

XCO.T473 サイバーセキュリティ概論

<https://www.coord.c.titech.ac.jp/c/cybersec>



Intrusion-Tolerant

Computer Systems

耐侵入コンピュータシステム



Xavier Défago

*School of Computing
Tokyo Institute of Technology*

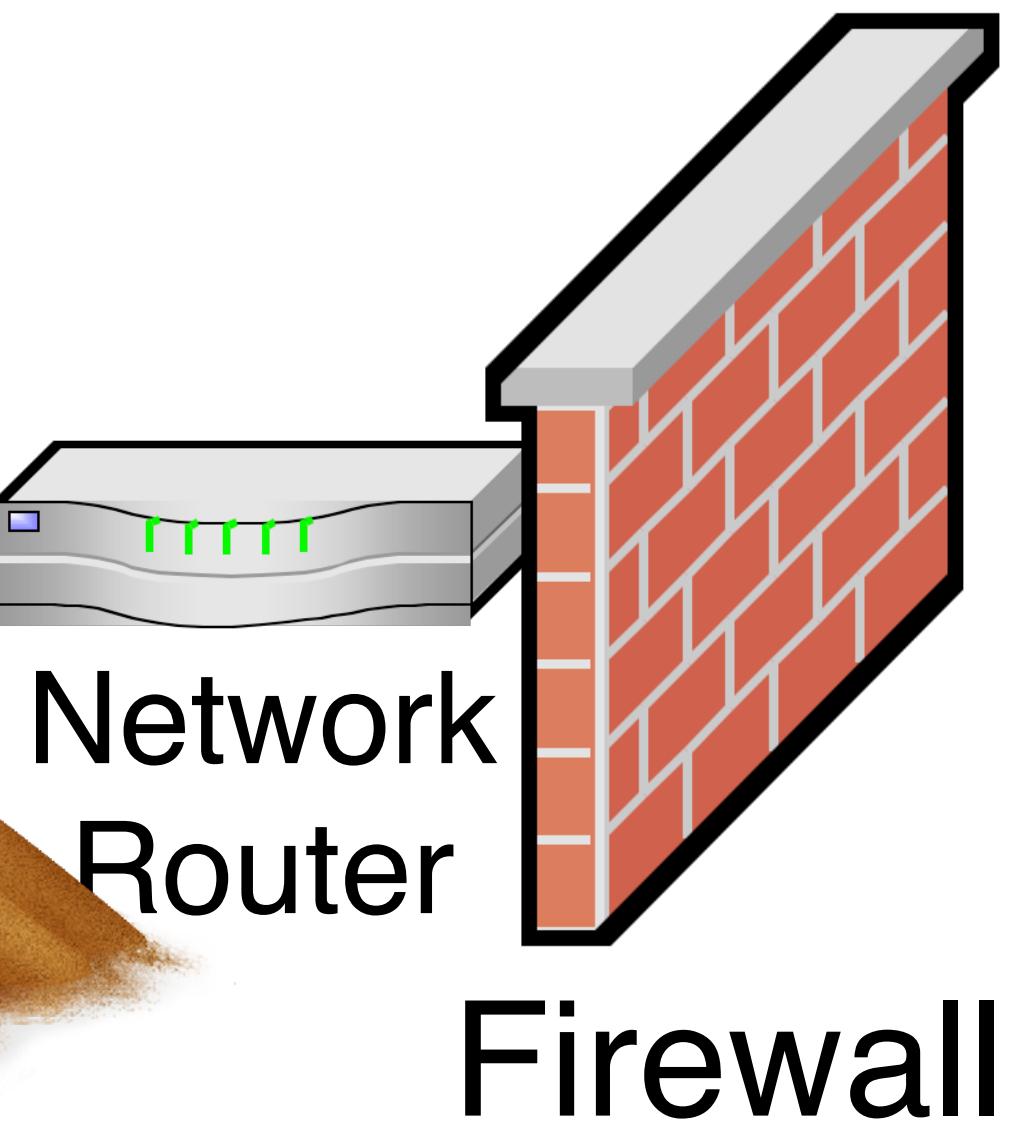
Why Intrusion-Tolerance?

Firewall Fallacy



“Inside”

comfy, safe



Network
Router

Firewall



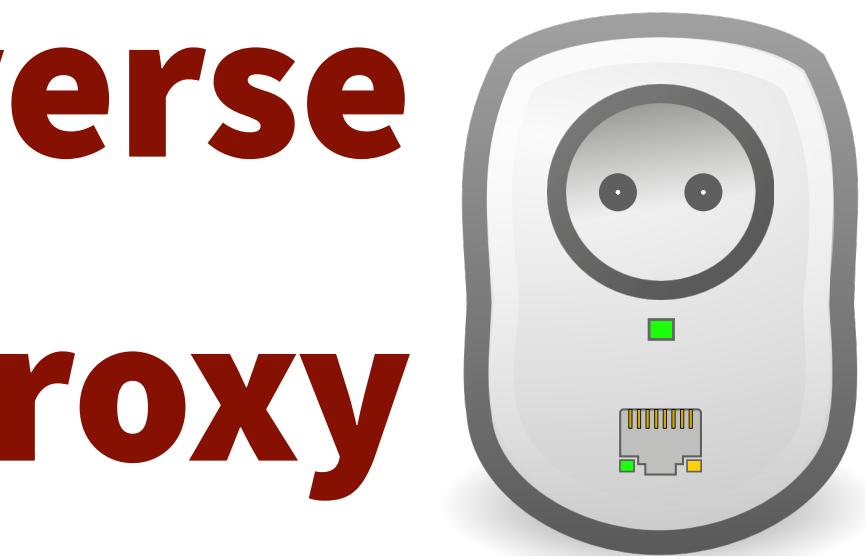
“Outside”

