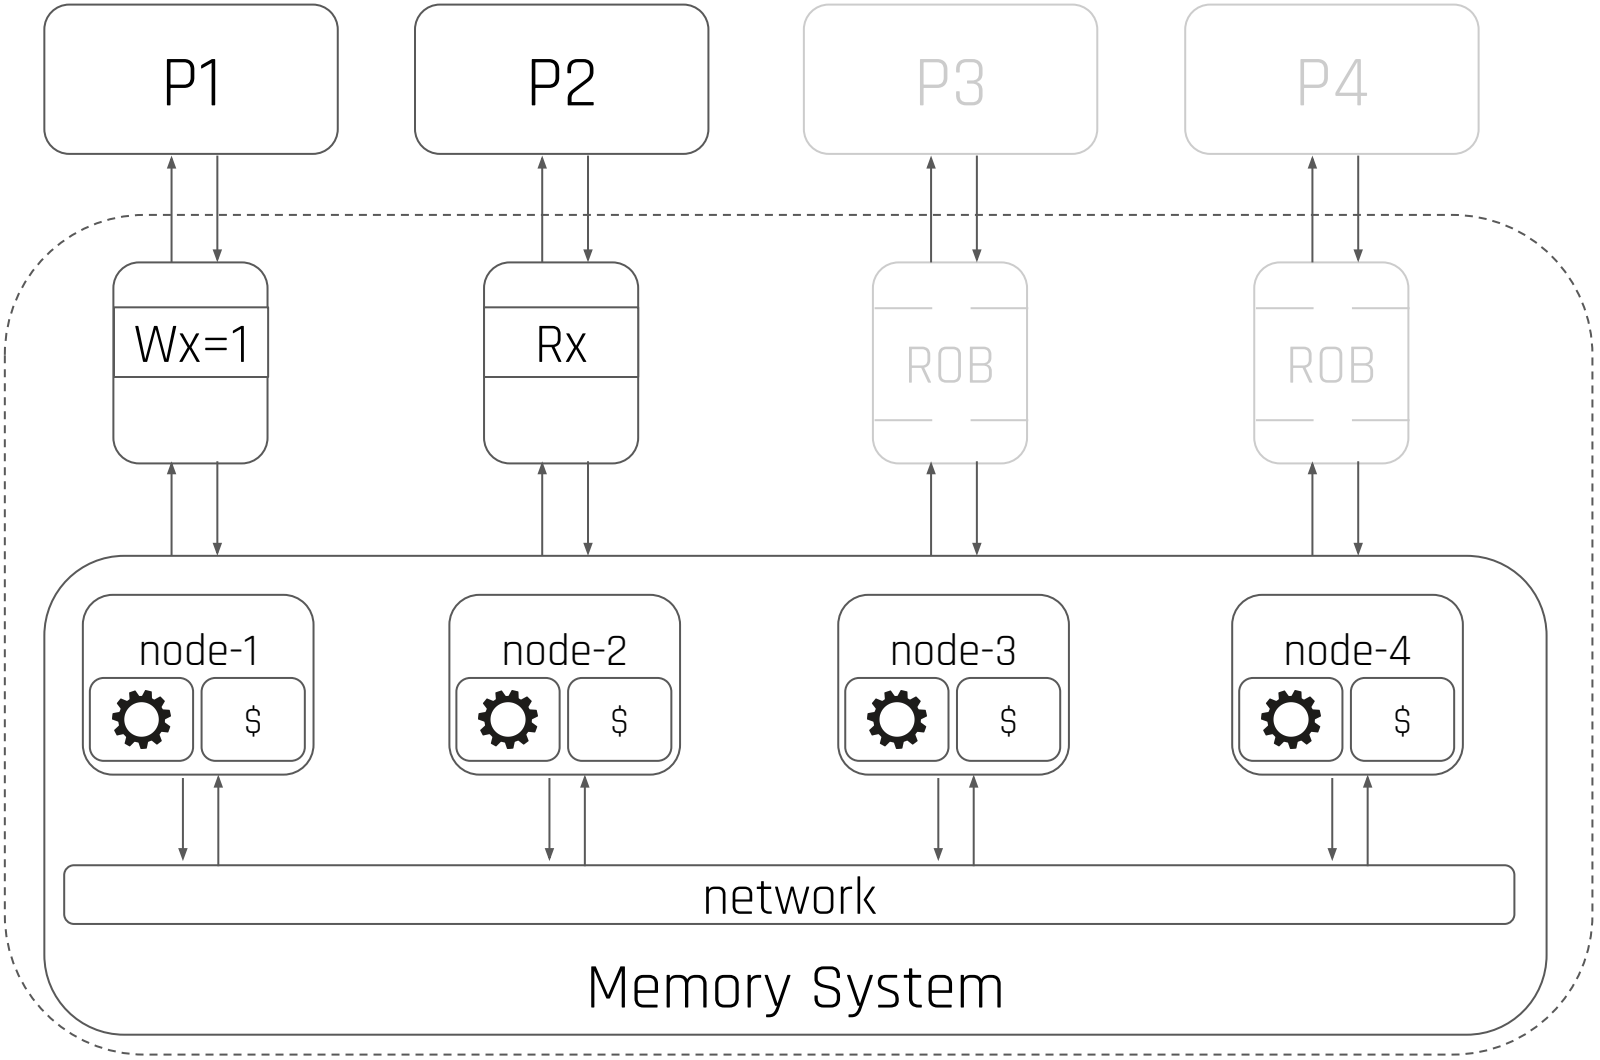
P1
P2
P3
P4
Wx=1
Rx
ROB
ROB
node-1
$
node-2
$
node-3
$
node-4
$
network
Memory System

