Efficient Counting with Optimal Resilience

Christoph Lenzen
MPI for Informatics

Joel Rybicki
MPI for Informatics
& Aalto University / HIIT

DISC 2015
October 7, Tokyo

Fault-tolerant counters

Deterministic ***round counters*** tolerating:

- **permanent** failures (*Byzantine faults*)
- **transient** failures (*self-stabilisation*)

that are

- **fast** to stabilise
- **communication-efficient**

Model of computing

n state machines

Synchronous rounds:

1. broadcast
2. receive
3. update state

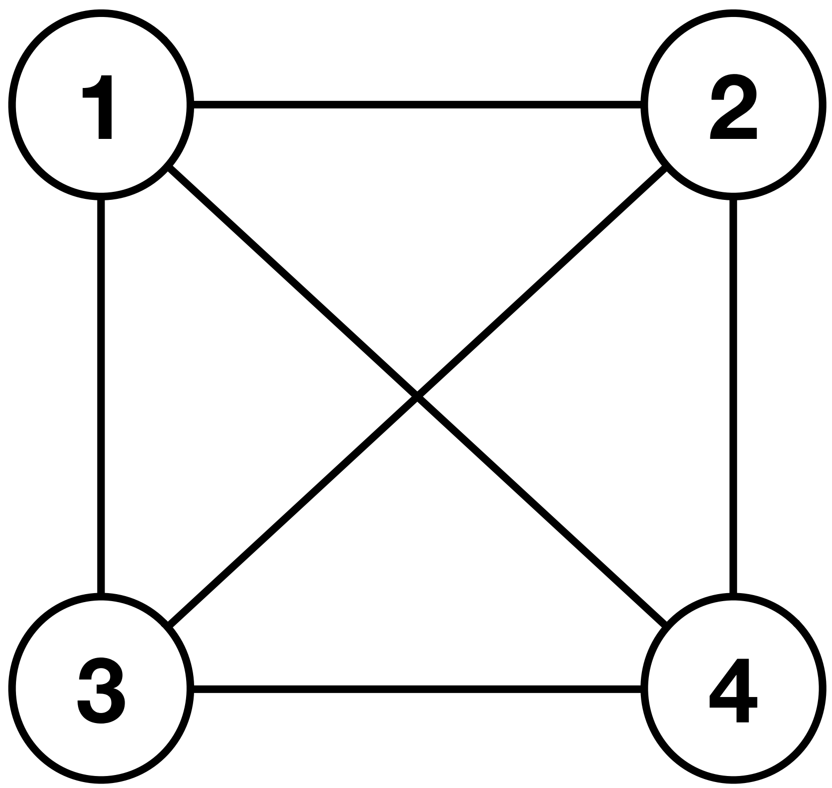

Model of computing

n state machines

Synchronous rounds:

1. broadcast

2. receive

3. update state

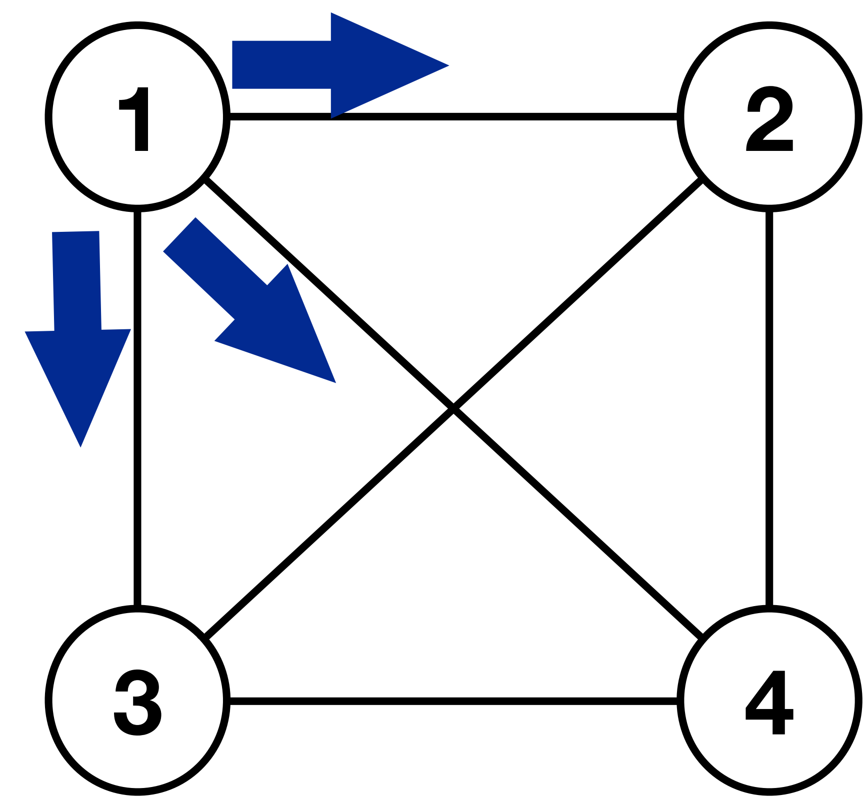

Model of computing

n state machines

Synchronous rounds:

1. broadcast
2. **receive**
3. update state

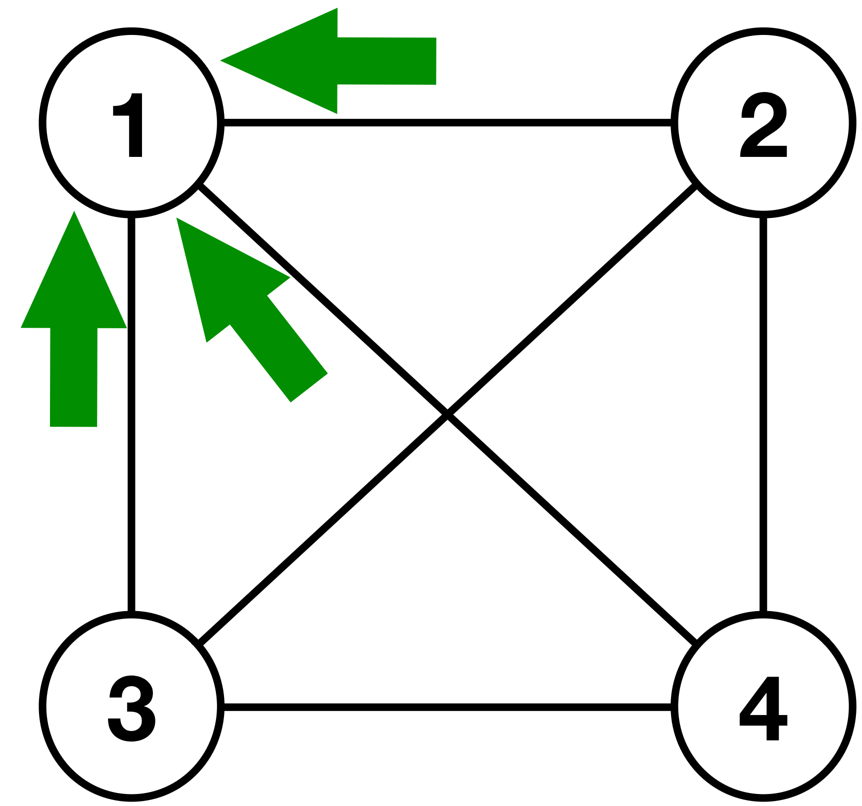

Model of computing

n state machines

Synchronous rounds:

1. broadcast
2. receive
3. **update state**

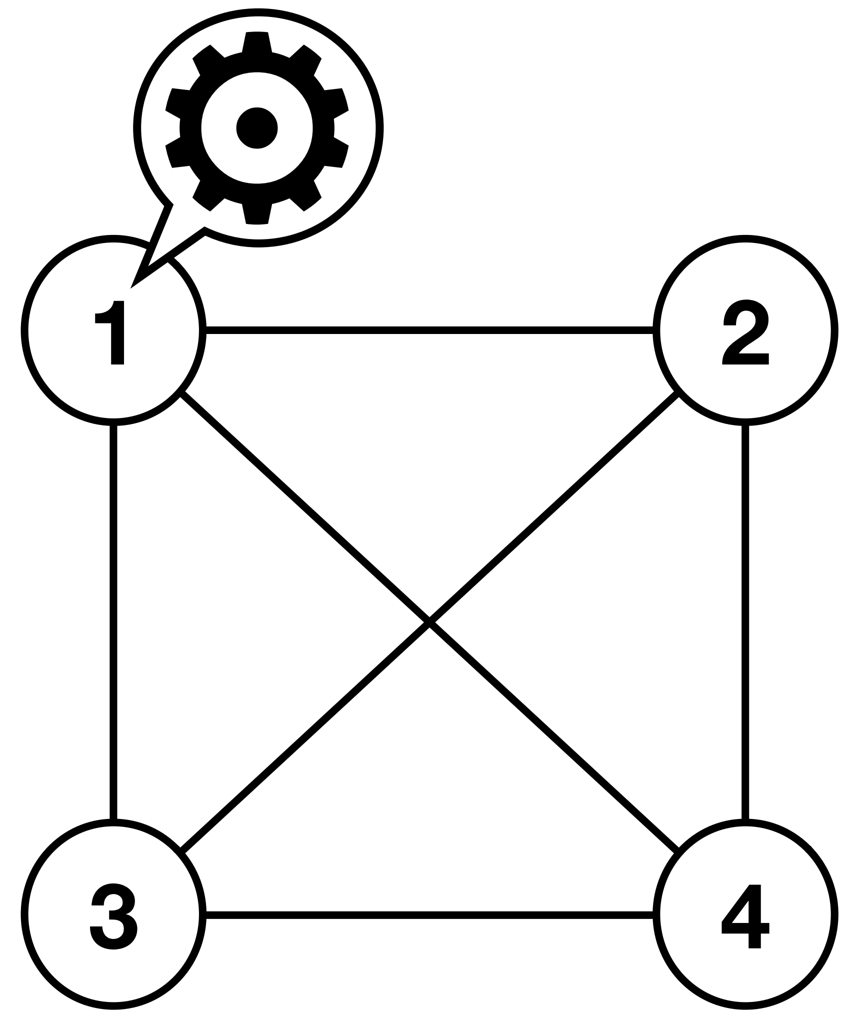

our adversary

Transient failures

arbitrary initial states

chosen by adversary!