threatening

(issue 1) Reverse Proxy

weak smart plug

Reverse
Proxy

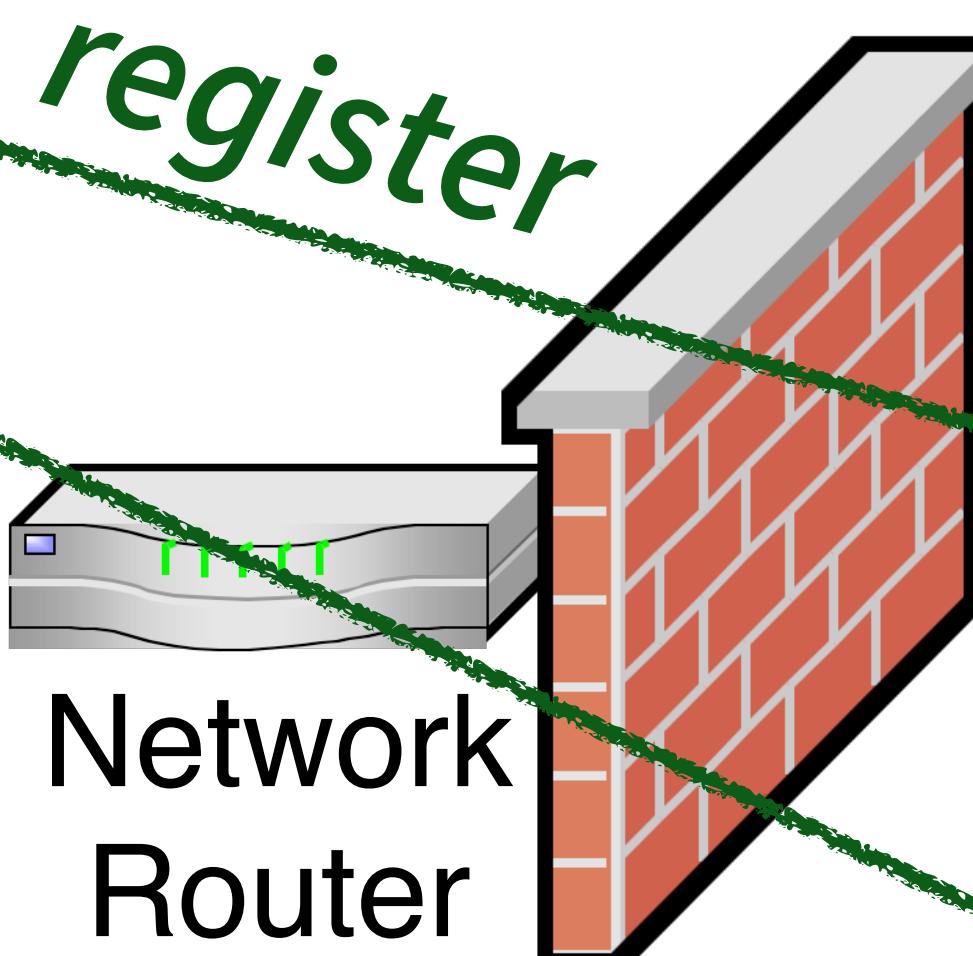