1. srt_{WR} (W⟶ R)
2. srt_{WW} (W⟶ W)
3. srt_{RR} (R⟶ R)
4. srt_{RW} (R⟶ W)

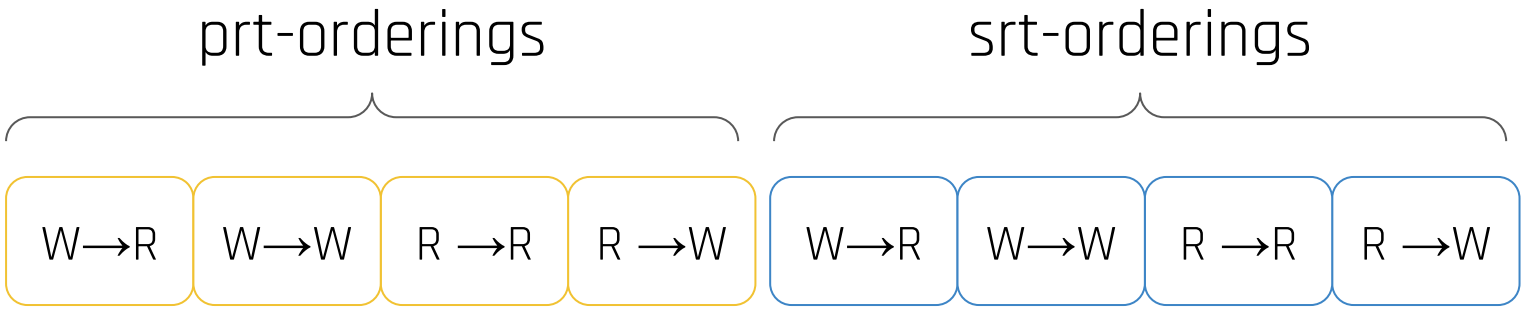
prt-orderings
srt-orderings
W⟶R
W⟶W
R ⟶R
R ⟶W
W⟶R
W⟶W
R ⟶R
R ⟶W