= self-stabilisation

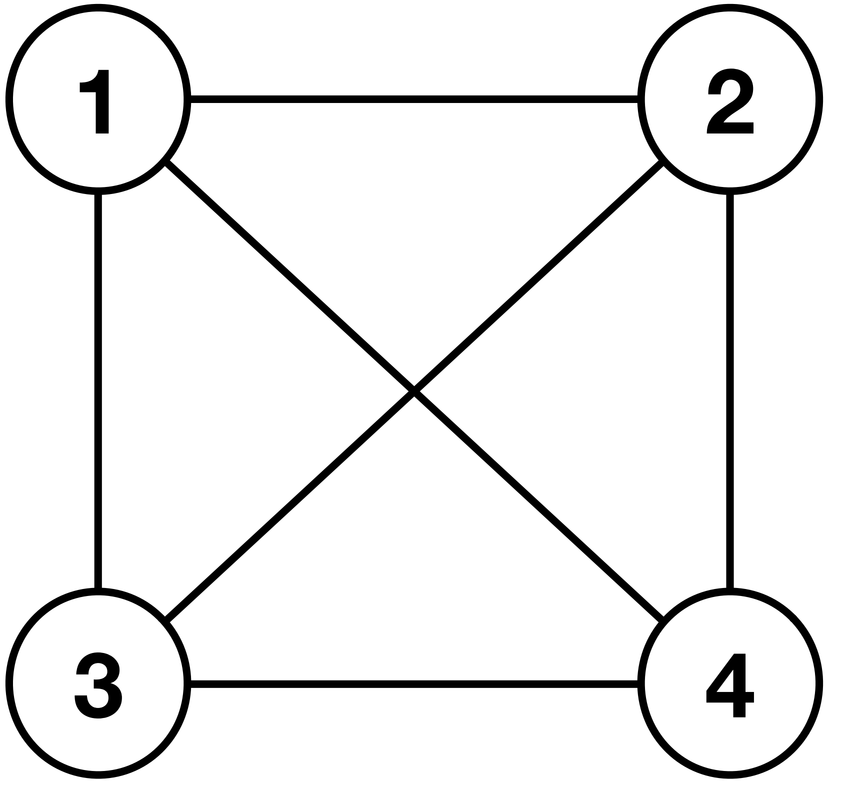

Byzantine failures

n state machines

f Byzantine failures

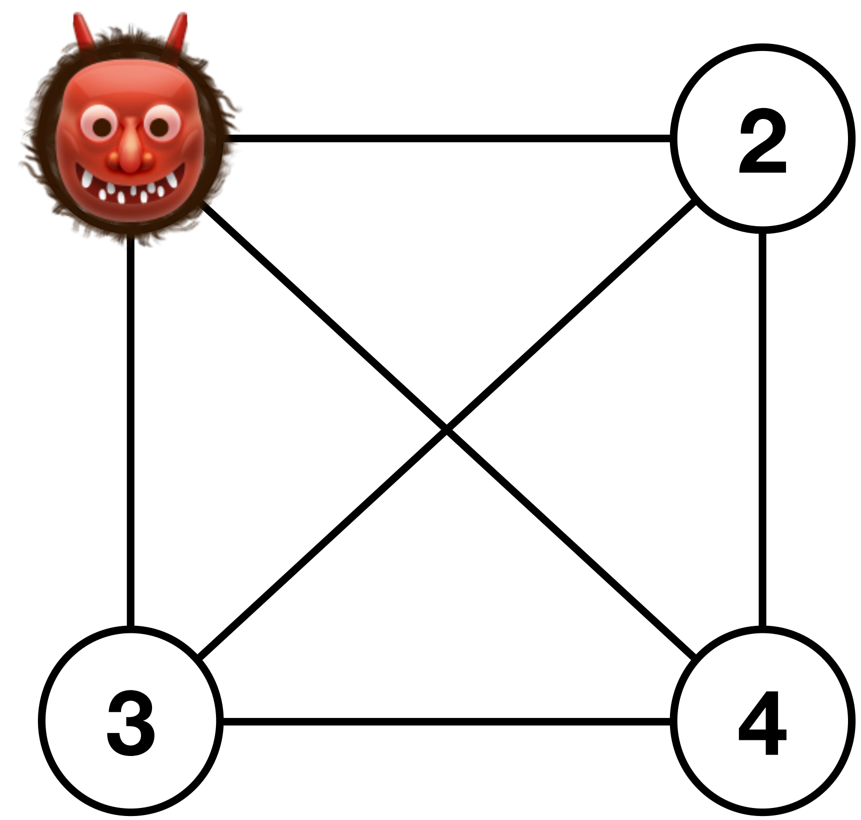

Byzantine failures

n state machines

f Byzantine failures

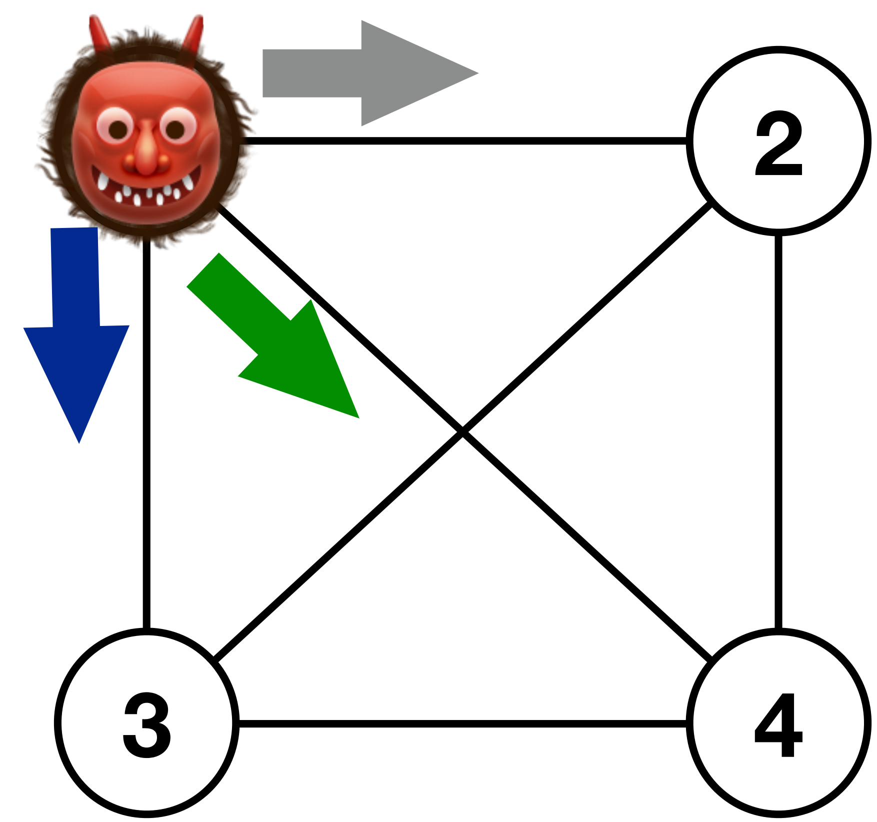

Byzantine failures

n state machines

f Byzantine failures

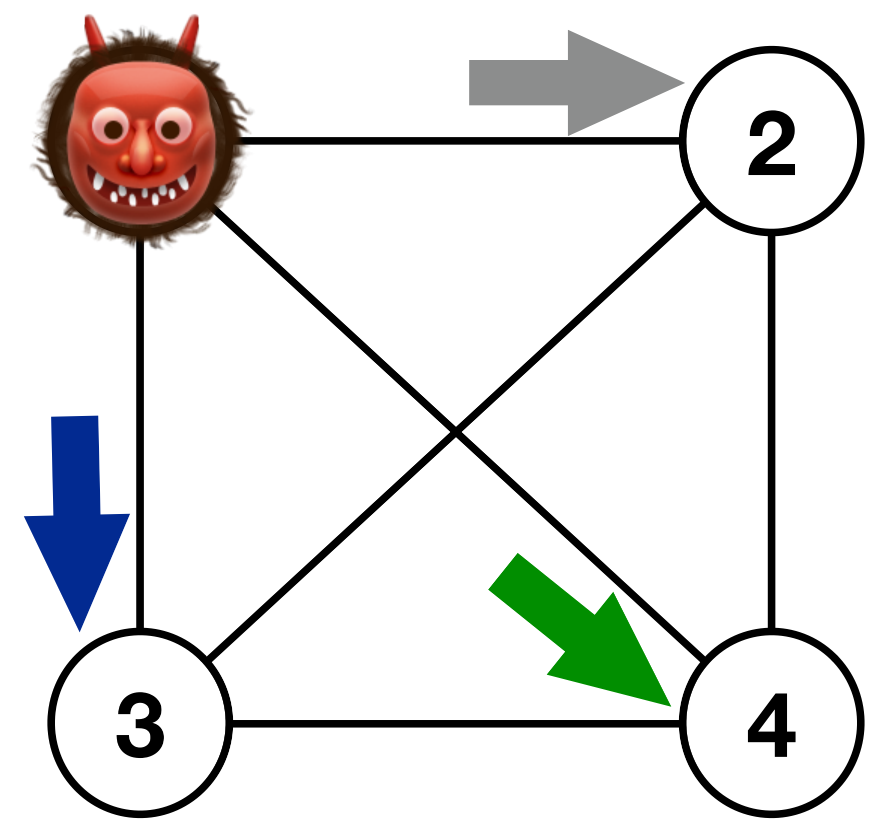

Correct nodes can observe *different* states for the system!