botnet

monitoring

attack

admin: admin
replay

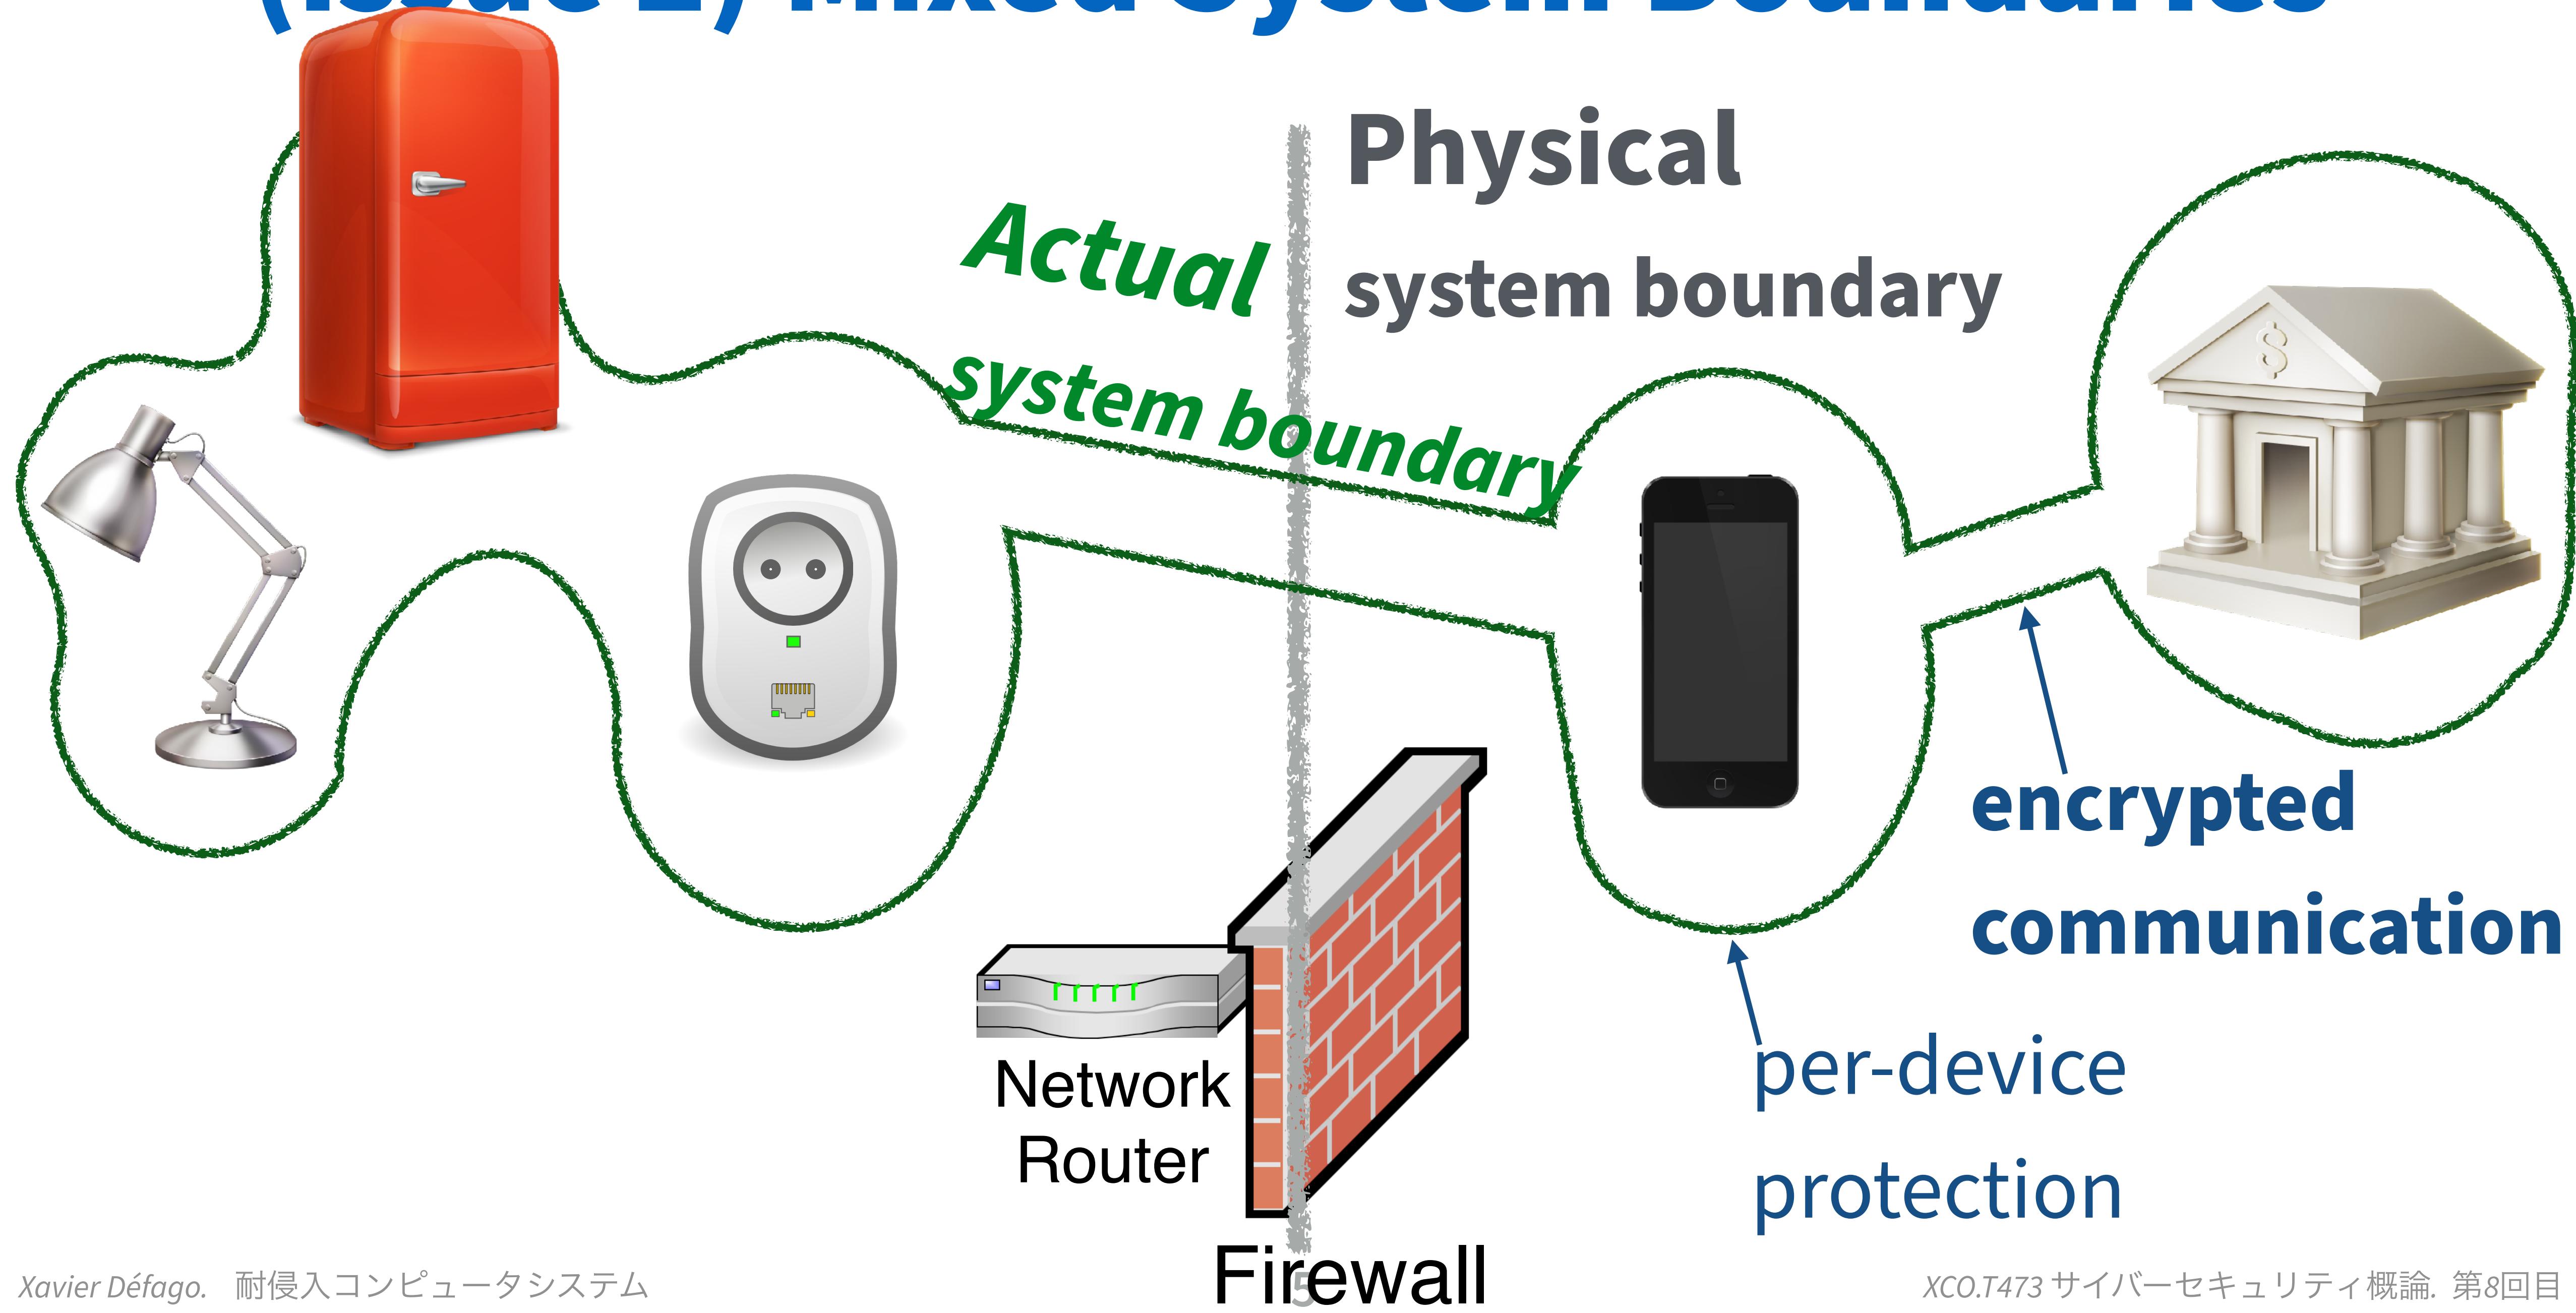


registry