Real Time Orderings

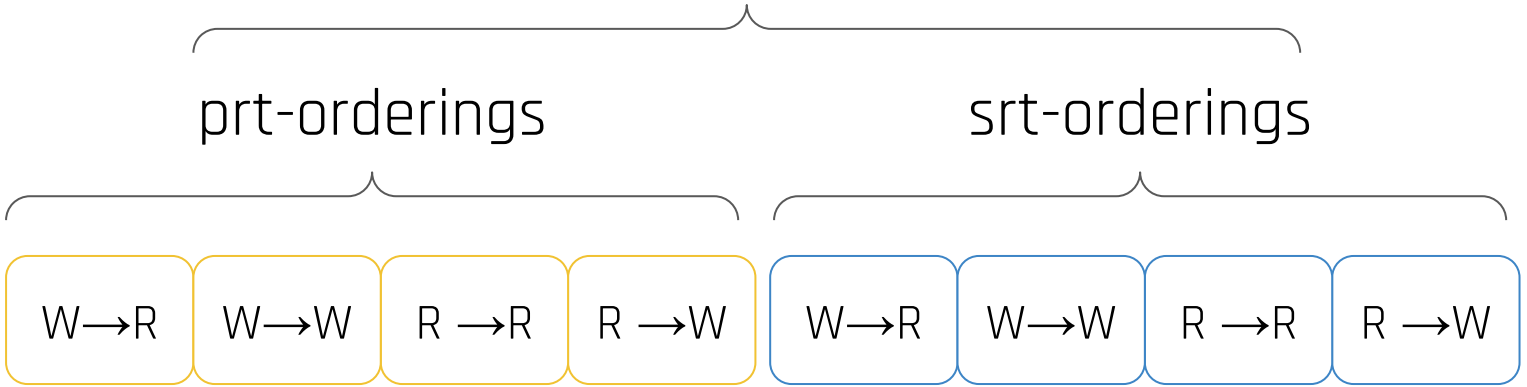

Real Time Orderings

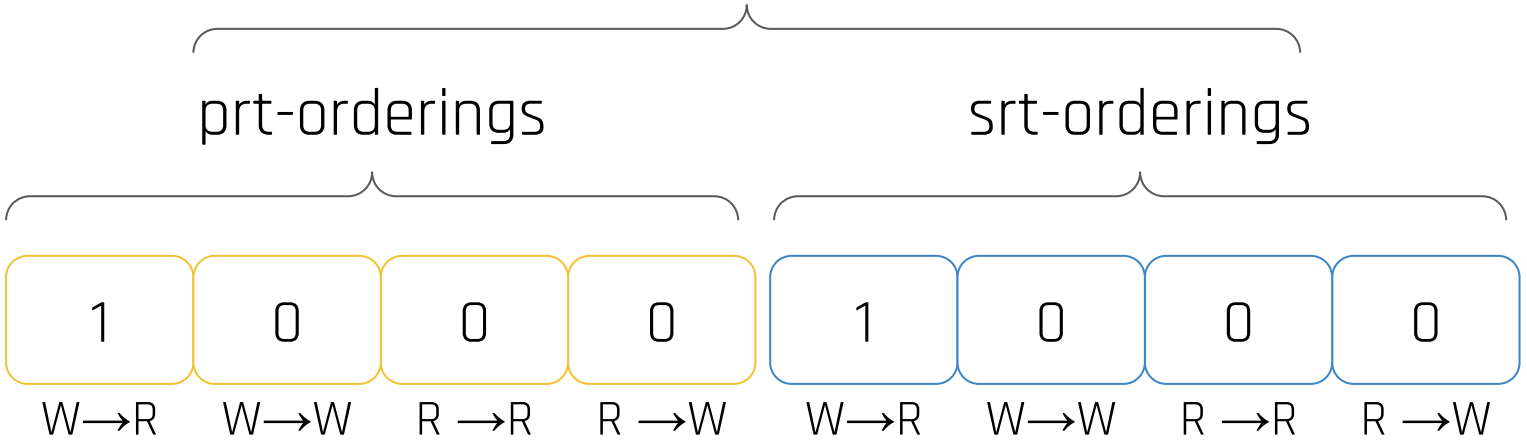

Consistency Model

Synchronization Patterns

Mapping

Protocol

Real time orderings

Initially X, Y = 0	
P1	P2
Write (x = 1)	
Read (y = 0)	
	Write (y = 1)
	Read (x)