Counting mod c

3-counting

0	**1**	**2**	**0**	**1**	**2**

increment counter +1 mod *c* each round

Synchronous counting

	Counting					
1	0	1	2	0	1	2
👹	?	?	?	?	?	?
3	0	1	2	0	1	2
4	0	1	2	0	1	2

Synchronous counting

	Stabilisation			Counting					
1	1	0	2	0	1	2	0	1	2
👹	?	?	?	?	?	?	?	?	?
3	2	2	2	0	1	2	0	1	2
4	1	2	1	0	1	2	0	1	2

Complexity measures

Time complexity: #rounds

Message size:
maximum number of bits broadcast
(per node, each round)

A related problem: Consensus

Input: private value **Output:** agreement

A related problem: Consensus

Input: private value **Output:** agreement

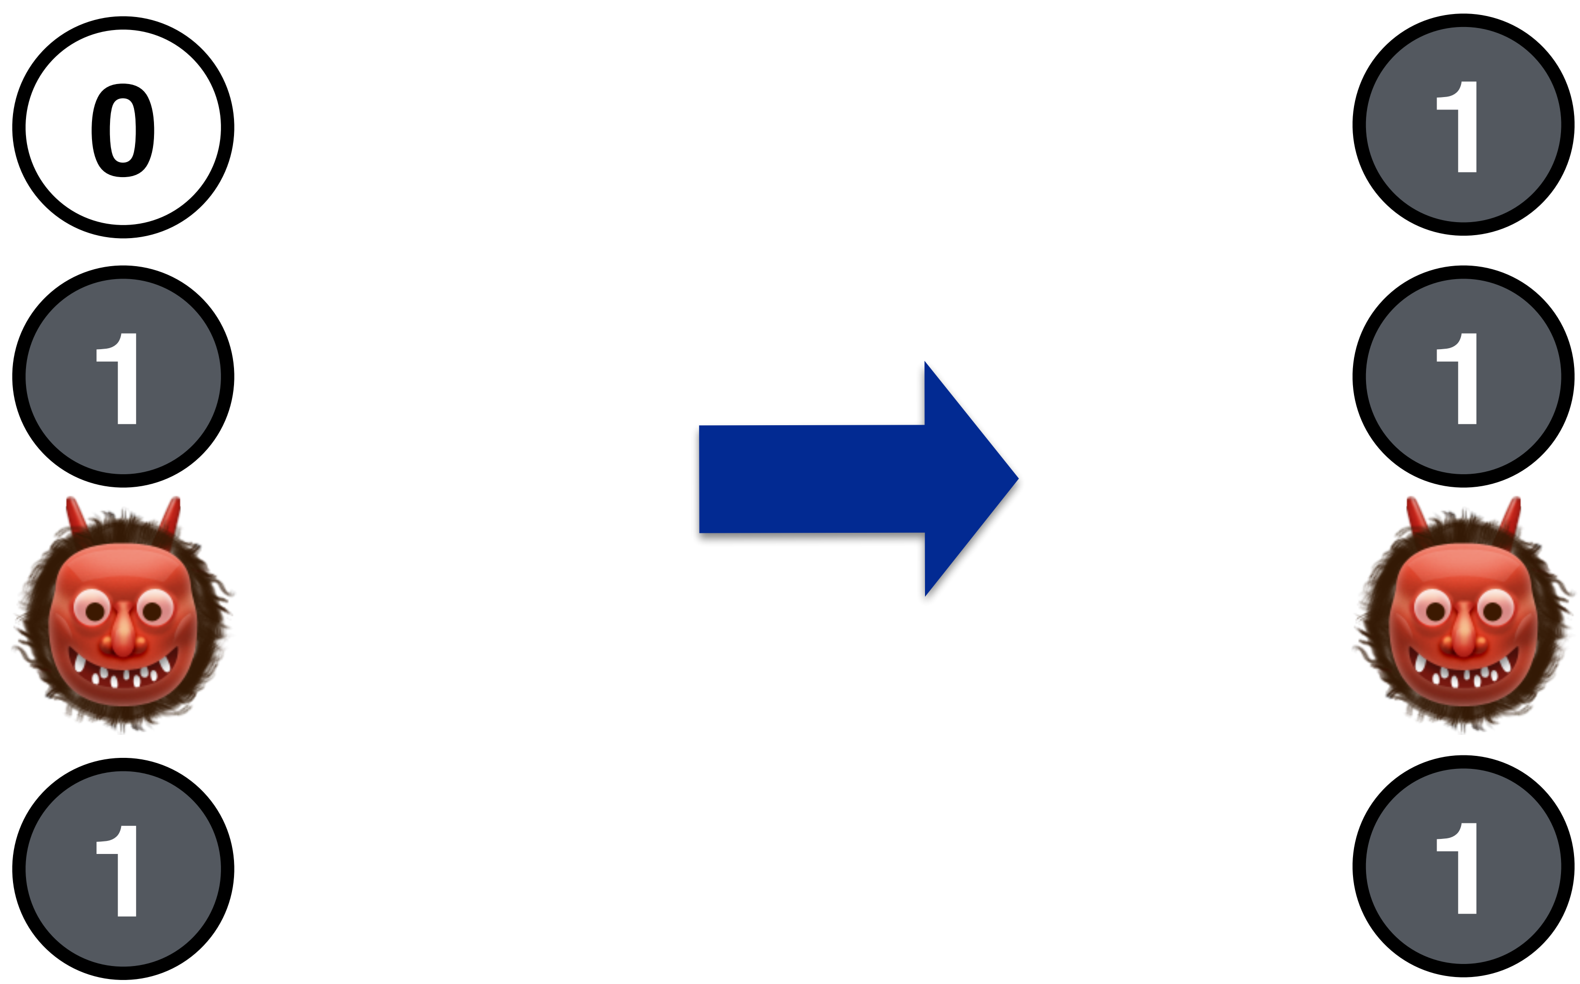

A related problem: Consensus

Input: private value **Output:** agreement

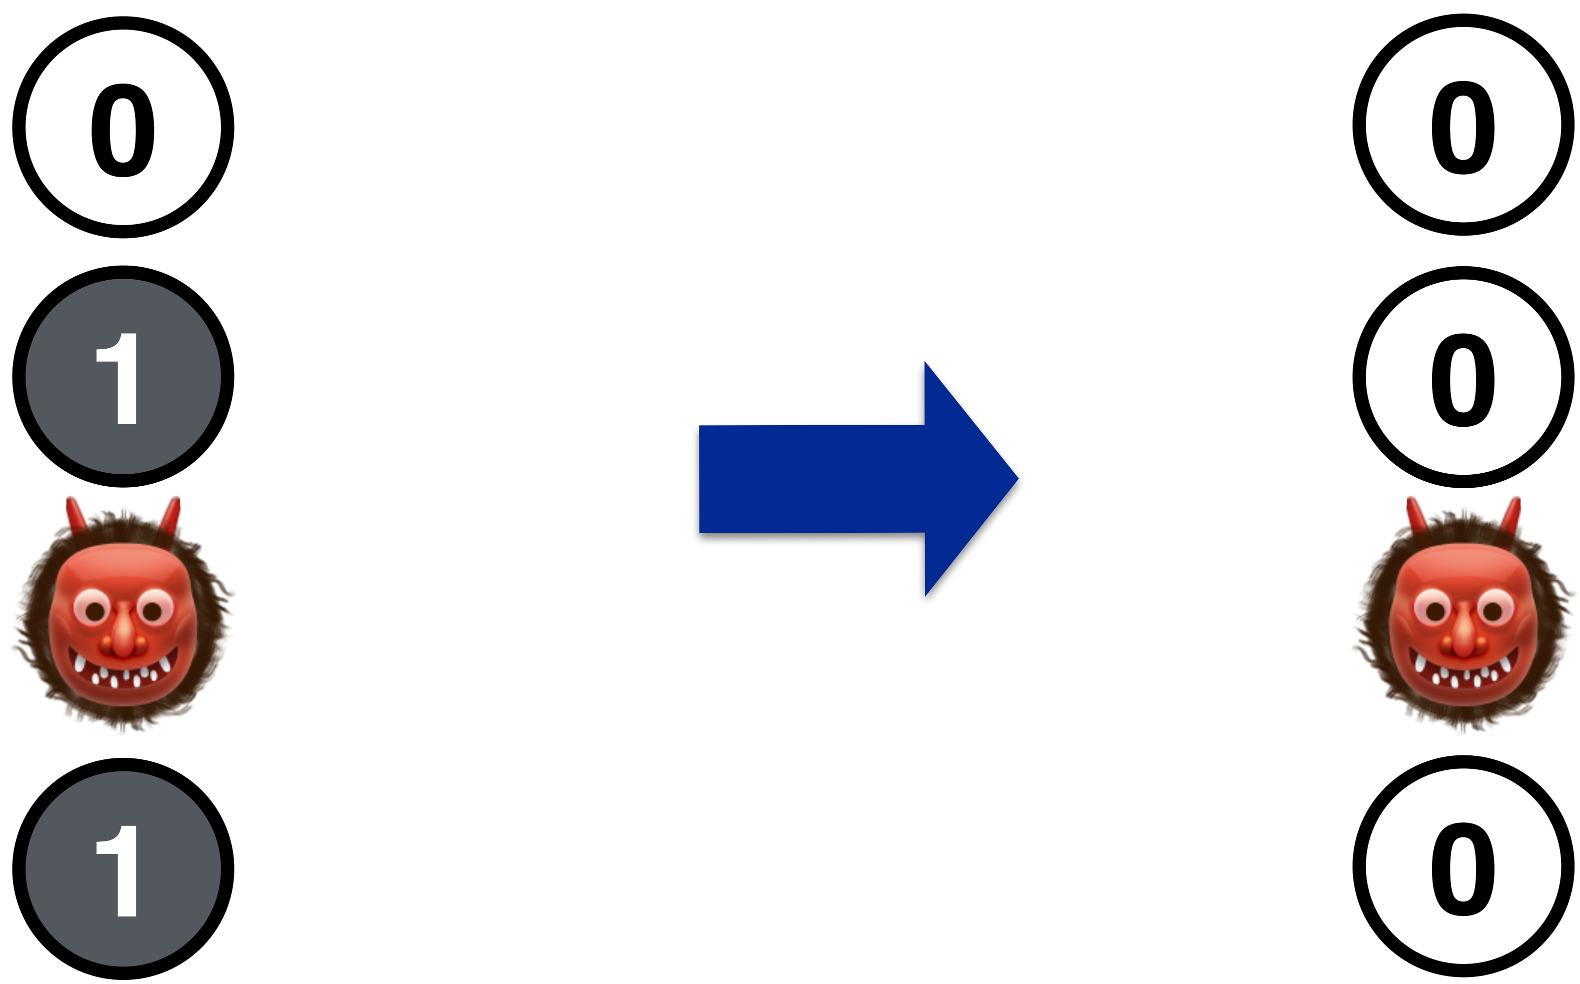

A related problem: Consensus

Input: private value **Output:** agreement

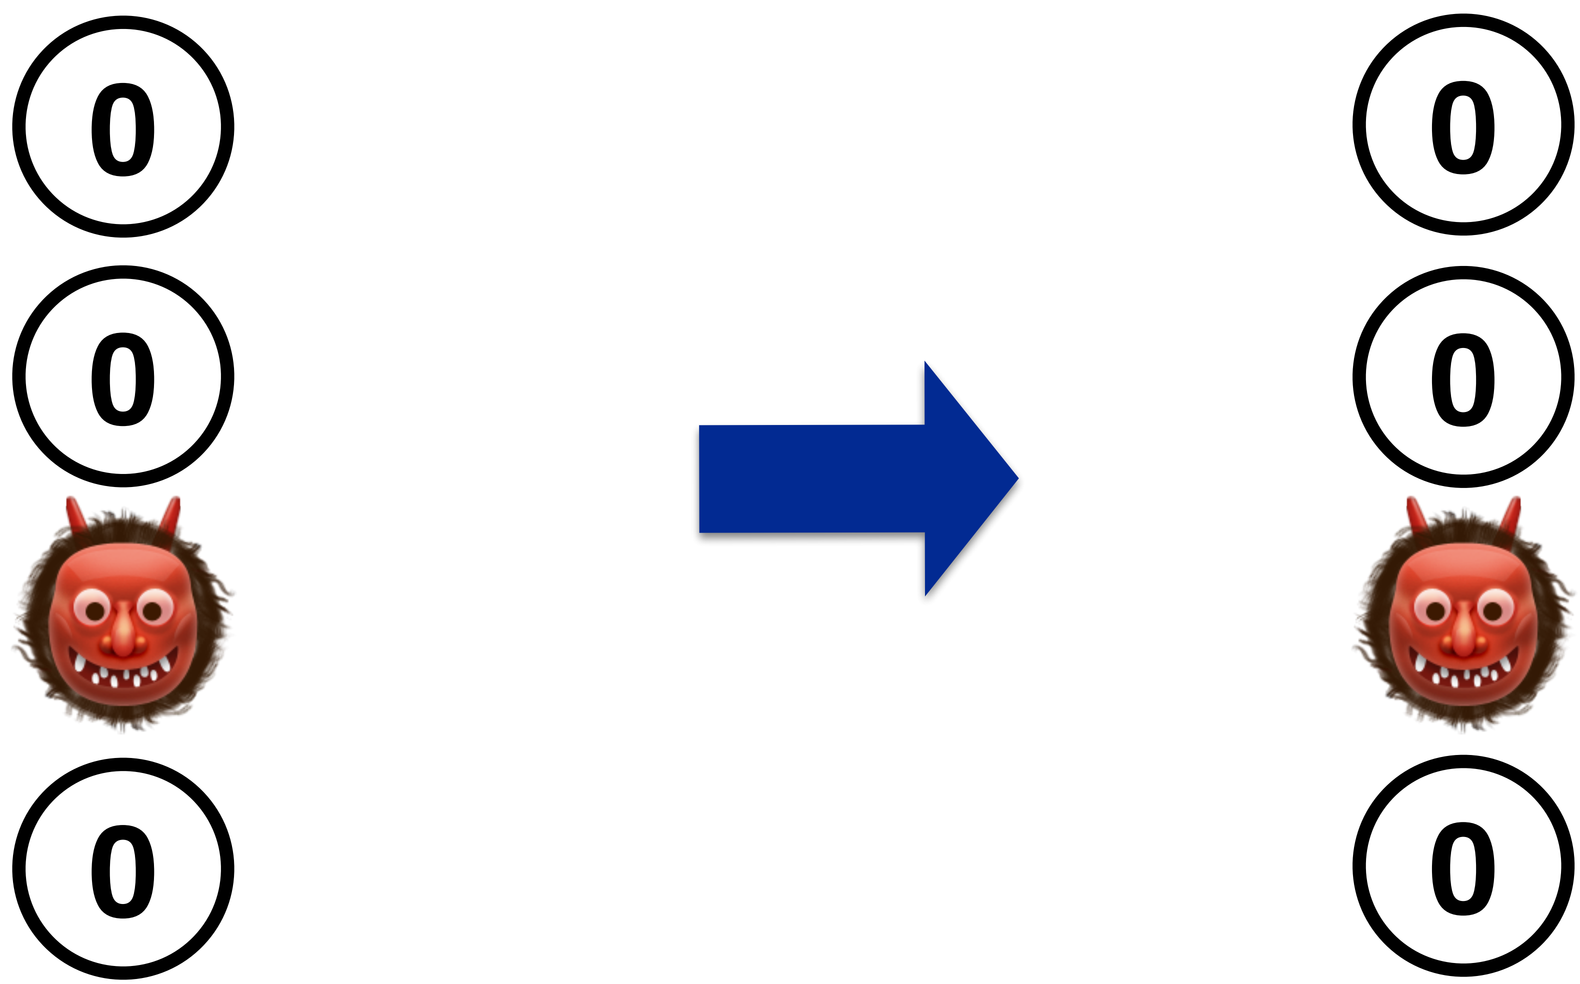

A related problem: Consensus

Input: private value **Output:** agreement

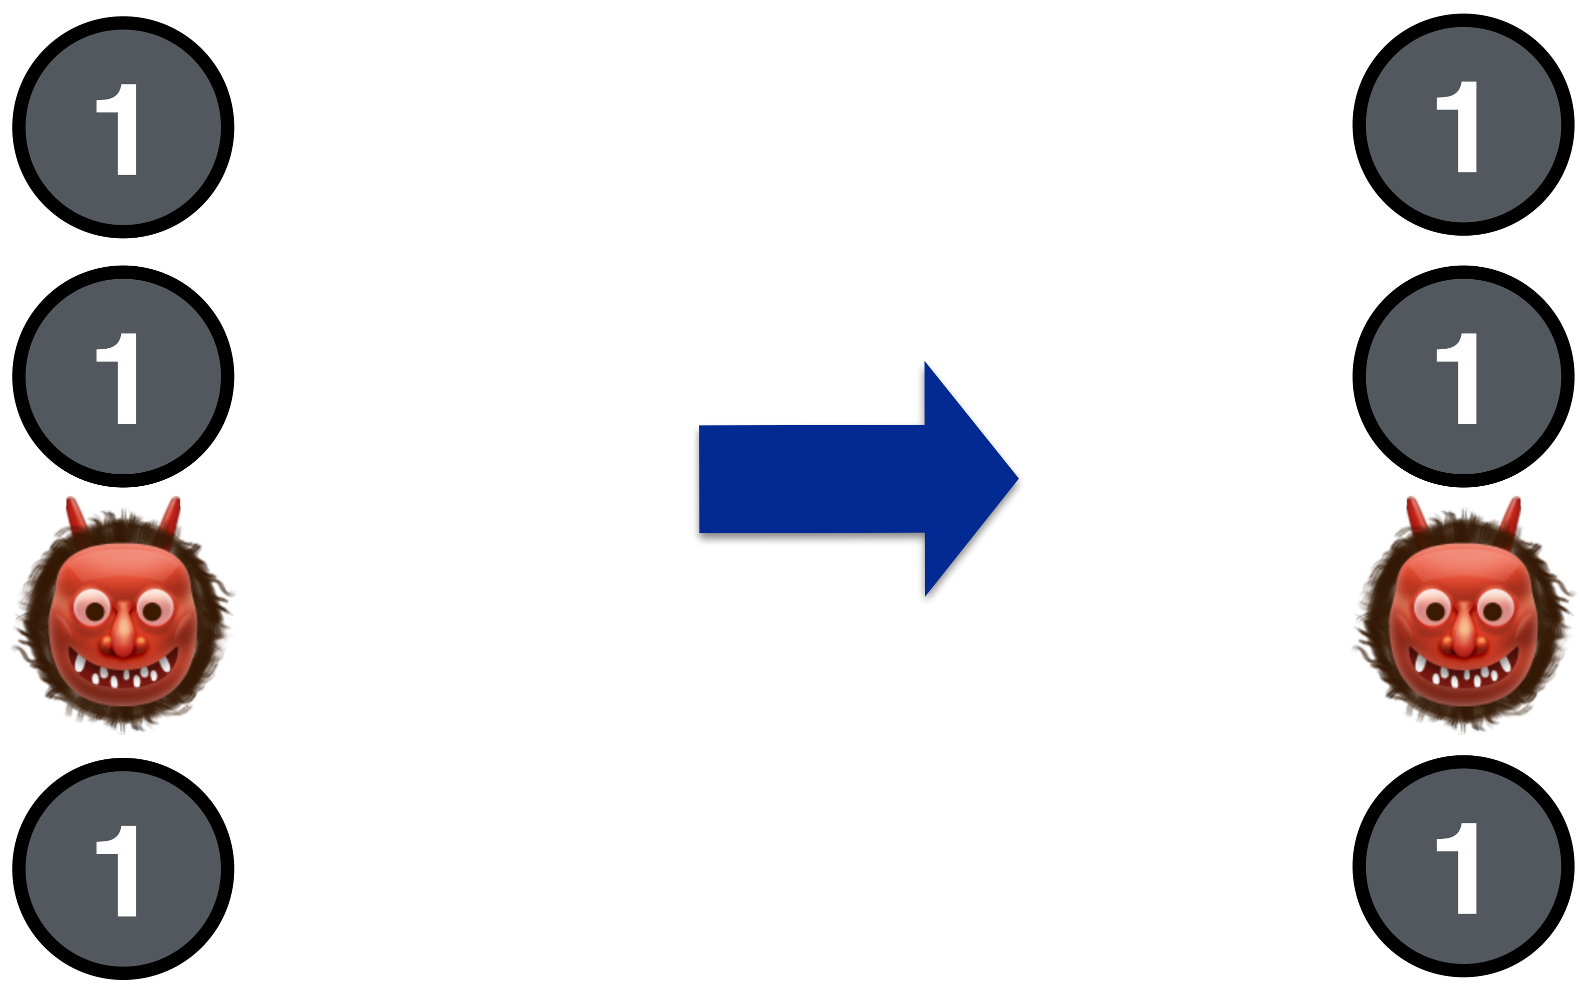

Consensus bounds*

Resilience

$$f < n/3$$

Pease et al. (1980)

Time

More than f rounds to reach agreement

Fischer & Lynch (1982)

***deterministic**

Counting bounds*

Resilience

$$f < n/3$$

Pease et al. (1980)

Time

More than f rounds to stabilise

Fischer & Lynch (1982)

***deterministic**

Upper bounds

 = randomised

	Time	Resilience	Bits
 S. Dolev & Welch 2004	$2^{2(n-f)}$	$< n/3$	$O(1)$
 Ben-Or *et al.* 2008	$O(1)$	$< n/3$	$n^{O(1)}$
D. Dolev & Hoch 2007	$\Theta(f)$	$< n/3$	$\Omega(f)$
PODC 2015	$\Theta(f)$	$n^{1-o(1)}$	$O(\log^2 f)$
This result	$\Theta(f)$	$< n/3$	$O(\log^2 f)$

	Time	Resilience	Bits
D. Dolev & Hoch 2007	$\Theta(f)$	$< n/3$	$\Omega(f)$
PODC 2015	$\Theta(f)$	$n^{1-o(1)}$	$O(\log^2 f)$