query

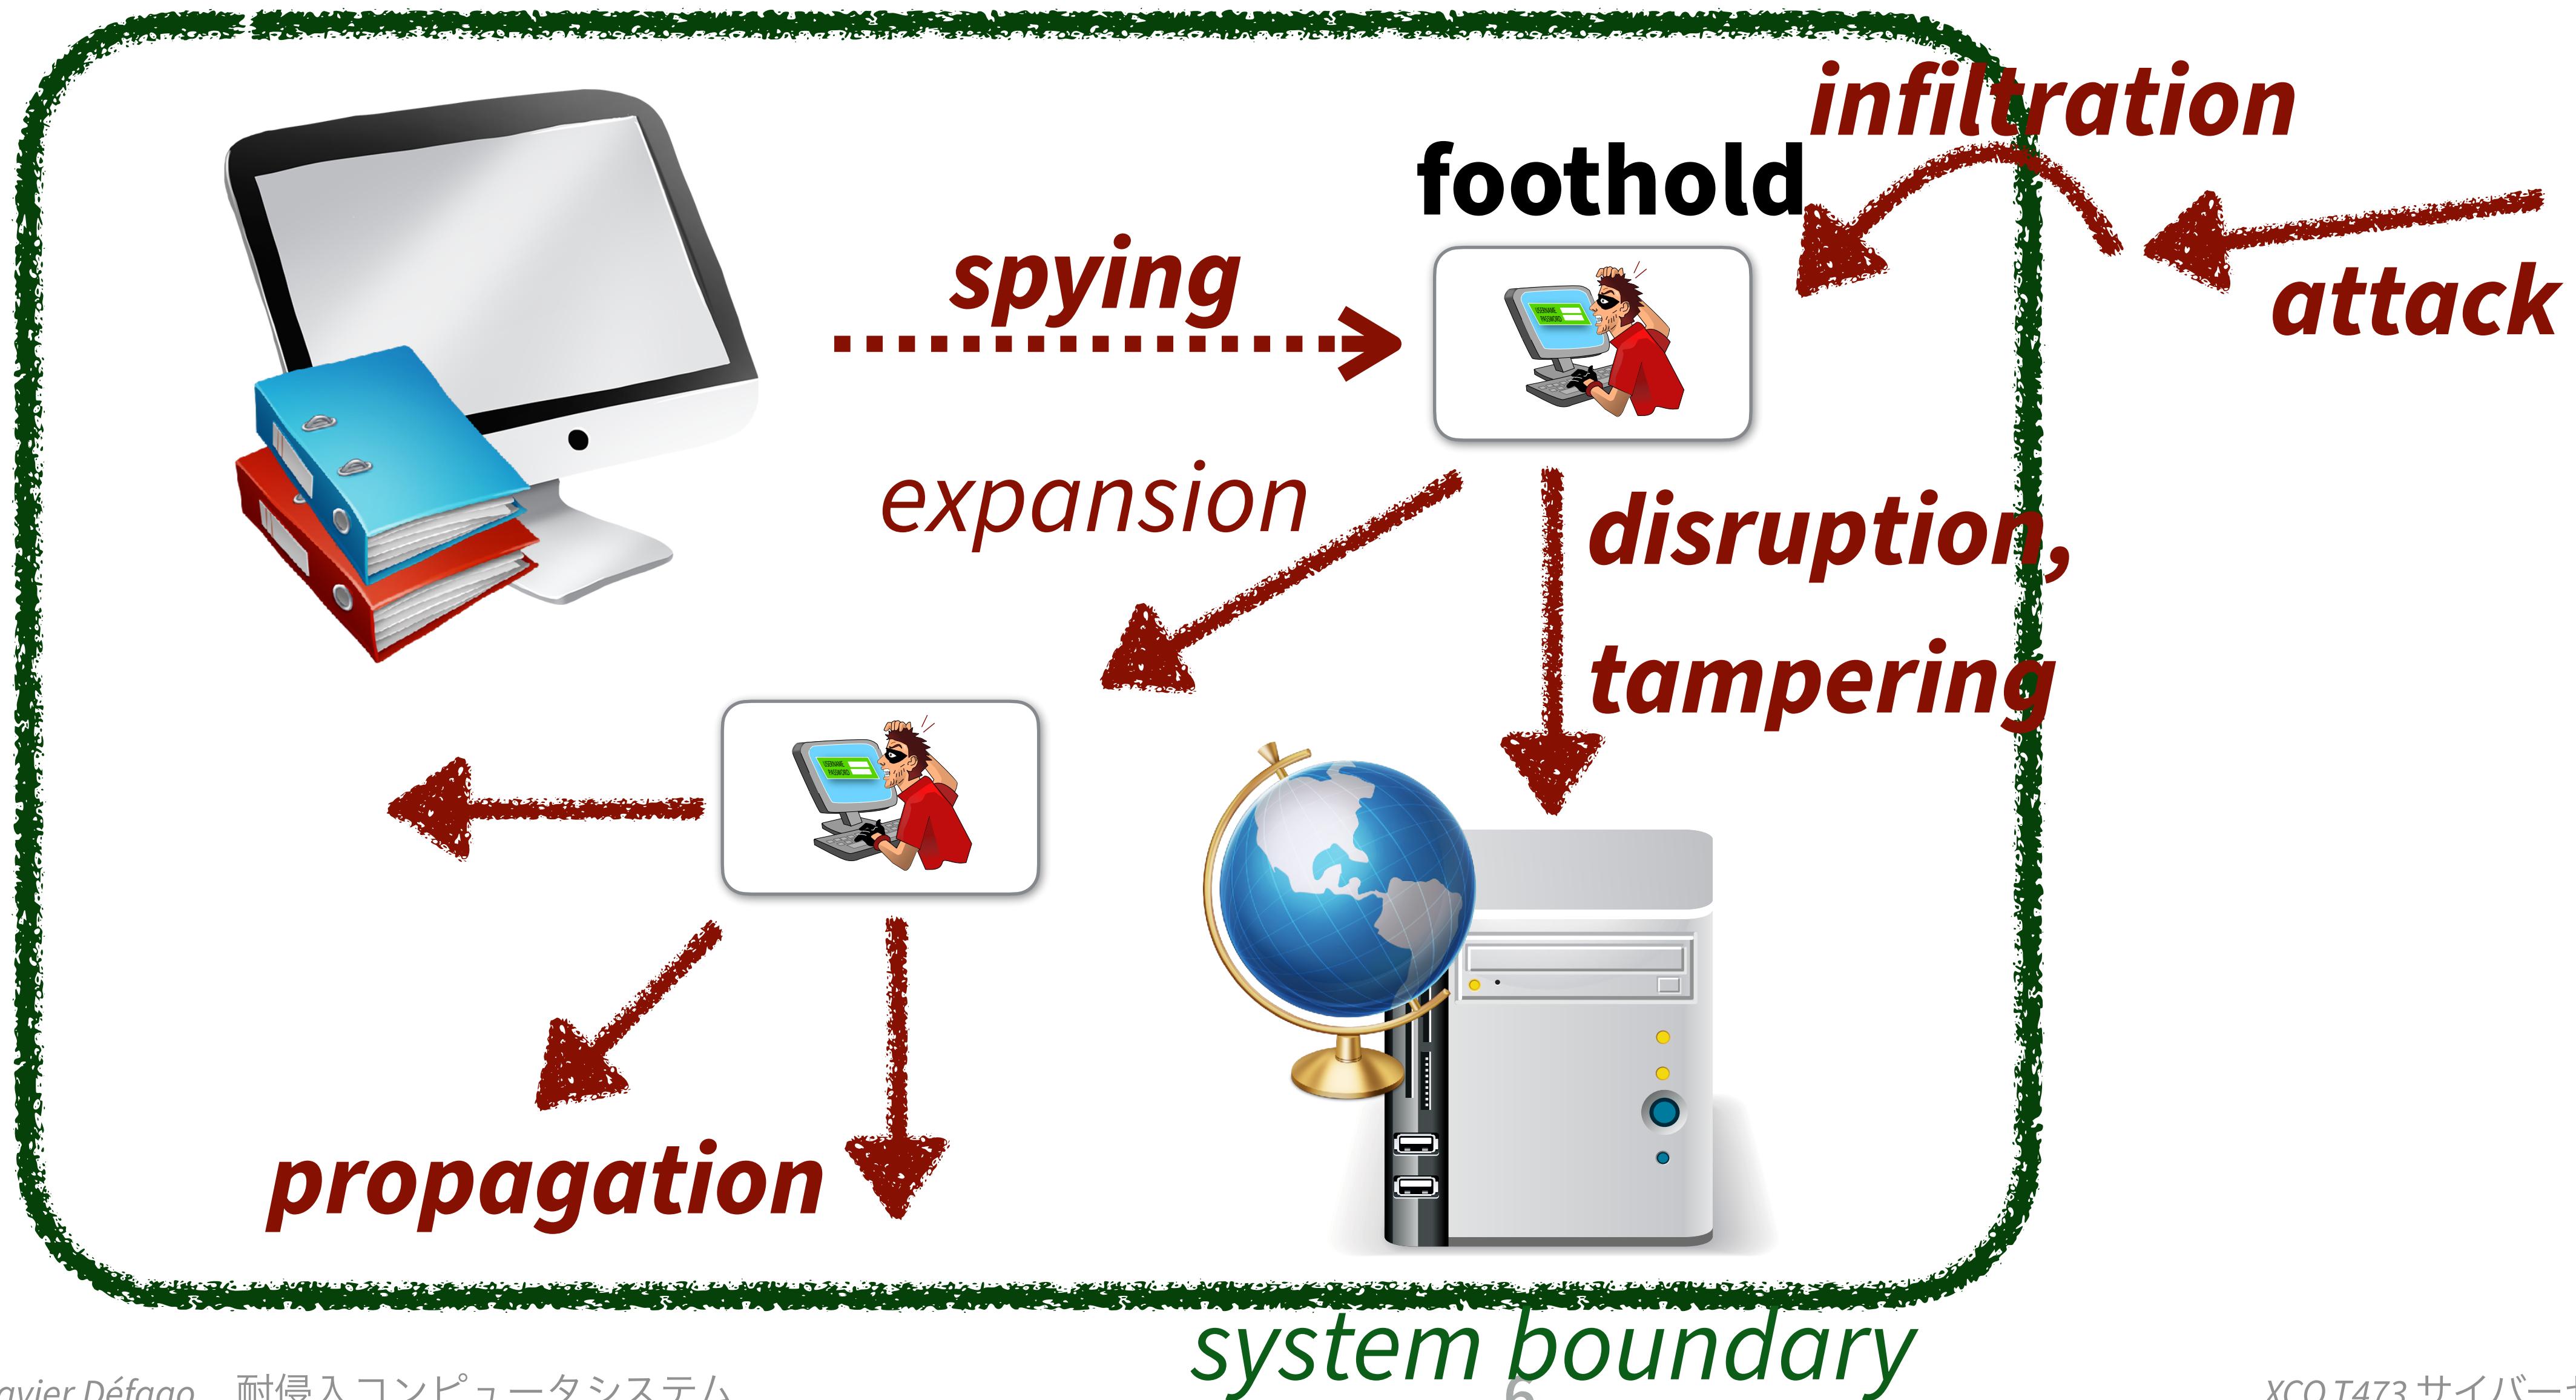


(issue 2) Mixed System Boundaries

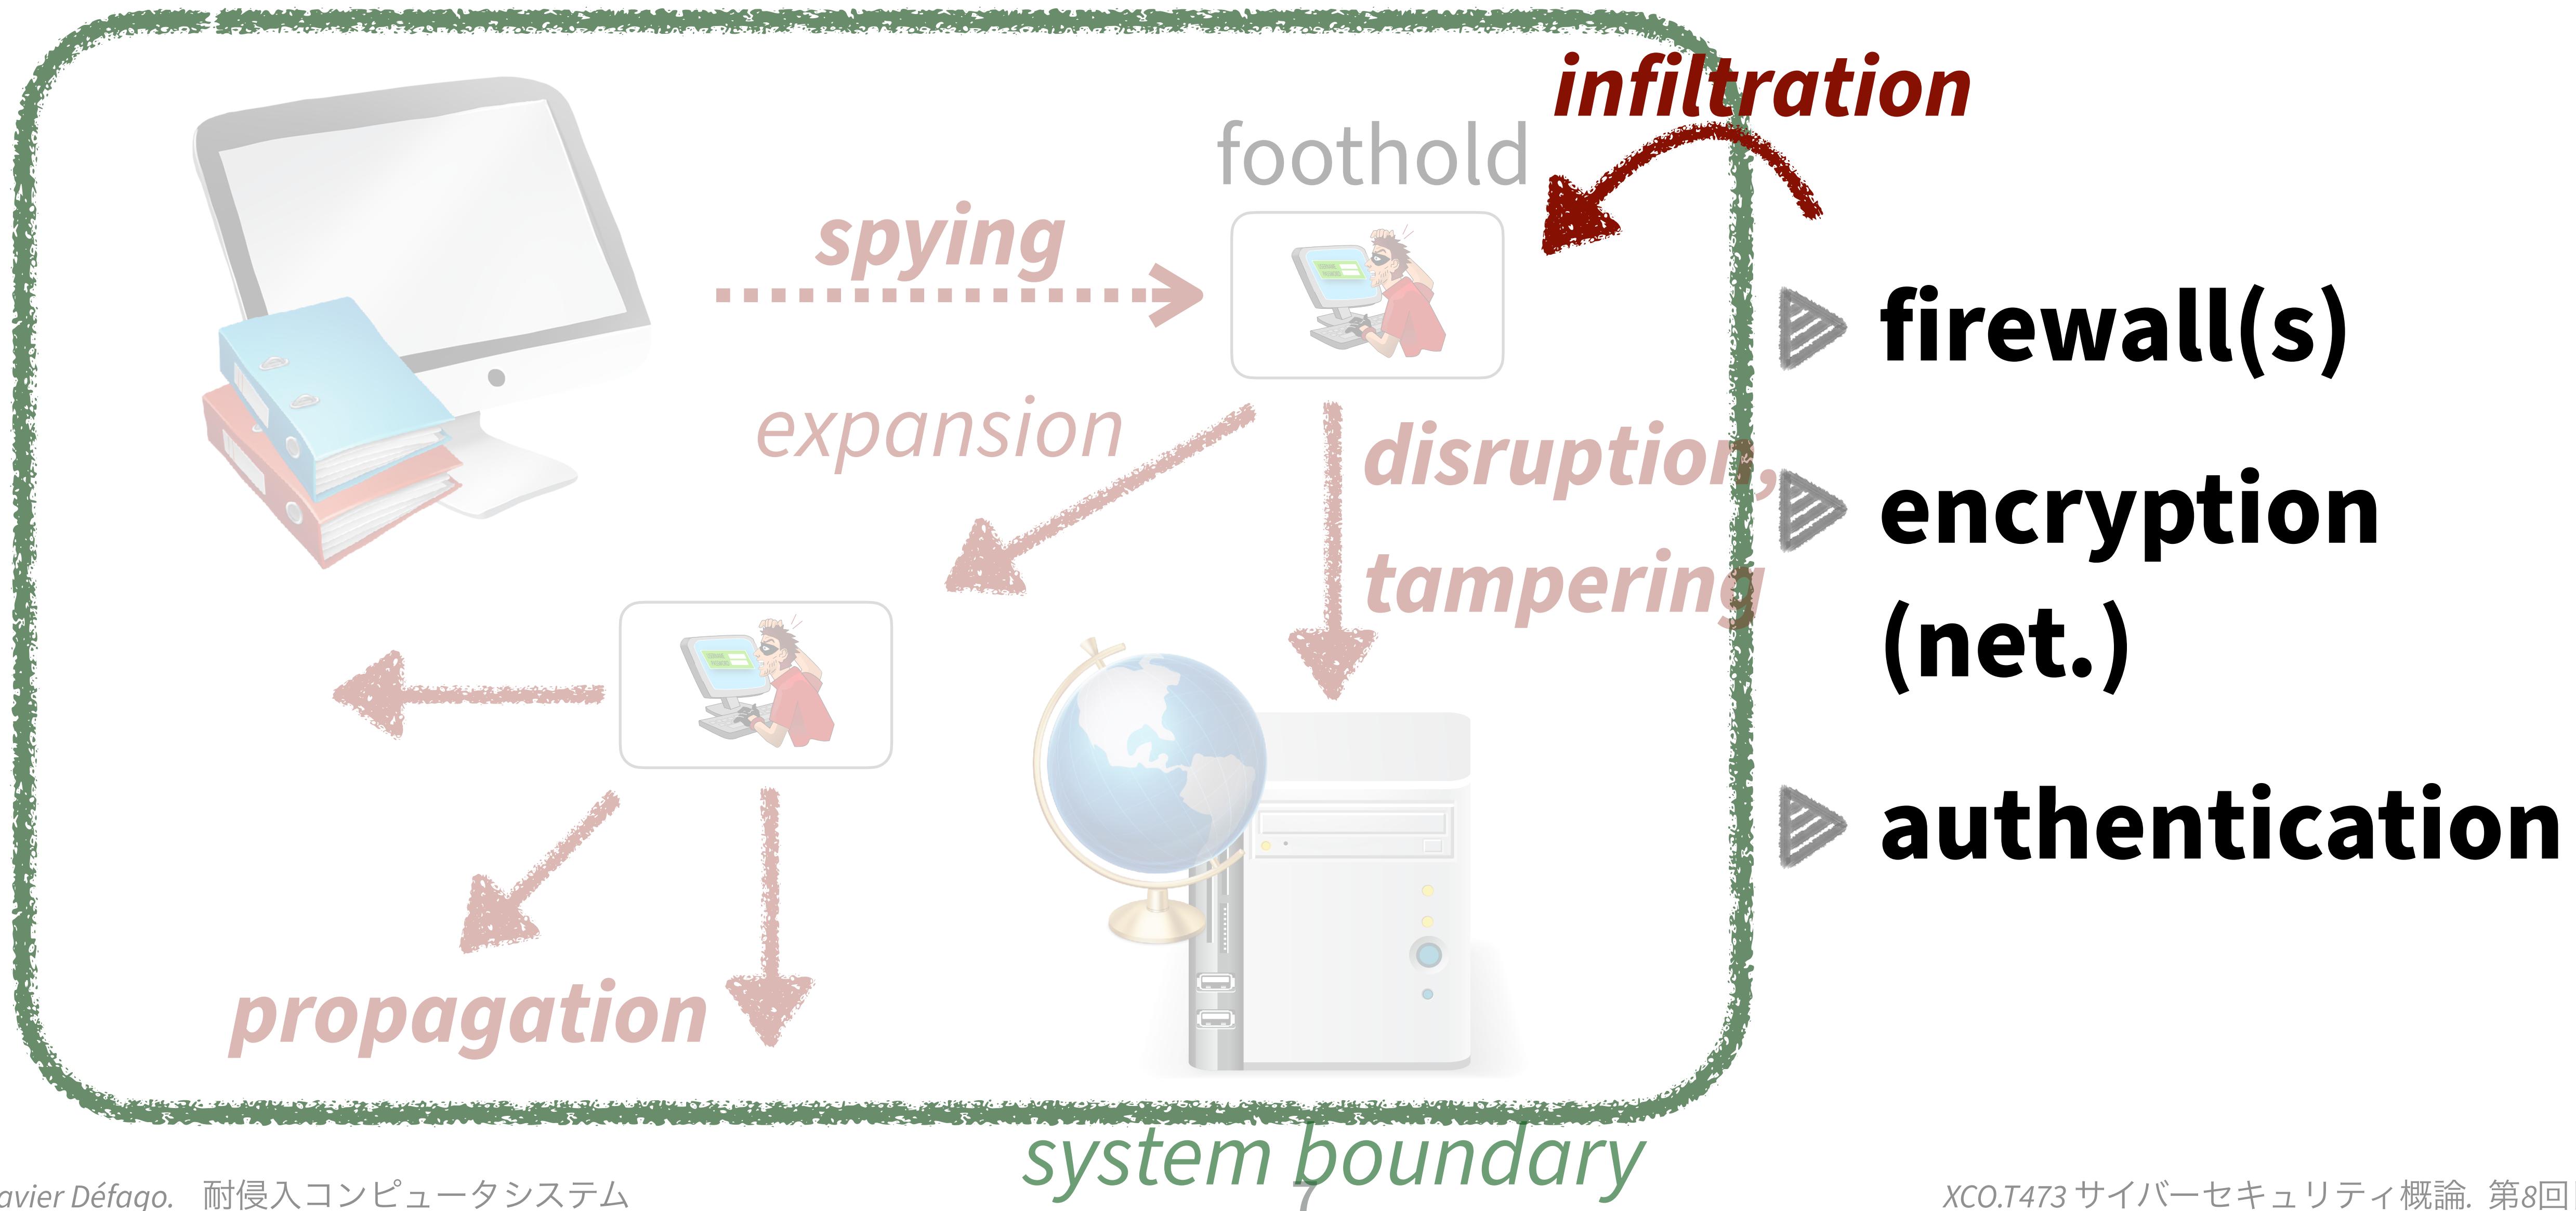