Initially X, Y = 0

P1	P2
Write (x = 1)	
Read (y = 0)	
	Write (y = 1)
	Read (x)

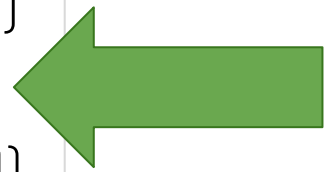

Initially X, Y = 0

P1	P2
Write (x = 1)	
Read (y = 0)	
	Write (y = 1)
	Read (x)

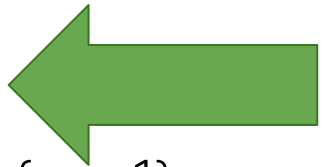

Initially X, Y = 0

P1	P2
Write (x = 1)	
↓	
Read (y = 0)	
→	Write (y = 1)
	↓
	Read (x)

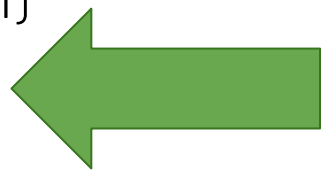

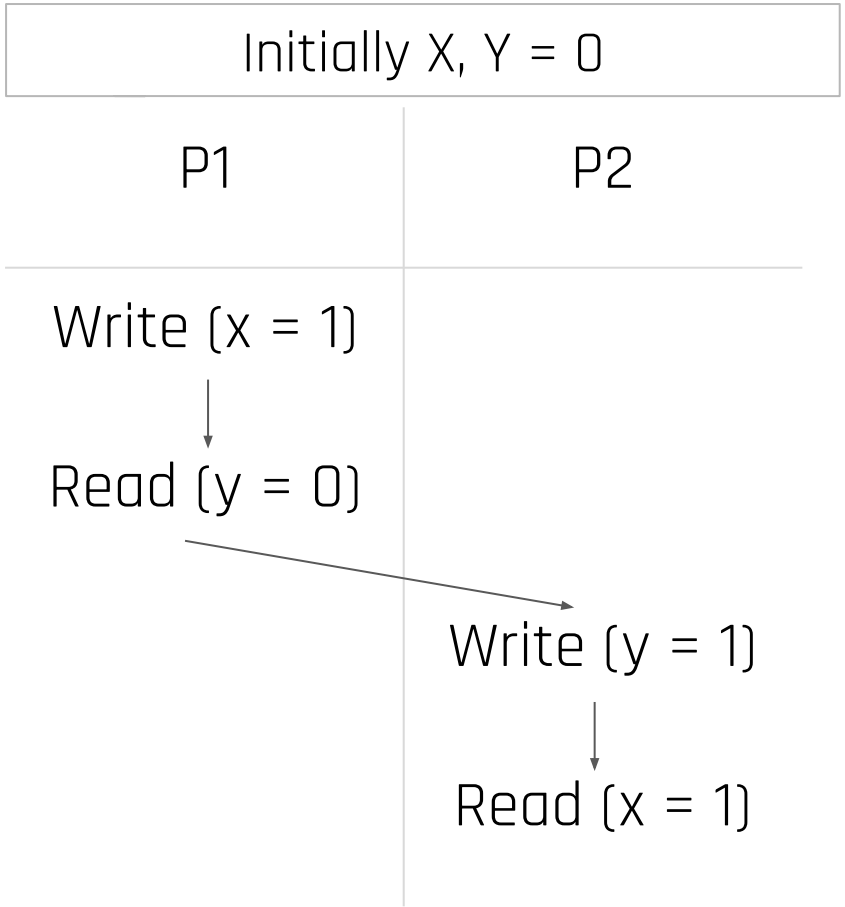
Initially X, Y = 0
P1
P2
Write (x = 1)
Read (y = 0)
Write (y = 1)
Read (x = 1)