High-level idea

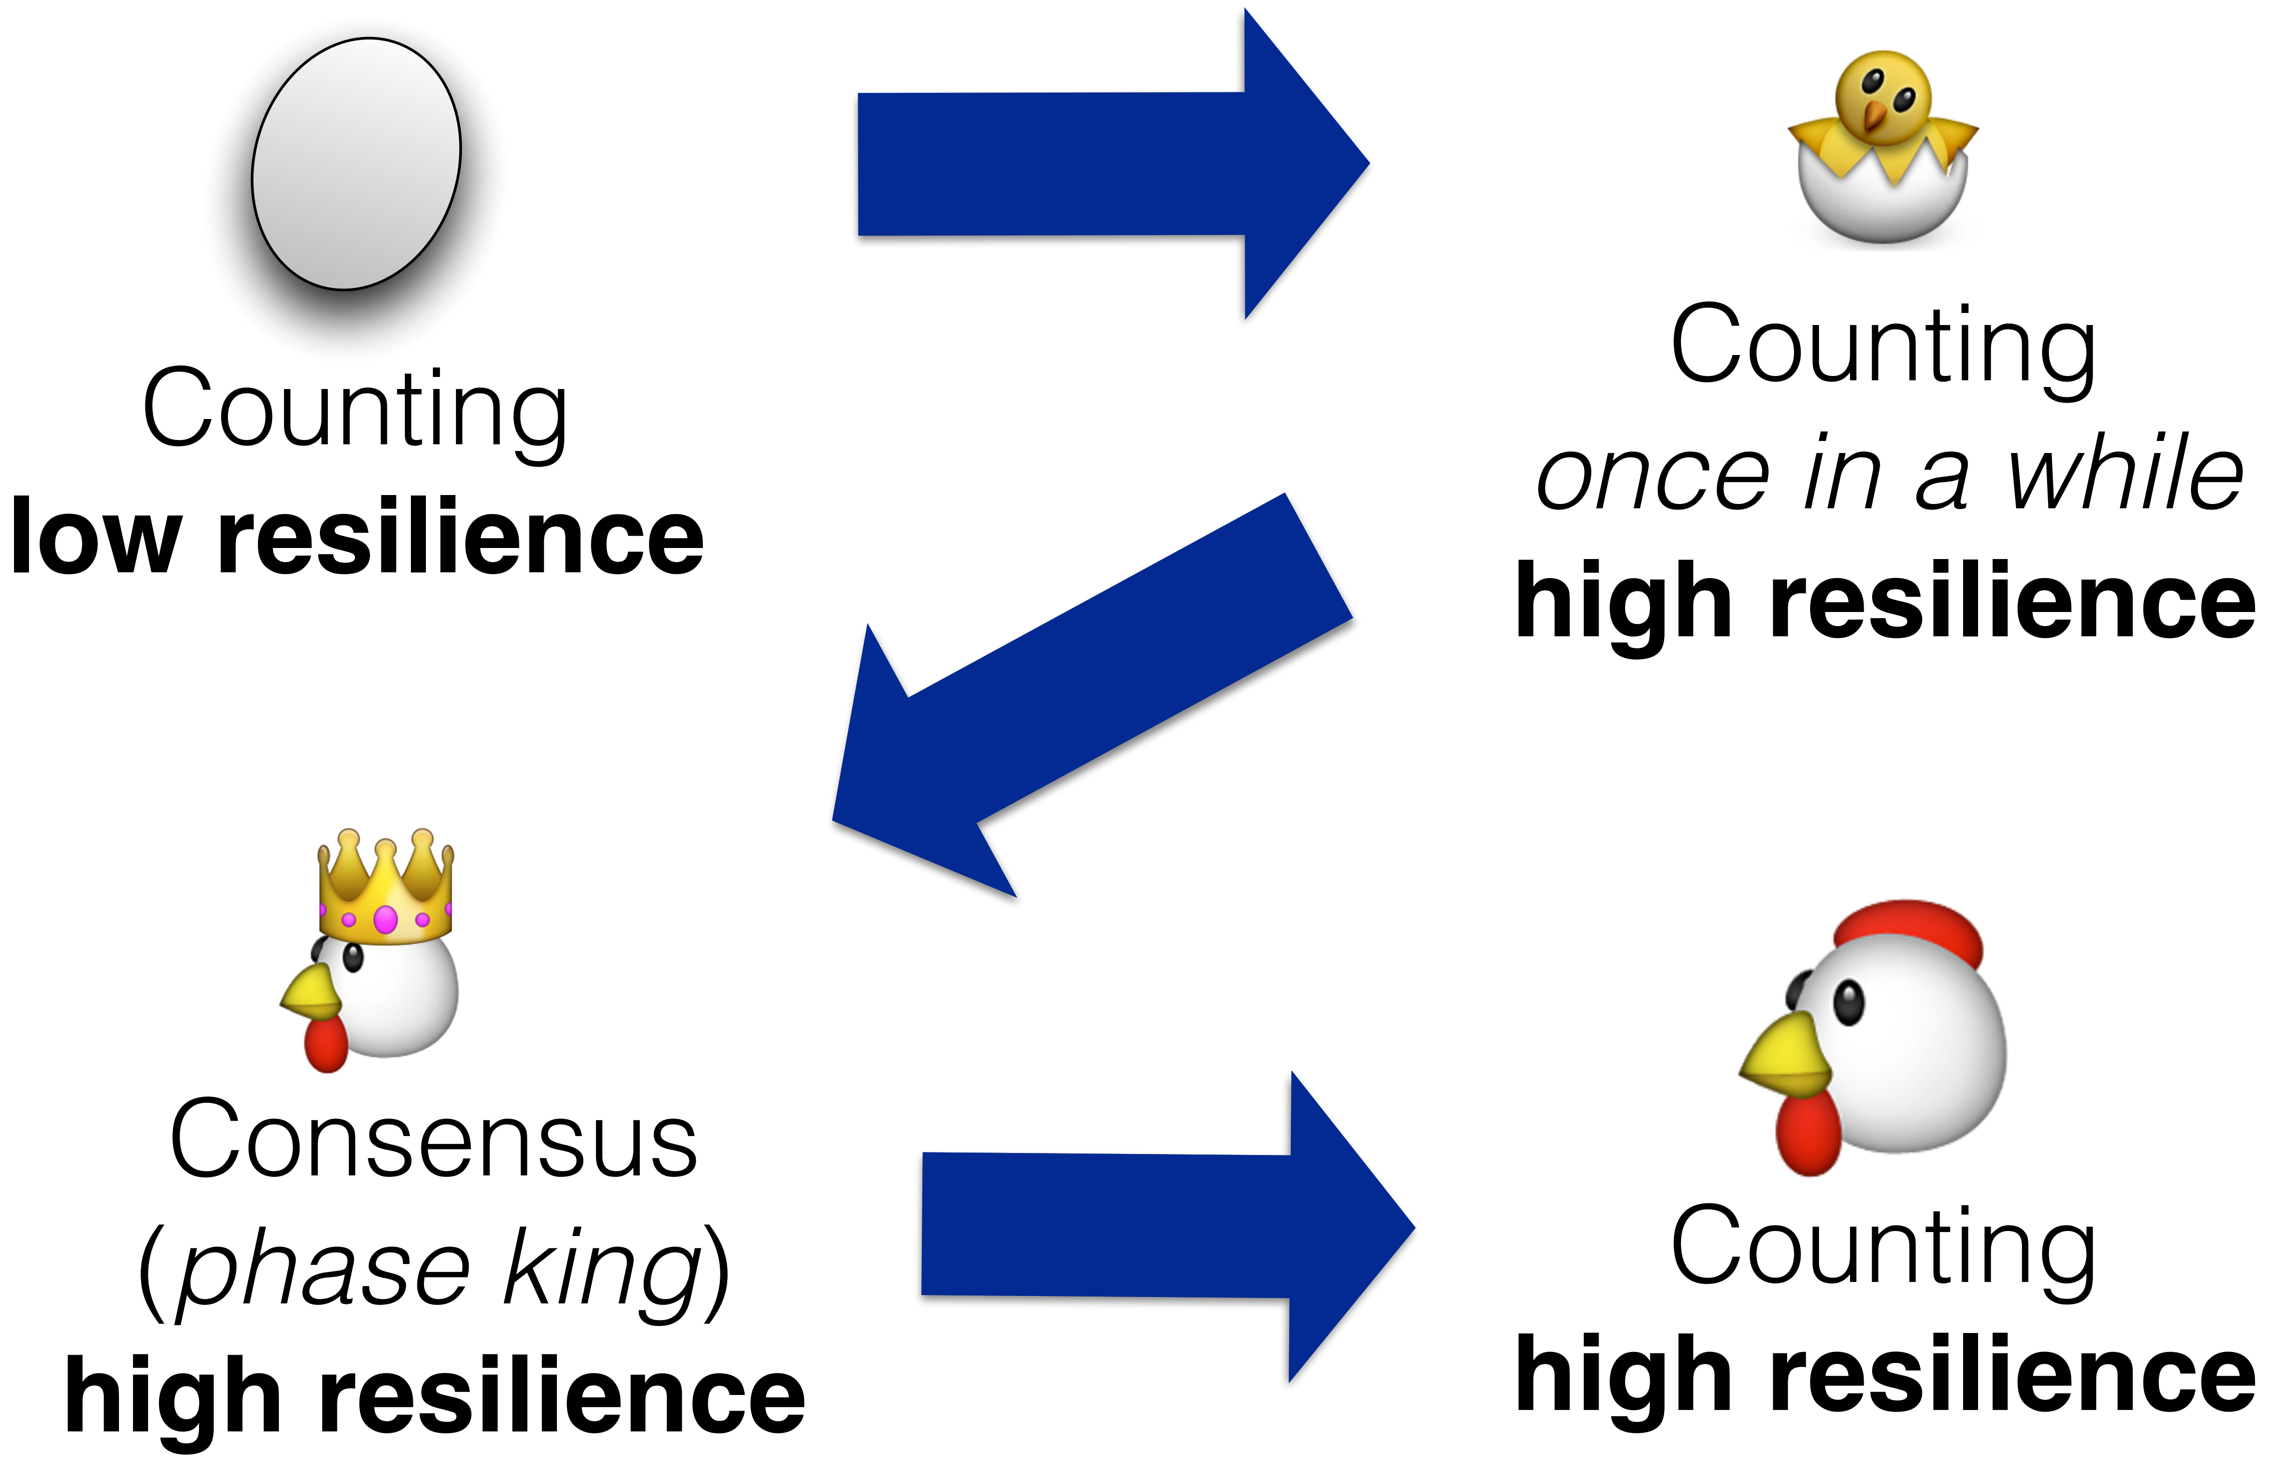

Counting once in a while

	Arbitrary			Counting			Arbitrary		
1	1	0	2	0	1	2	2	1	0
	?	?	?	?	?	?	?	?	?
3	2	2	2	0	1	2	0	2	2
4	1	2	1	0	1	2	1	1	2