Phases of an Attack

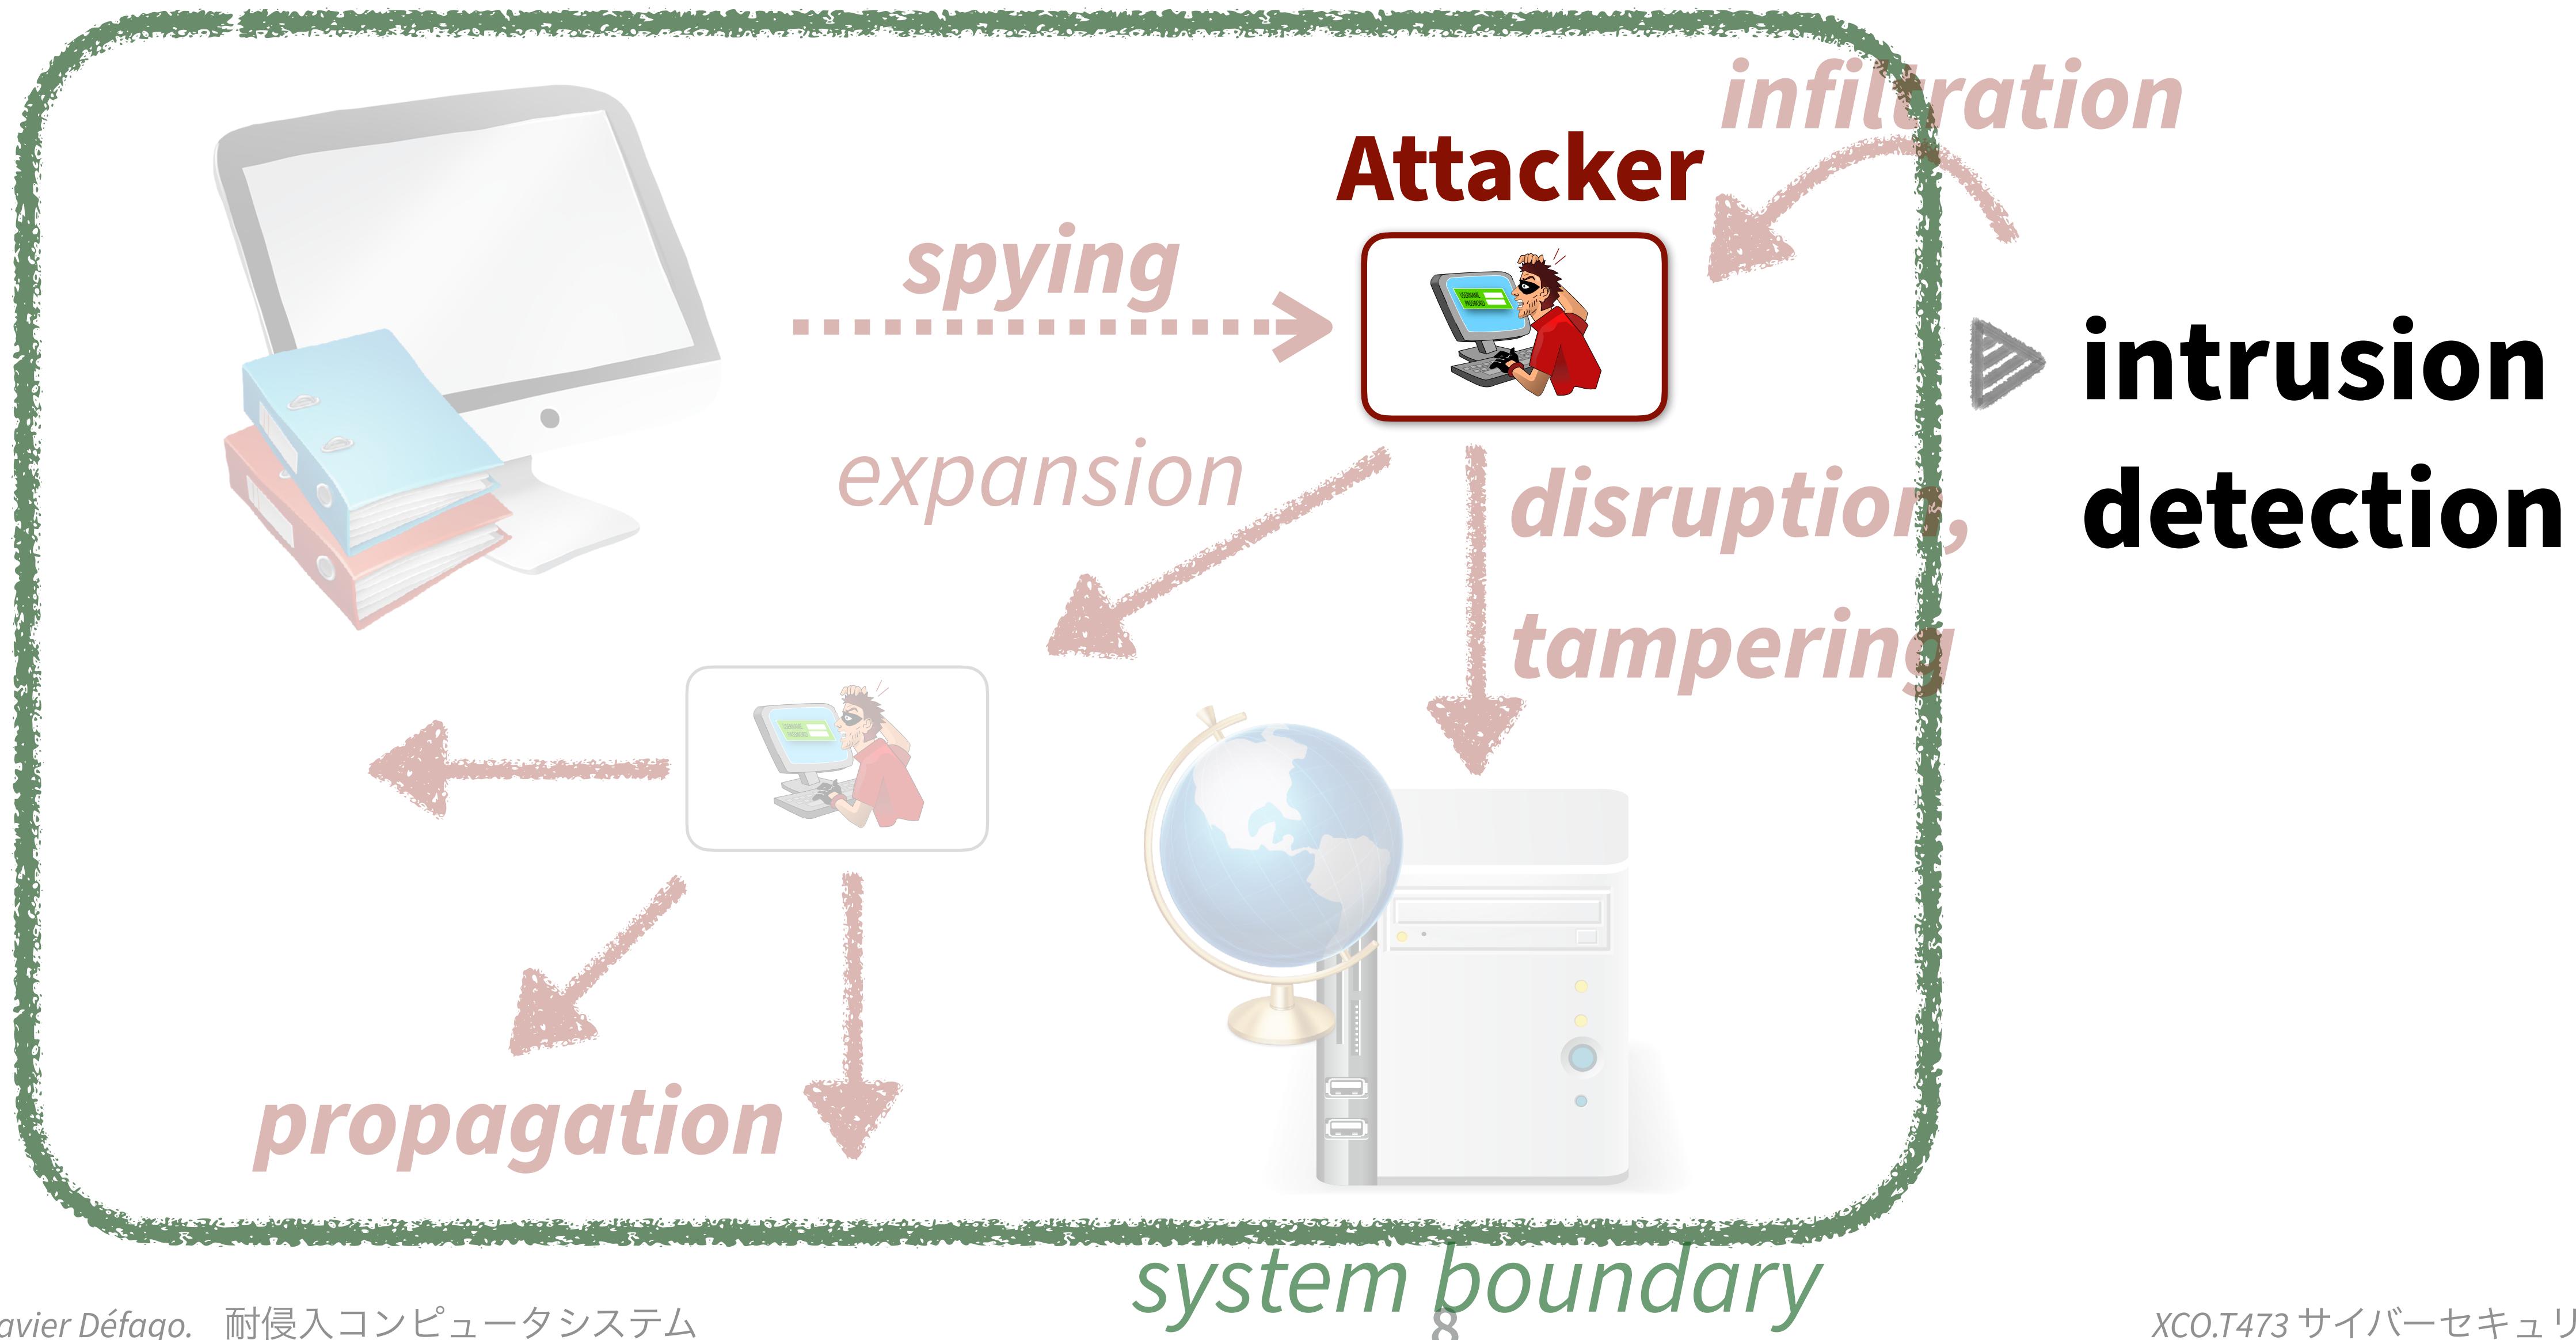
Attacker



Typical Countermeasures