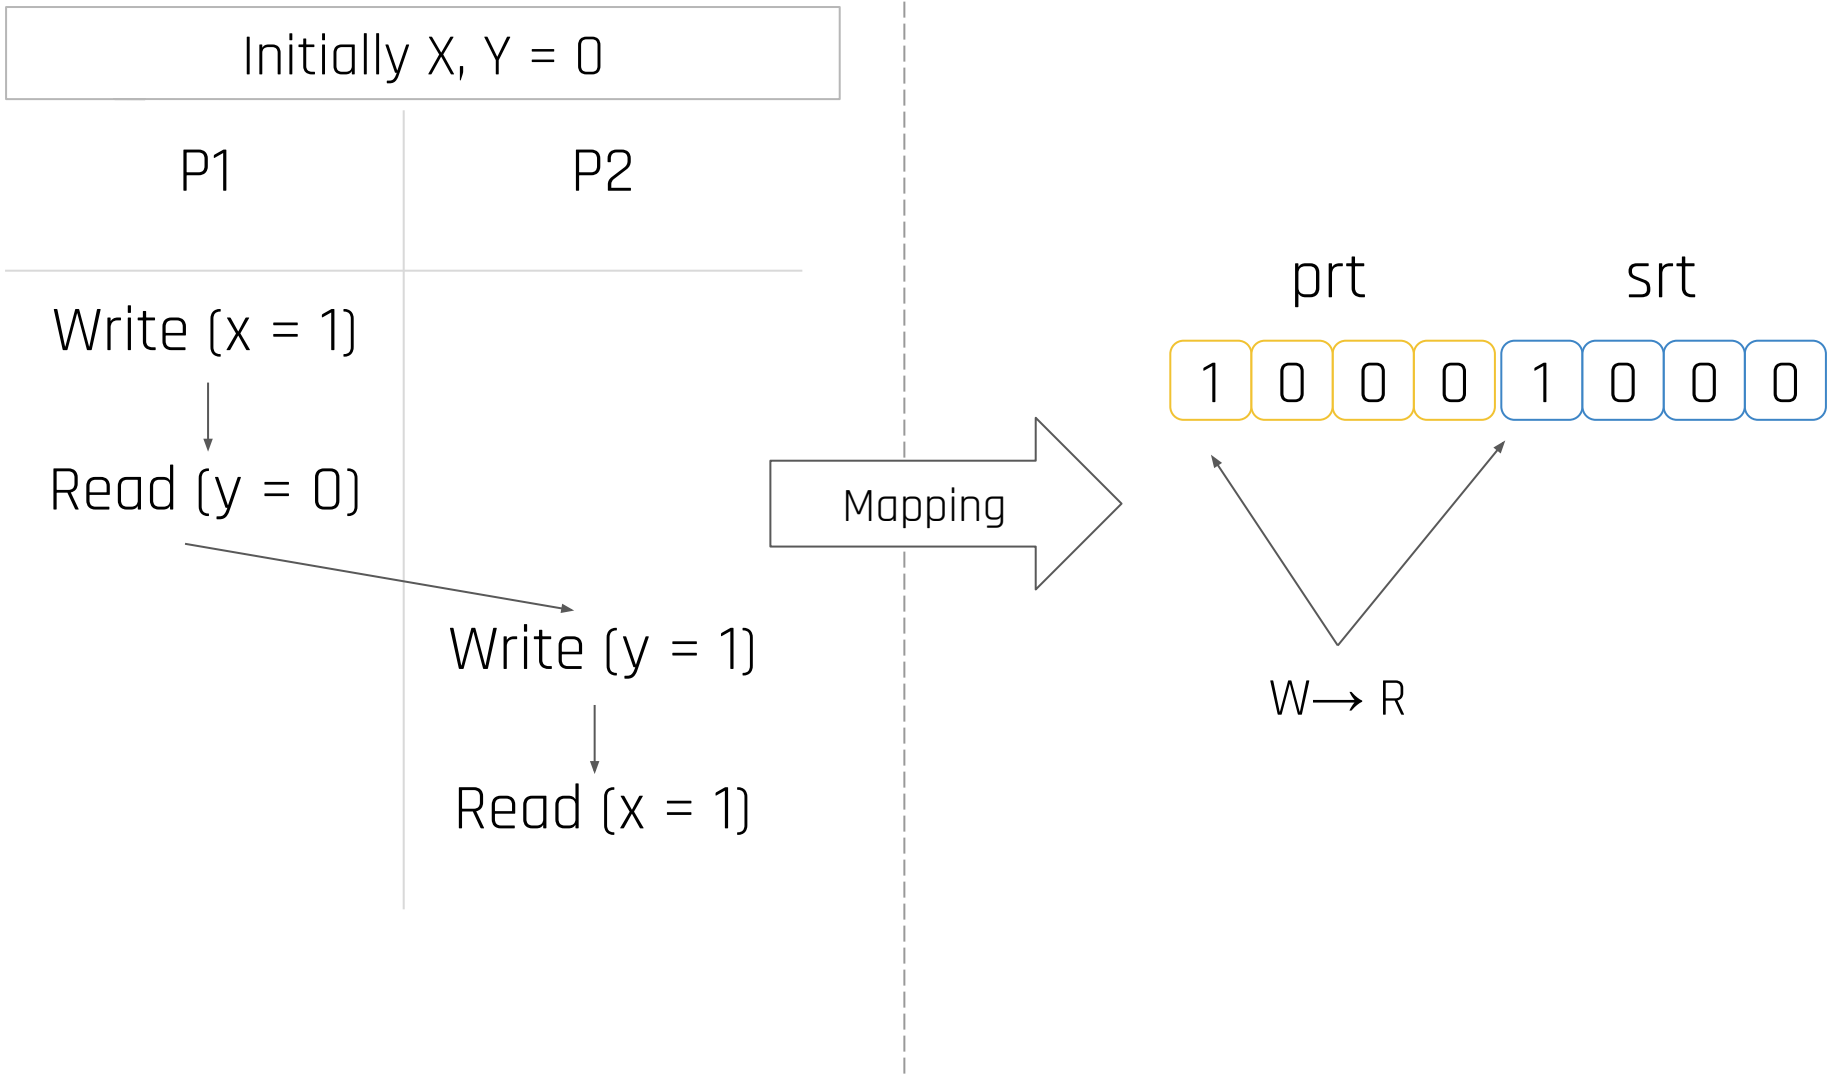
Initially X, Y = 0
P1
P2
Write (x = 1)
Read (y = 0)
Write (y = 1)
Read (x = 1)
Mapping
prt
srt
1 0 0 0 1 0 0 0
W⟶ R