High-level idea

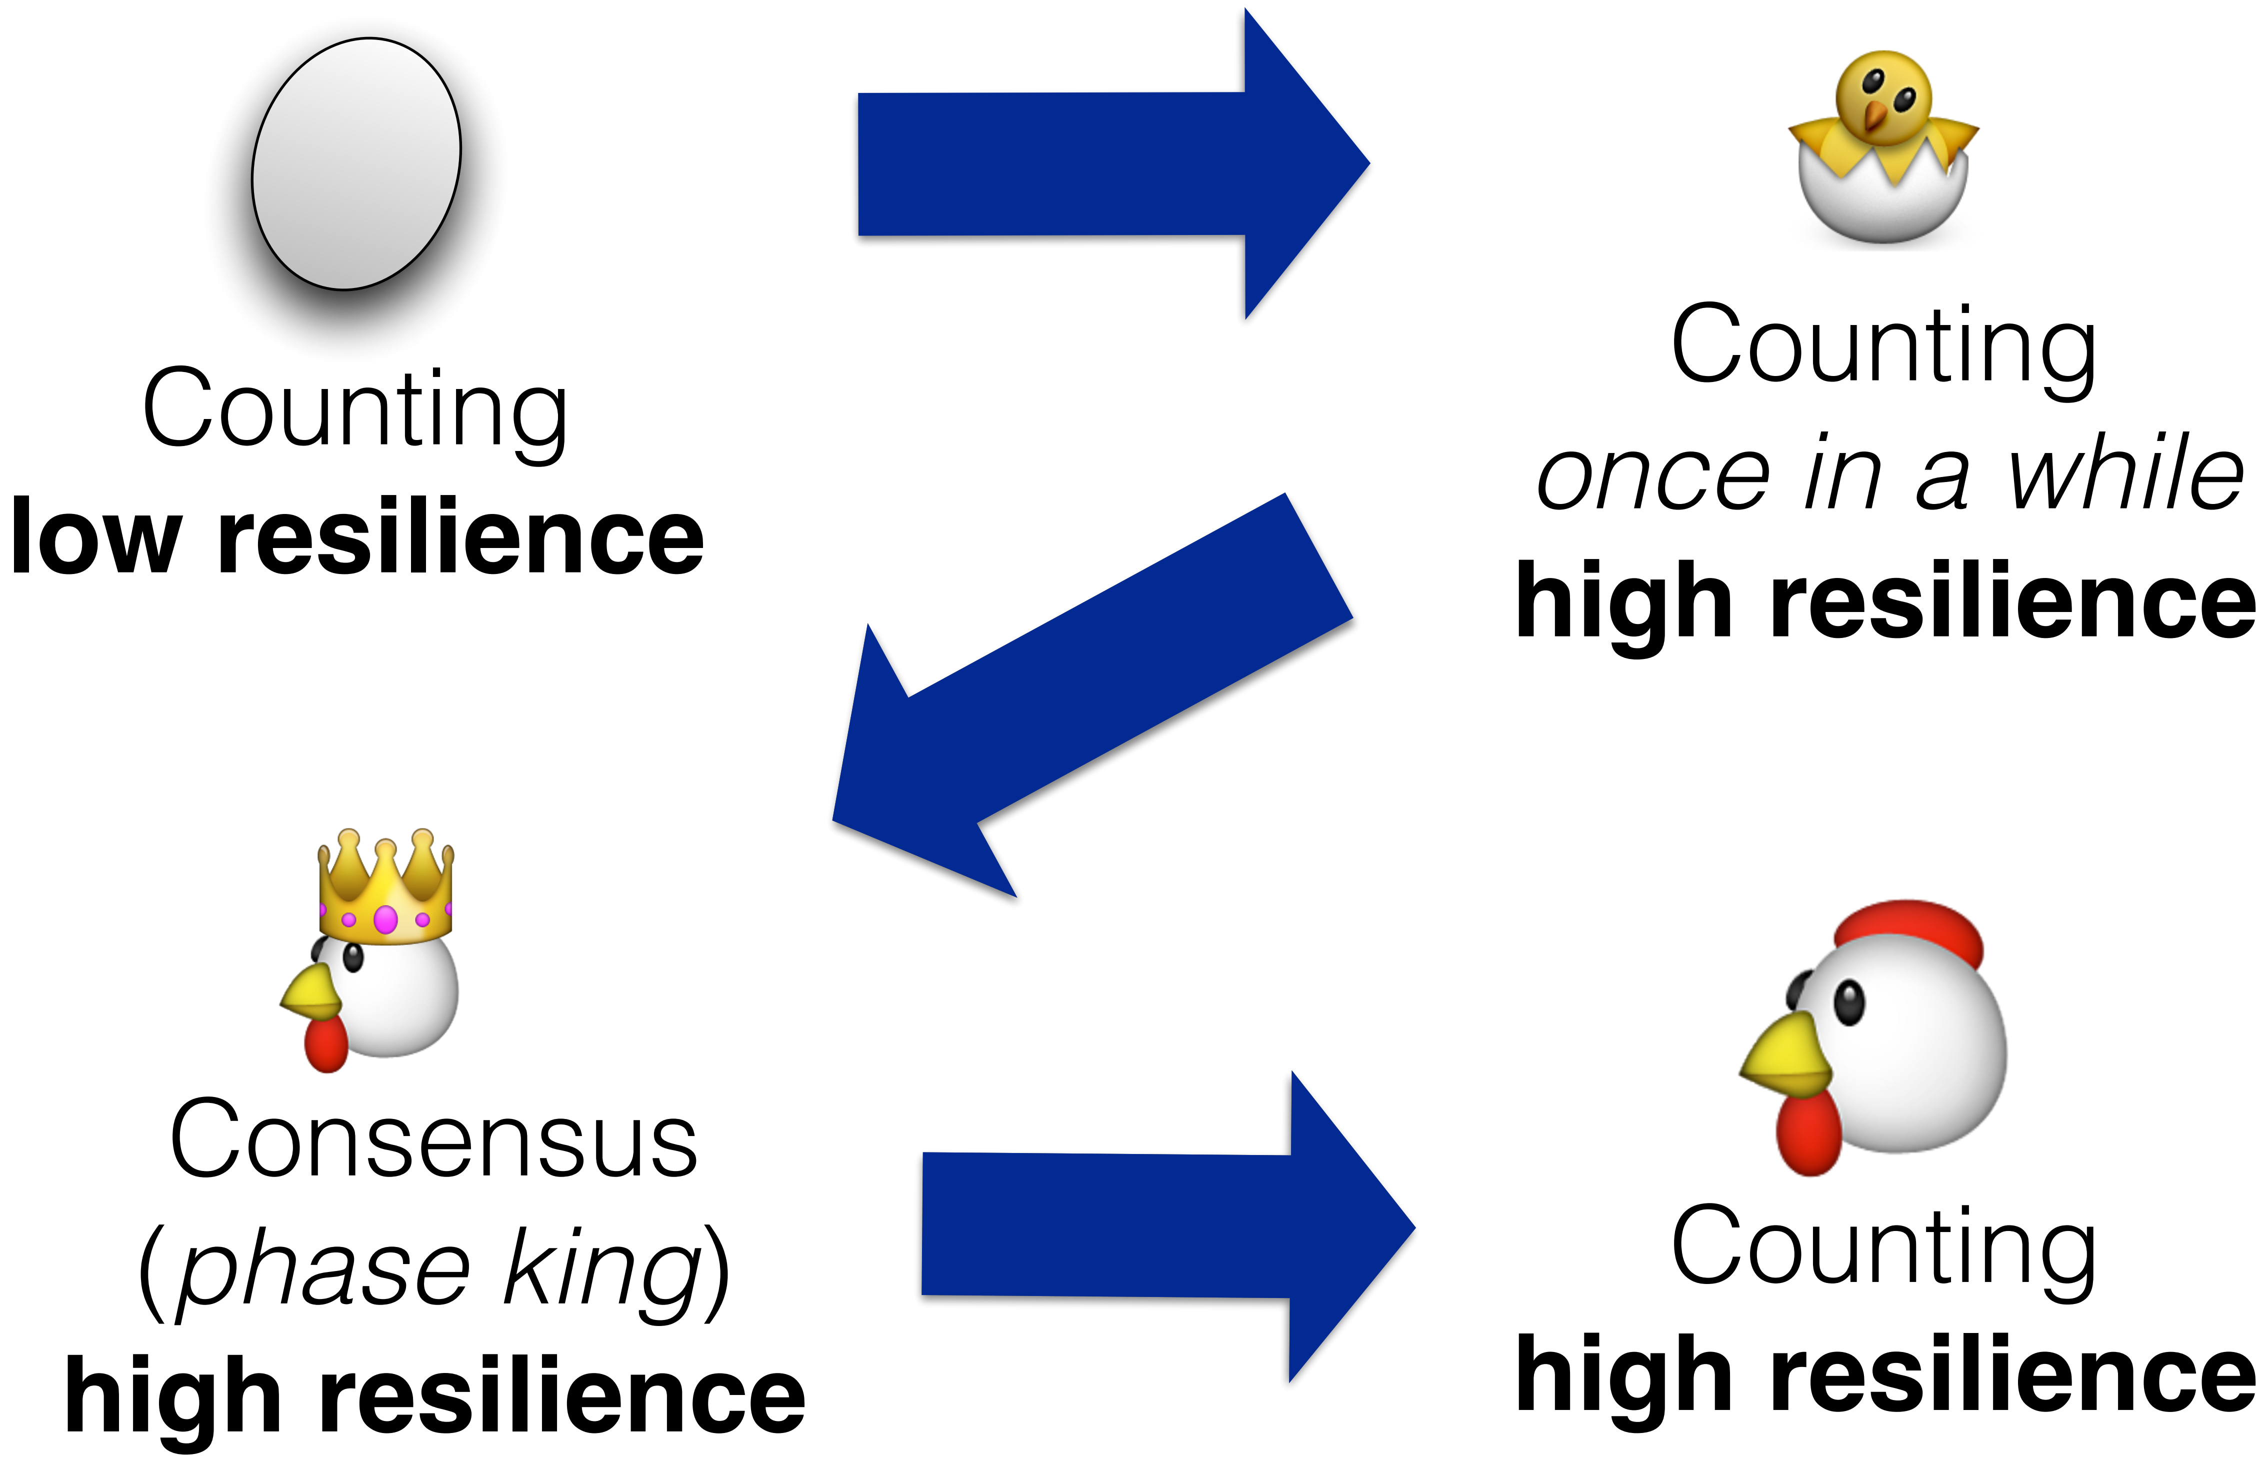

	Time	**Resilience**	**Bits**
D. Dolev & Hoch 2007	$\Theta(f)$	$< n/3$	$\Omega(f)$
PODC 2015	$\Theta(f)$	$n^{1-o(1)}$	$O(\log^2 f)$
This result	$\Theta(f)$	$< n/3$	$O(\log^2 f)$