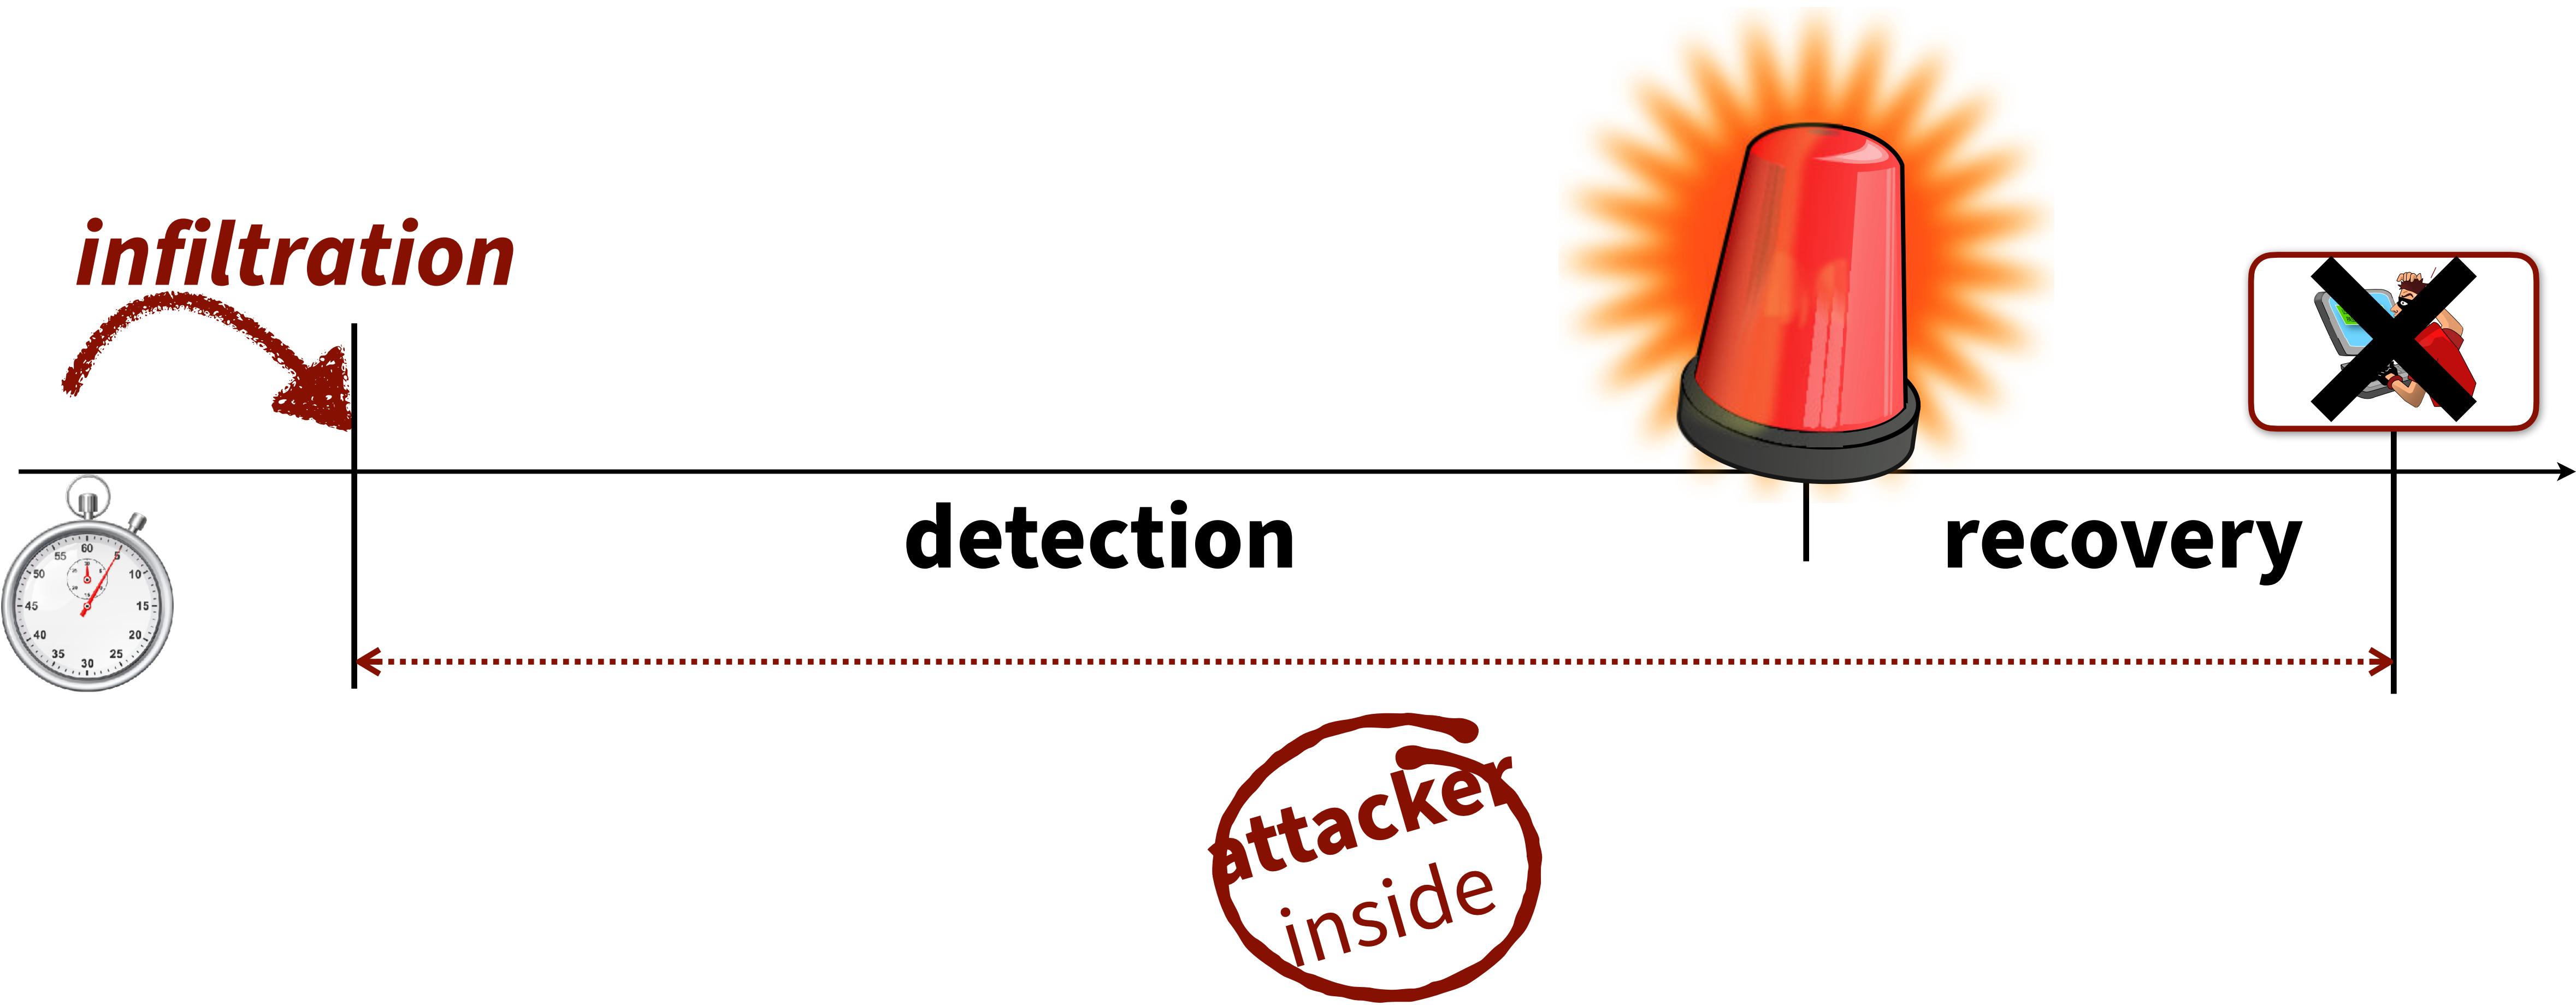


Typical Countermeasures

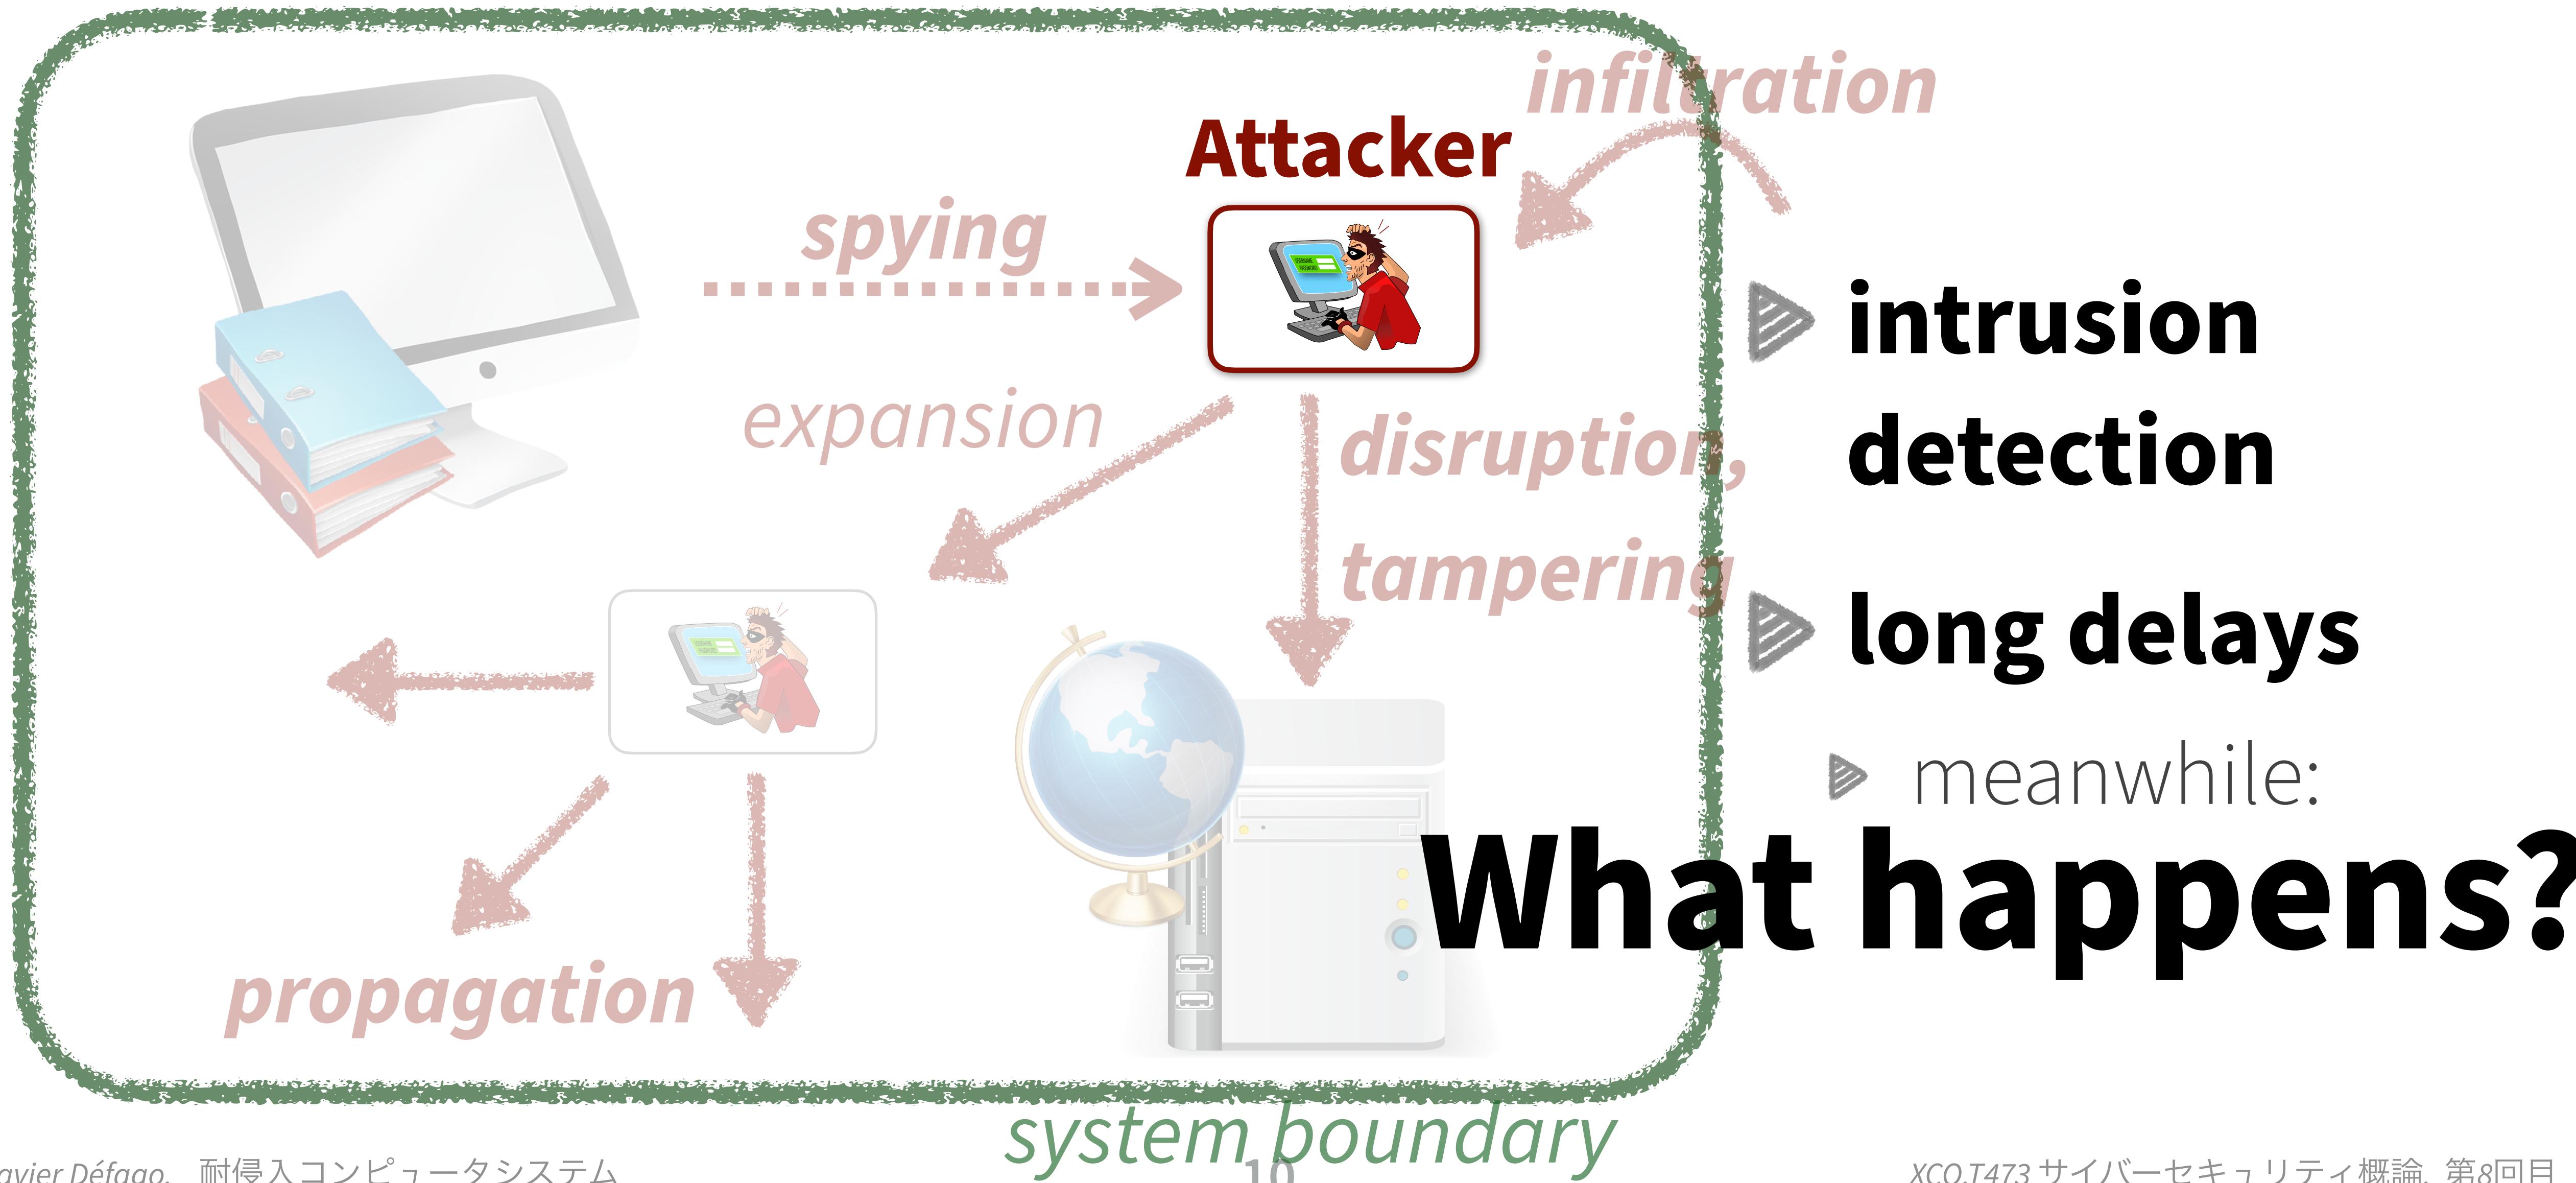