Consistency Model

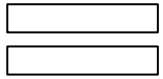

Synchronization Patterns

Mapping

Protocol

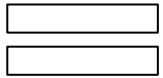

Real time orderings

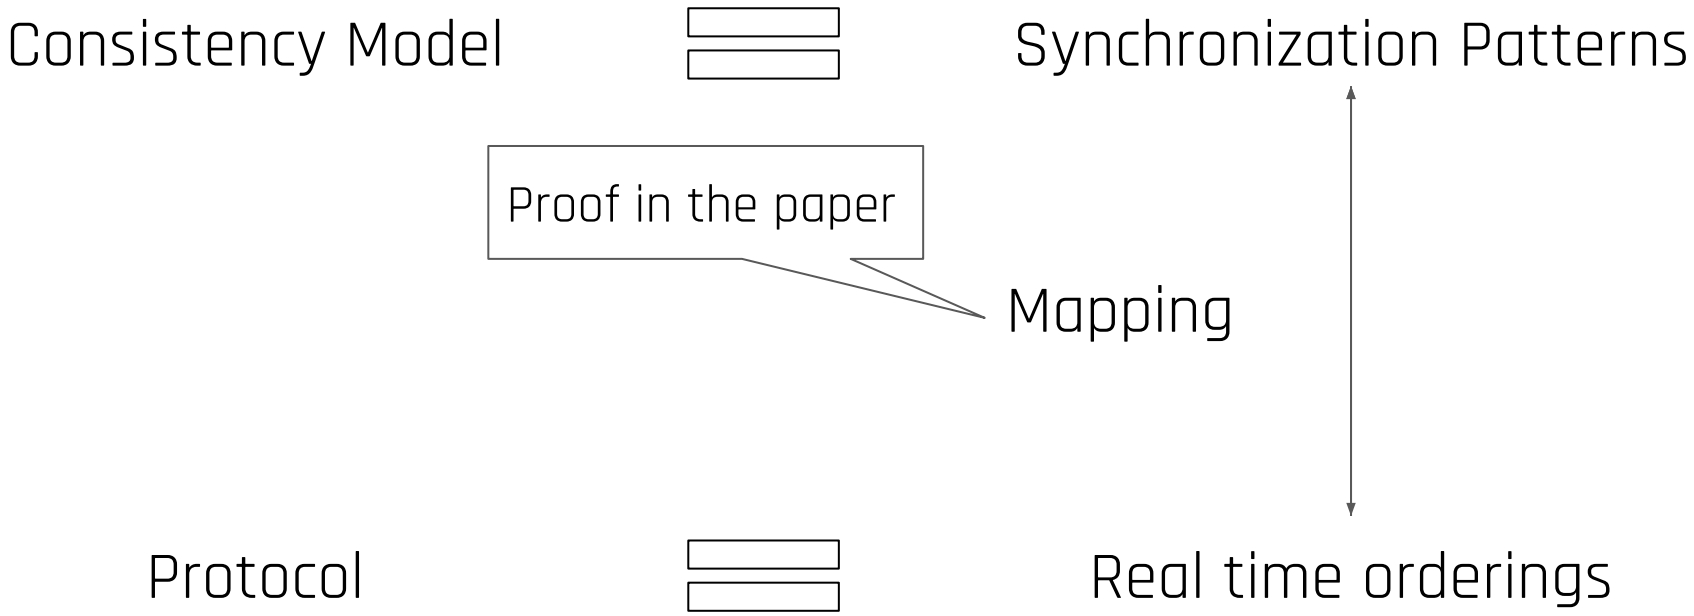
Consistency Model
Synchronization Patterns
Proof in the paper
Mapping
Protocol
Real time orderings