Previous boosting lemma

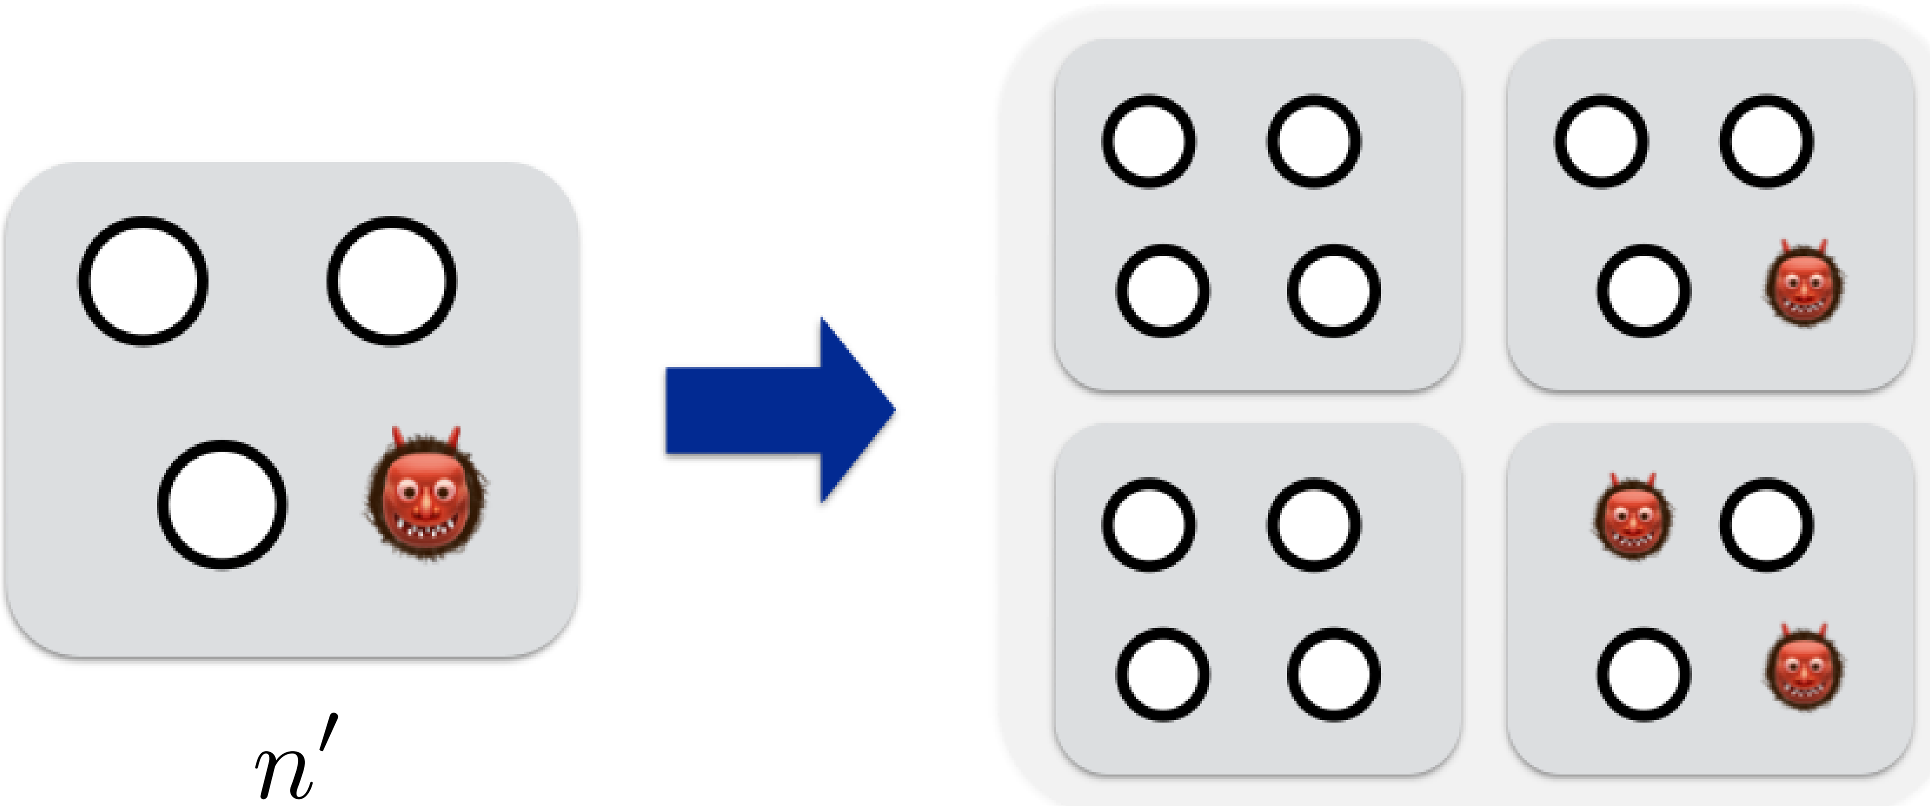

n'

f'

$n = kn'$

$f \approx (f' + 1)k/2$

Stabilisation time: T → $T + O(k^k \cdot f)$

Message size: S → $S + O(\log f)$

New boosting lemma

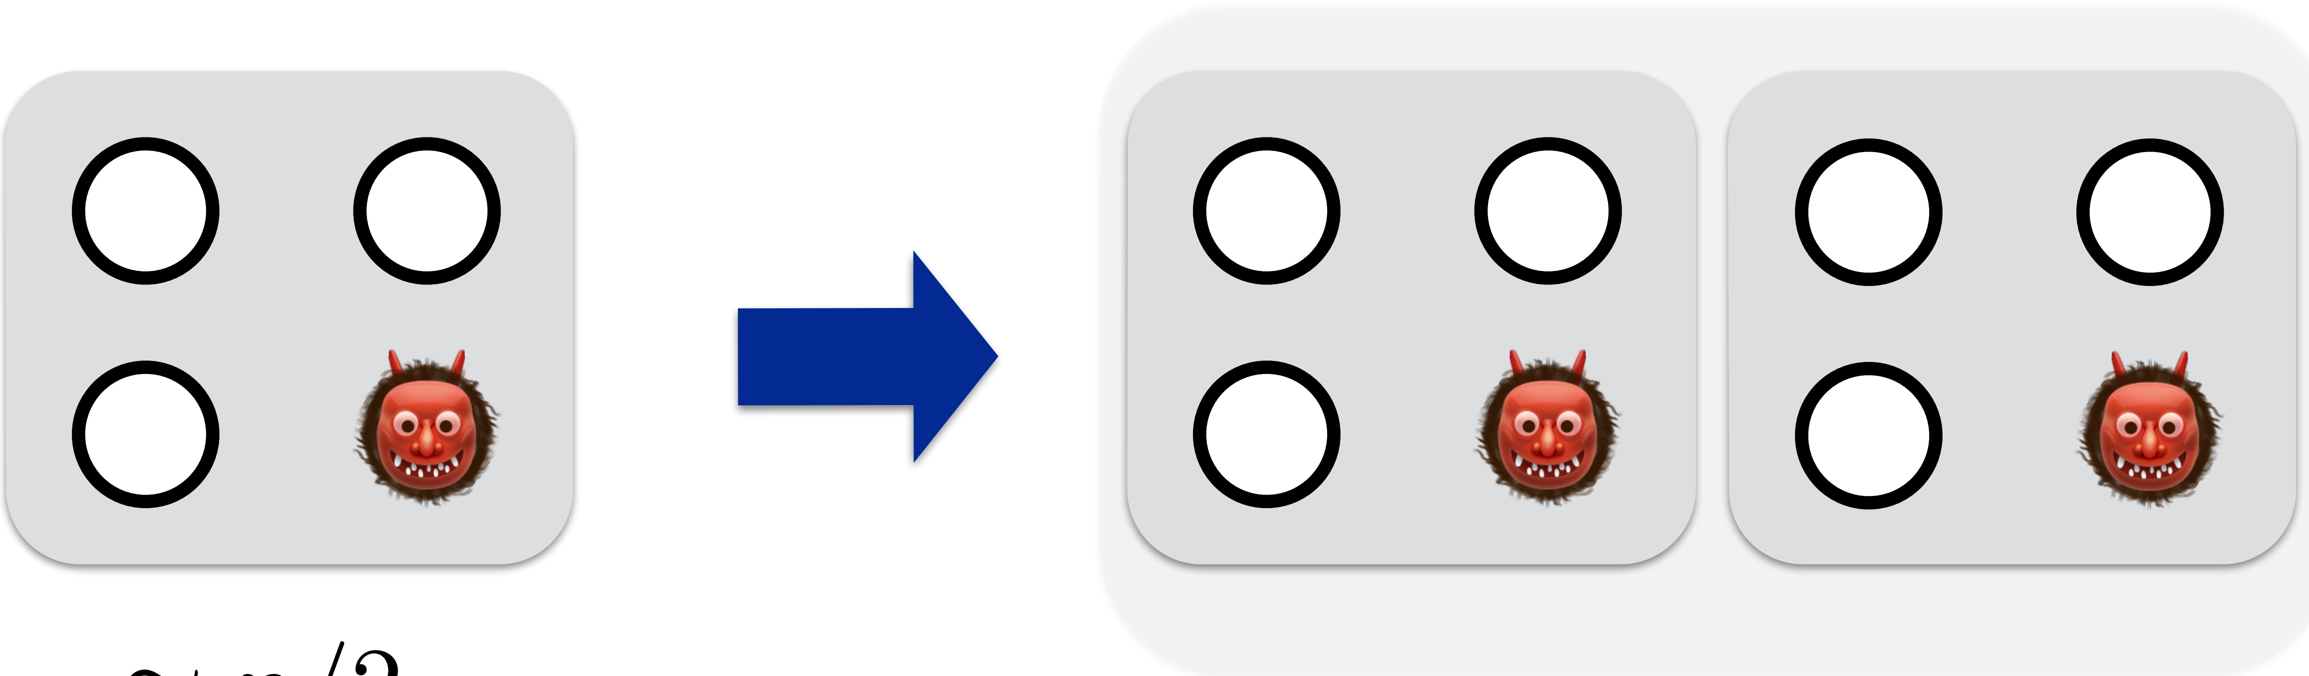

$\approx n/2$
$\approx f/2$

n, f

Stabilisation time: T → $T + O(f)$

Message size: S → $S + O(\log f)$

Resilience does not degrade!

Boosting resilience

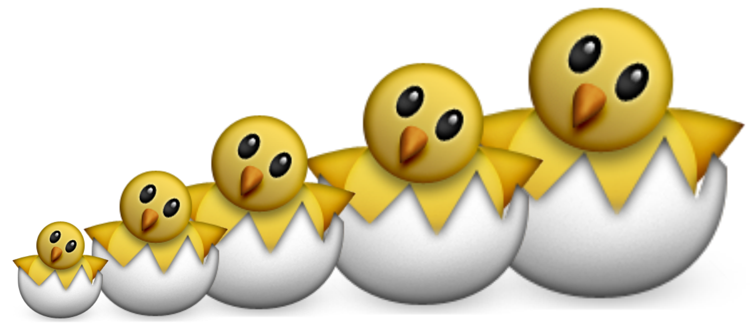

Boost resilience recursively
for log f steps

Stabilisation time: $O(f)$
Message size: $O(\log^2 f + \log c)$

Summary

	Time	Bits
 S. Dolev & Welch 2004	$2^{2(n-f)}$	$O(1)$
 Ben-Or *et al.* 2008	$O(1)$	$n^{O(1)}$
This paper (deterministic)	$\Theta(f)$	$O(\log^2 f)$

Reducing communication

Each node broadcasts

$$O(\log^2 f + \log c)$$

bits *during* stabilisation.

What about *after* stabilisation?

Quiet poly-counters

If $c = n^{O(1)}$ is a multiple of n then we get

- **optimal** stabilisation time
- **optimal** resilience

..and *after stabilisation* each node broadcasts **optimal** $O(1)$ bits every $\Theta(n)$ rounds.

Summary

	Time	Bits
 S. Dolev & Welch 2004	$2^{2(n-f)}$	$O(1)$
 Ben-Or *et al.* 2008	$O(1)$	$n^{O(1)}$
This paper (deterministic)	$\Theta(f)$	$O(\log^2 f)$

Thanks!