Intrusion Detection

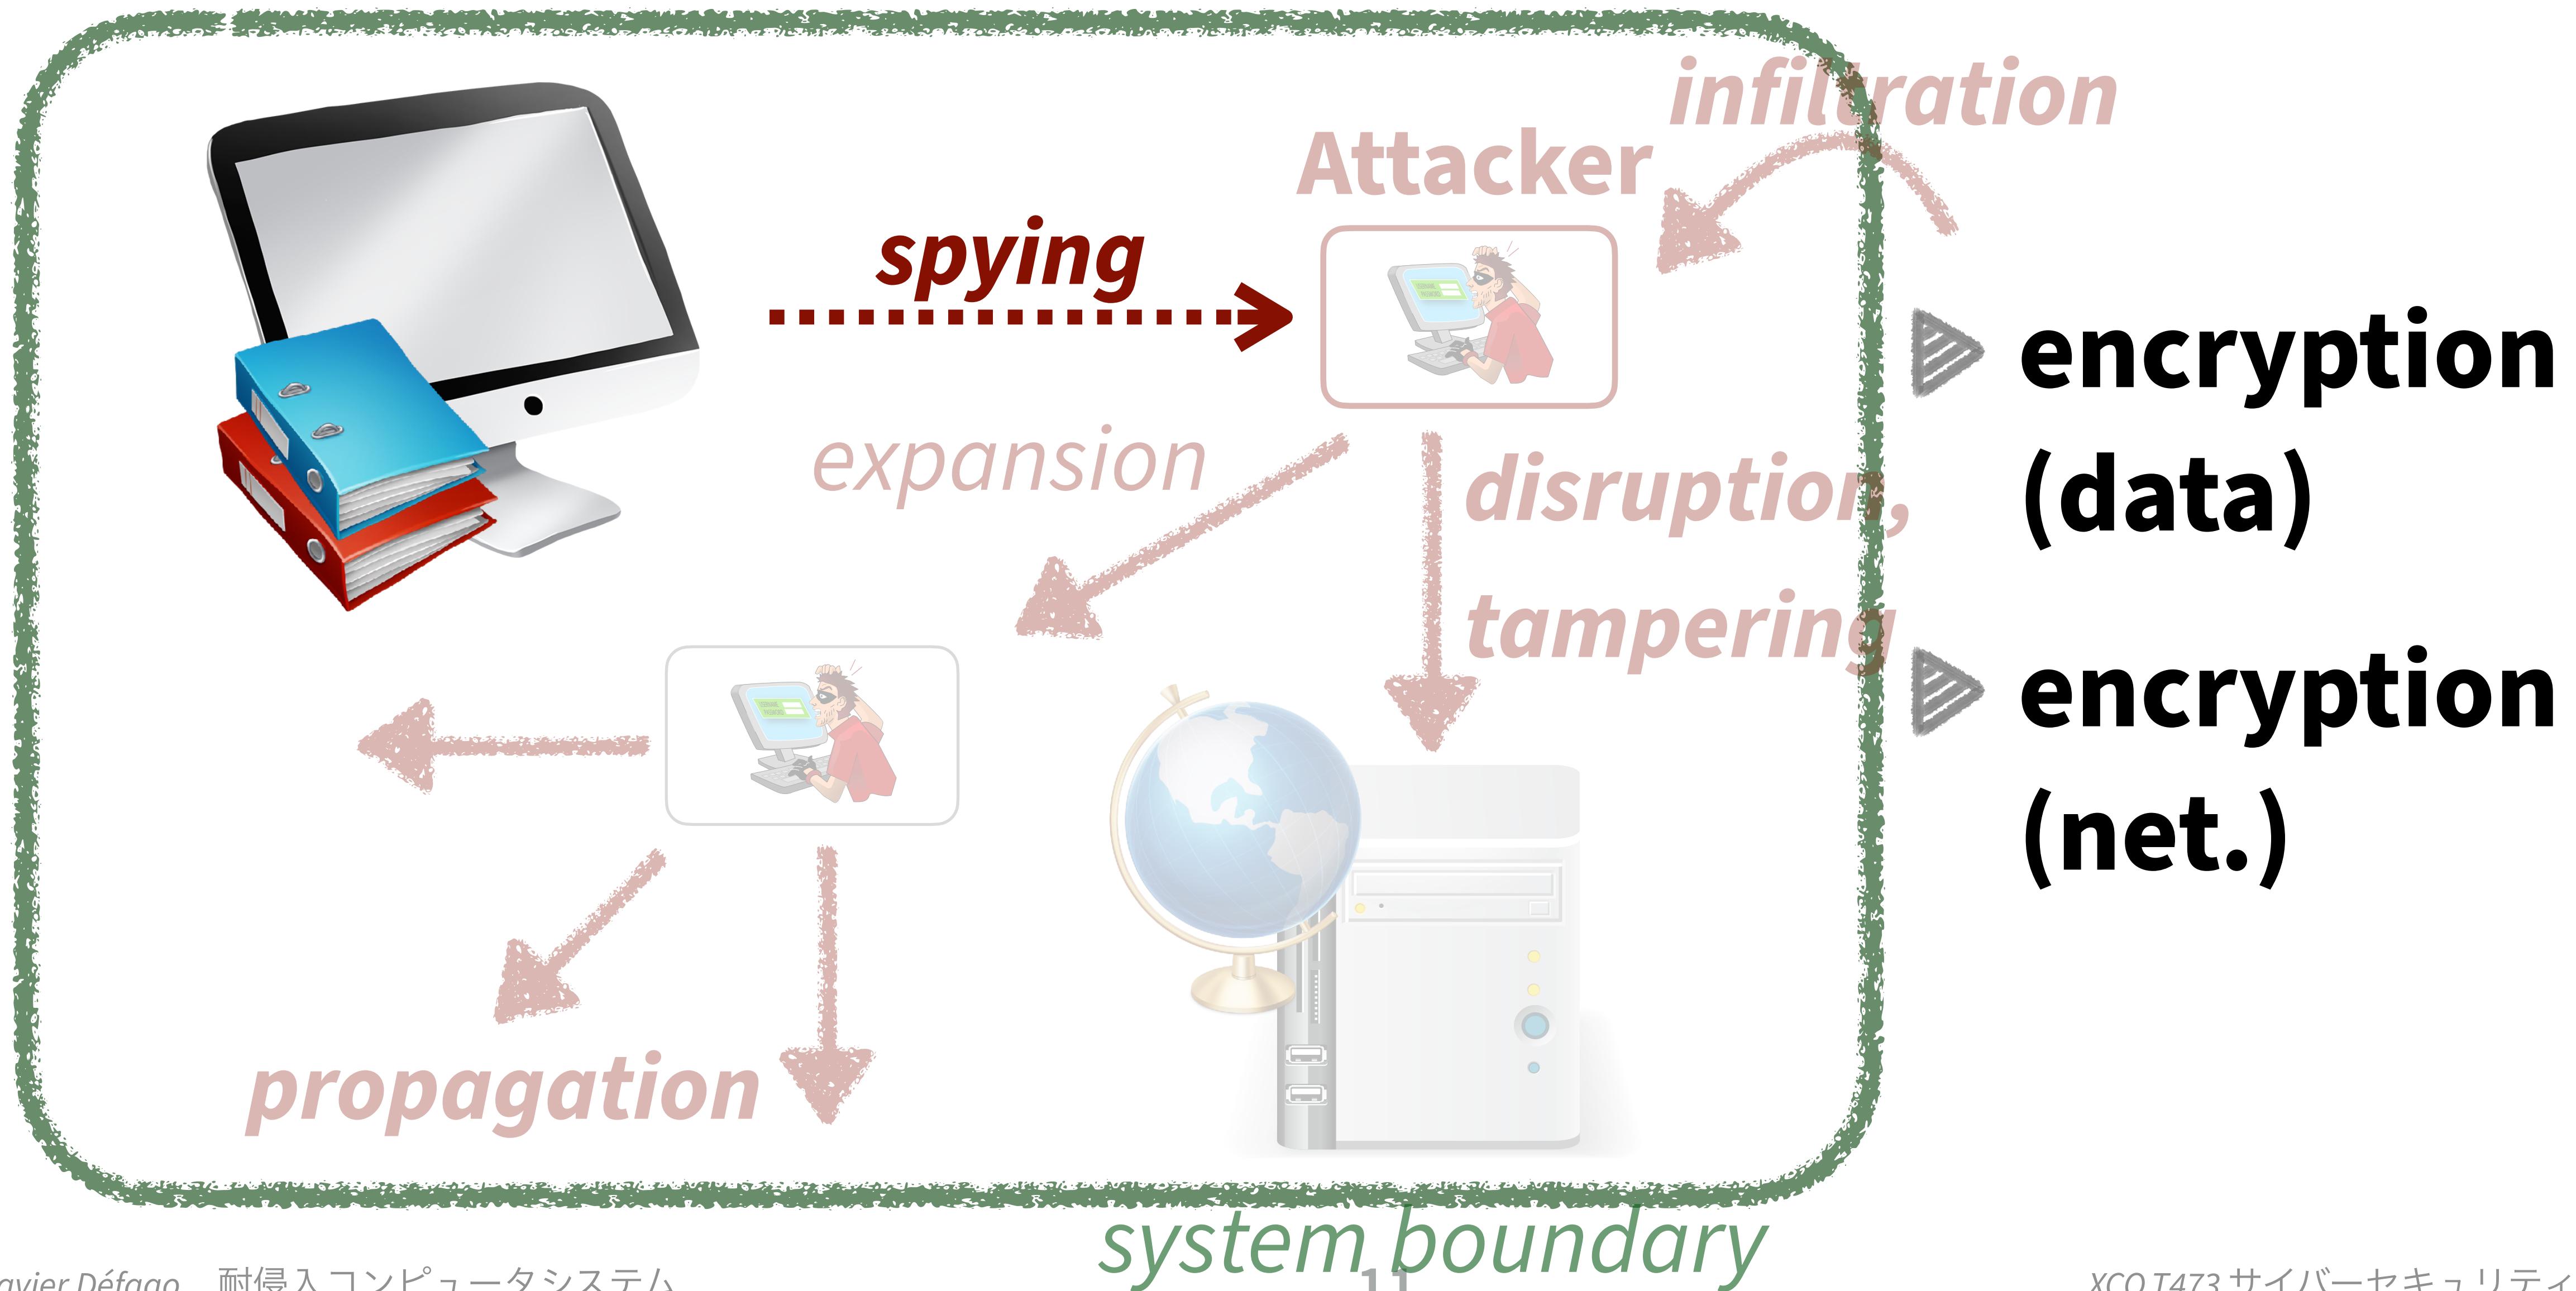
Detection time is very long!