Abstraction

Consistency Model

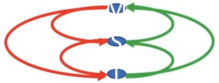

Protocol High-level

Protocol Low-level

Verification

Simulation

Implementation

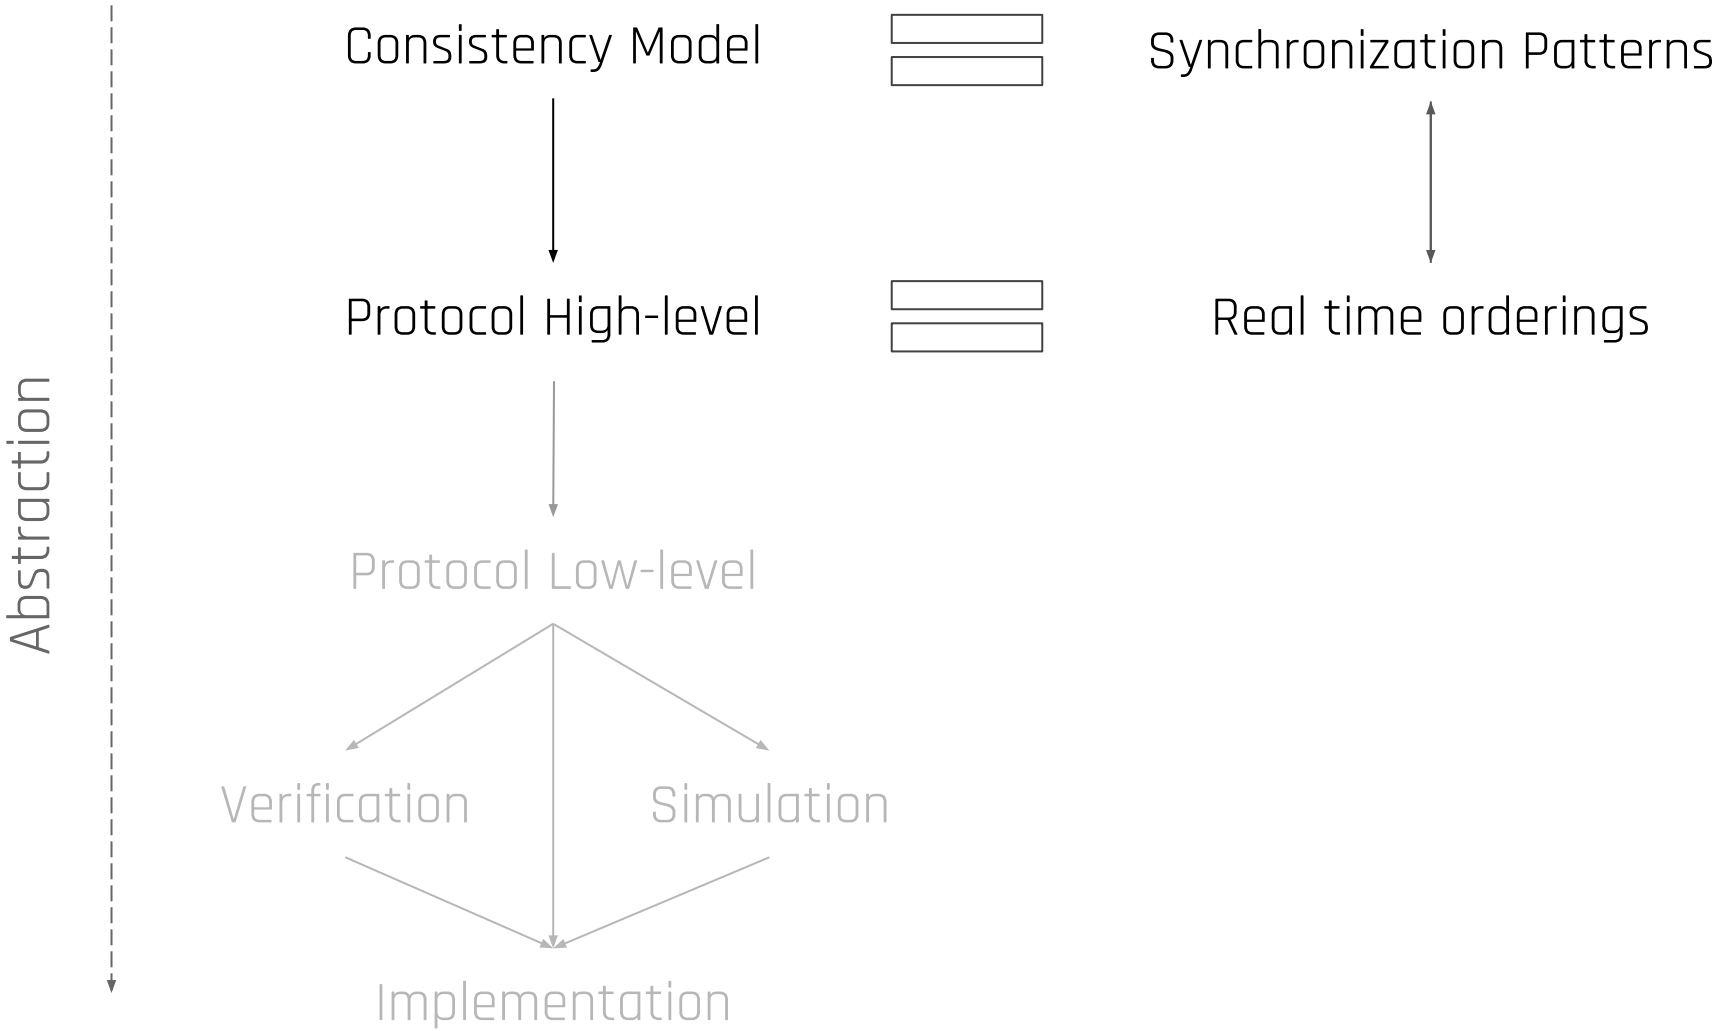
Abstraction
Consistency Model
Synchronization Patterns
Protocol High-level
Real time orderings
Protocol Low-level
Verification
Simulation
Implementation

Backup Slides