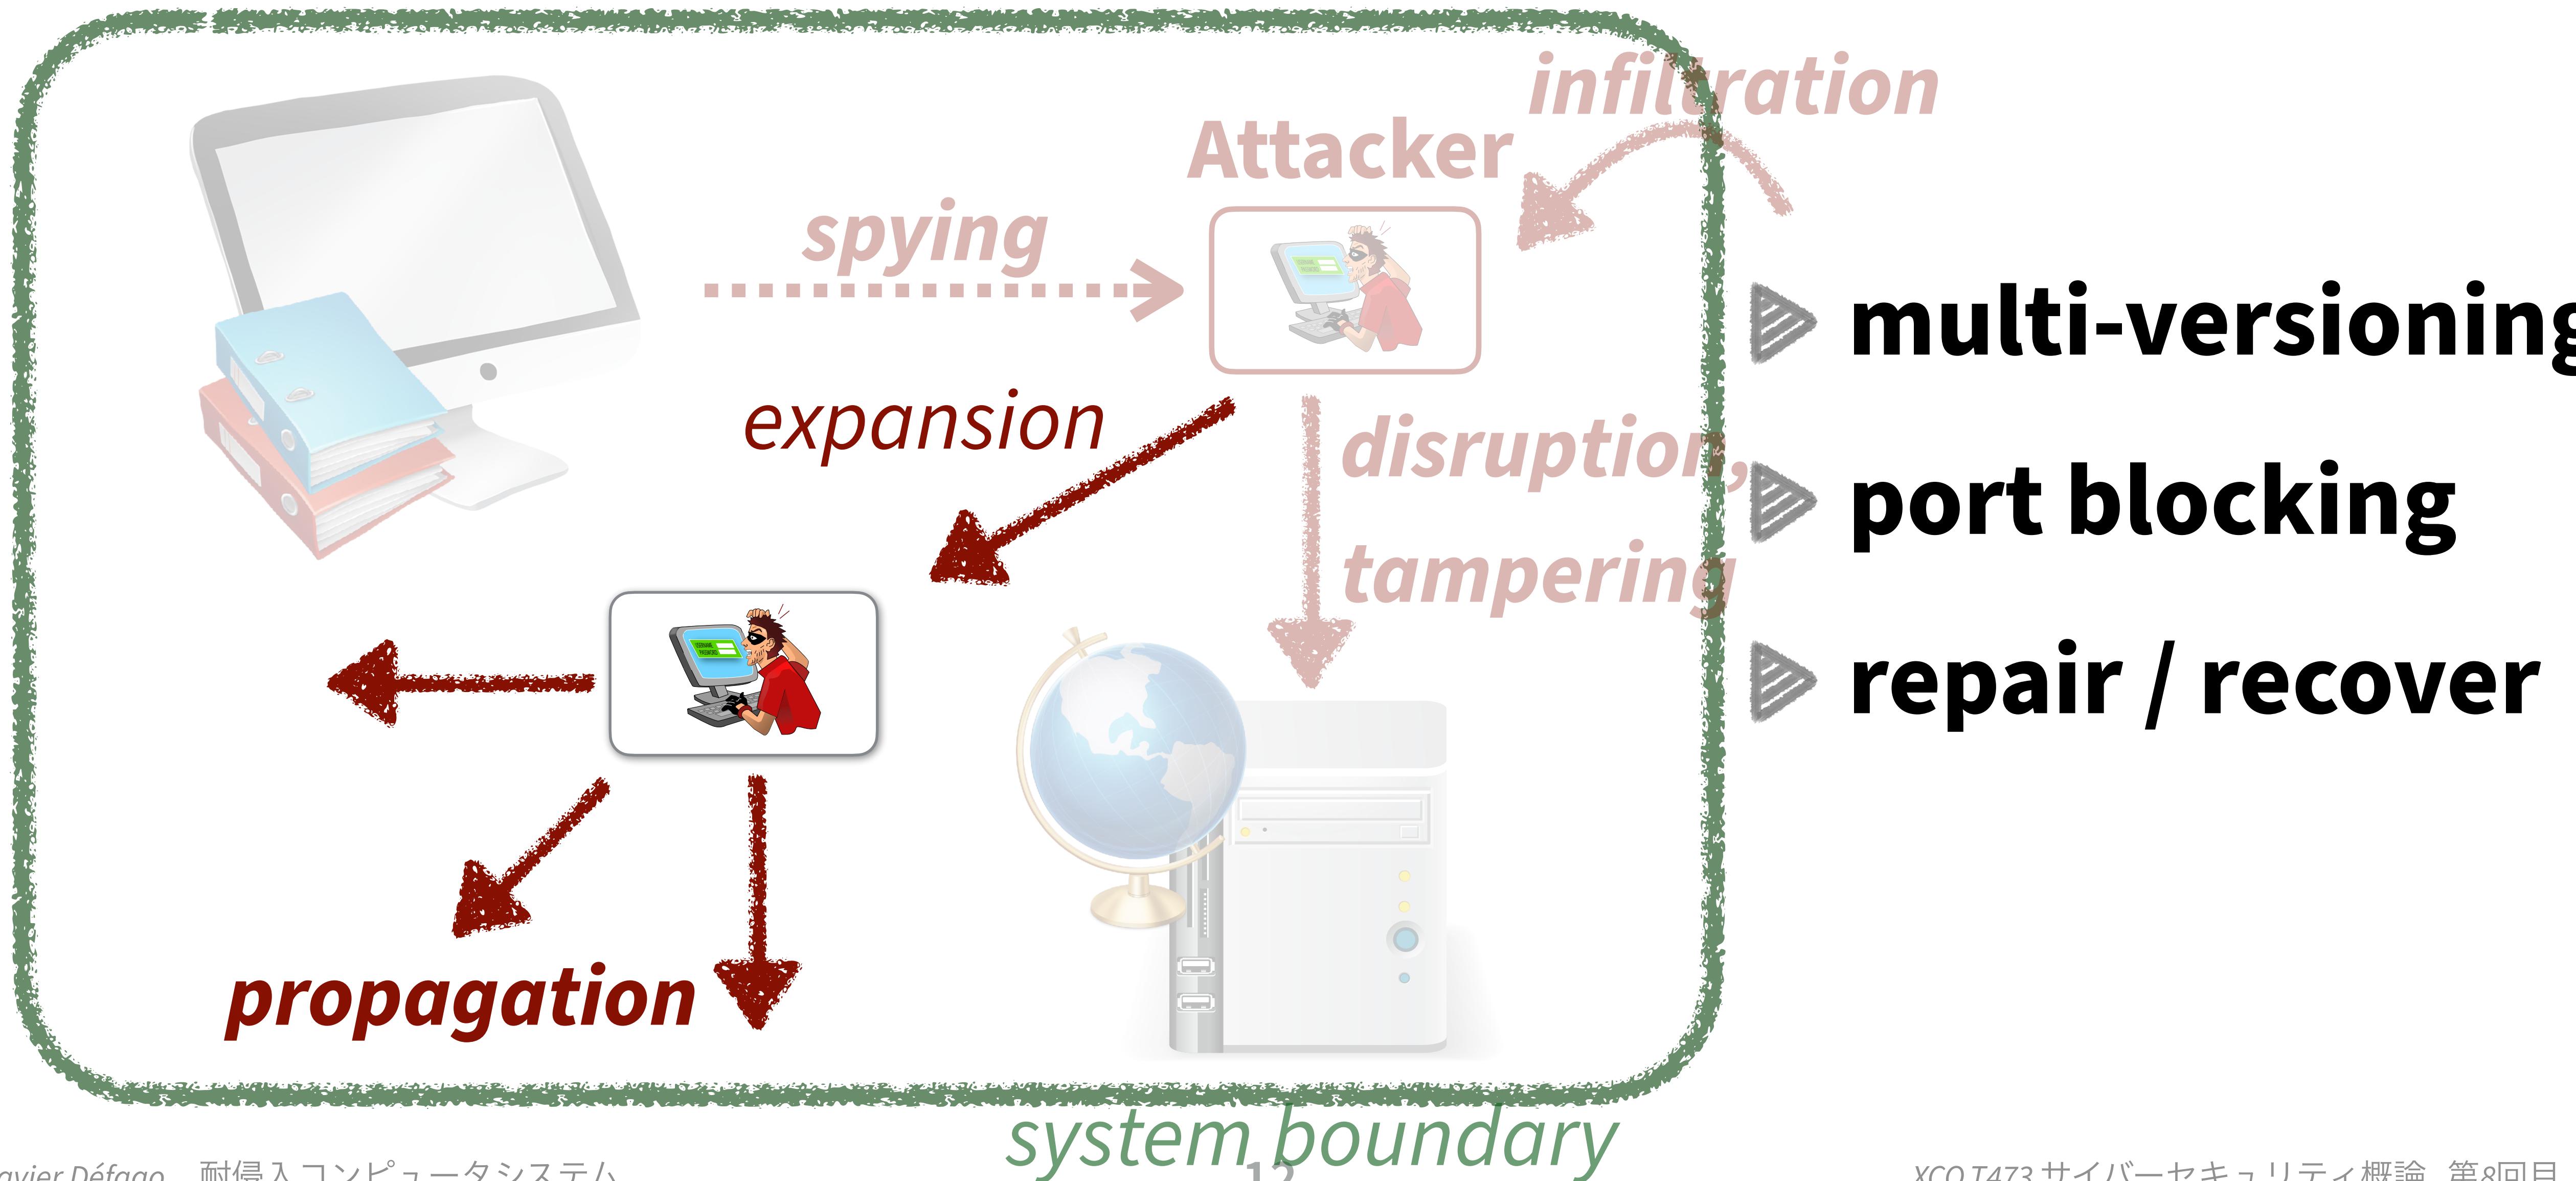
Typical Countermeasures



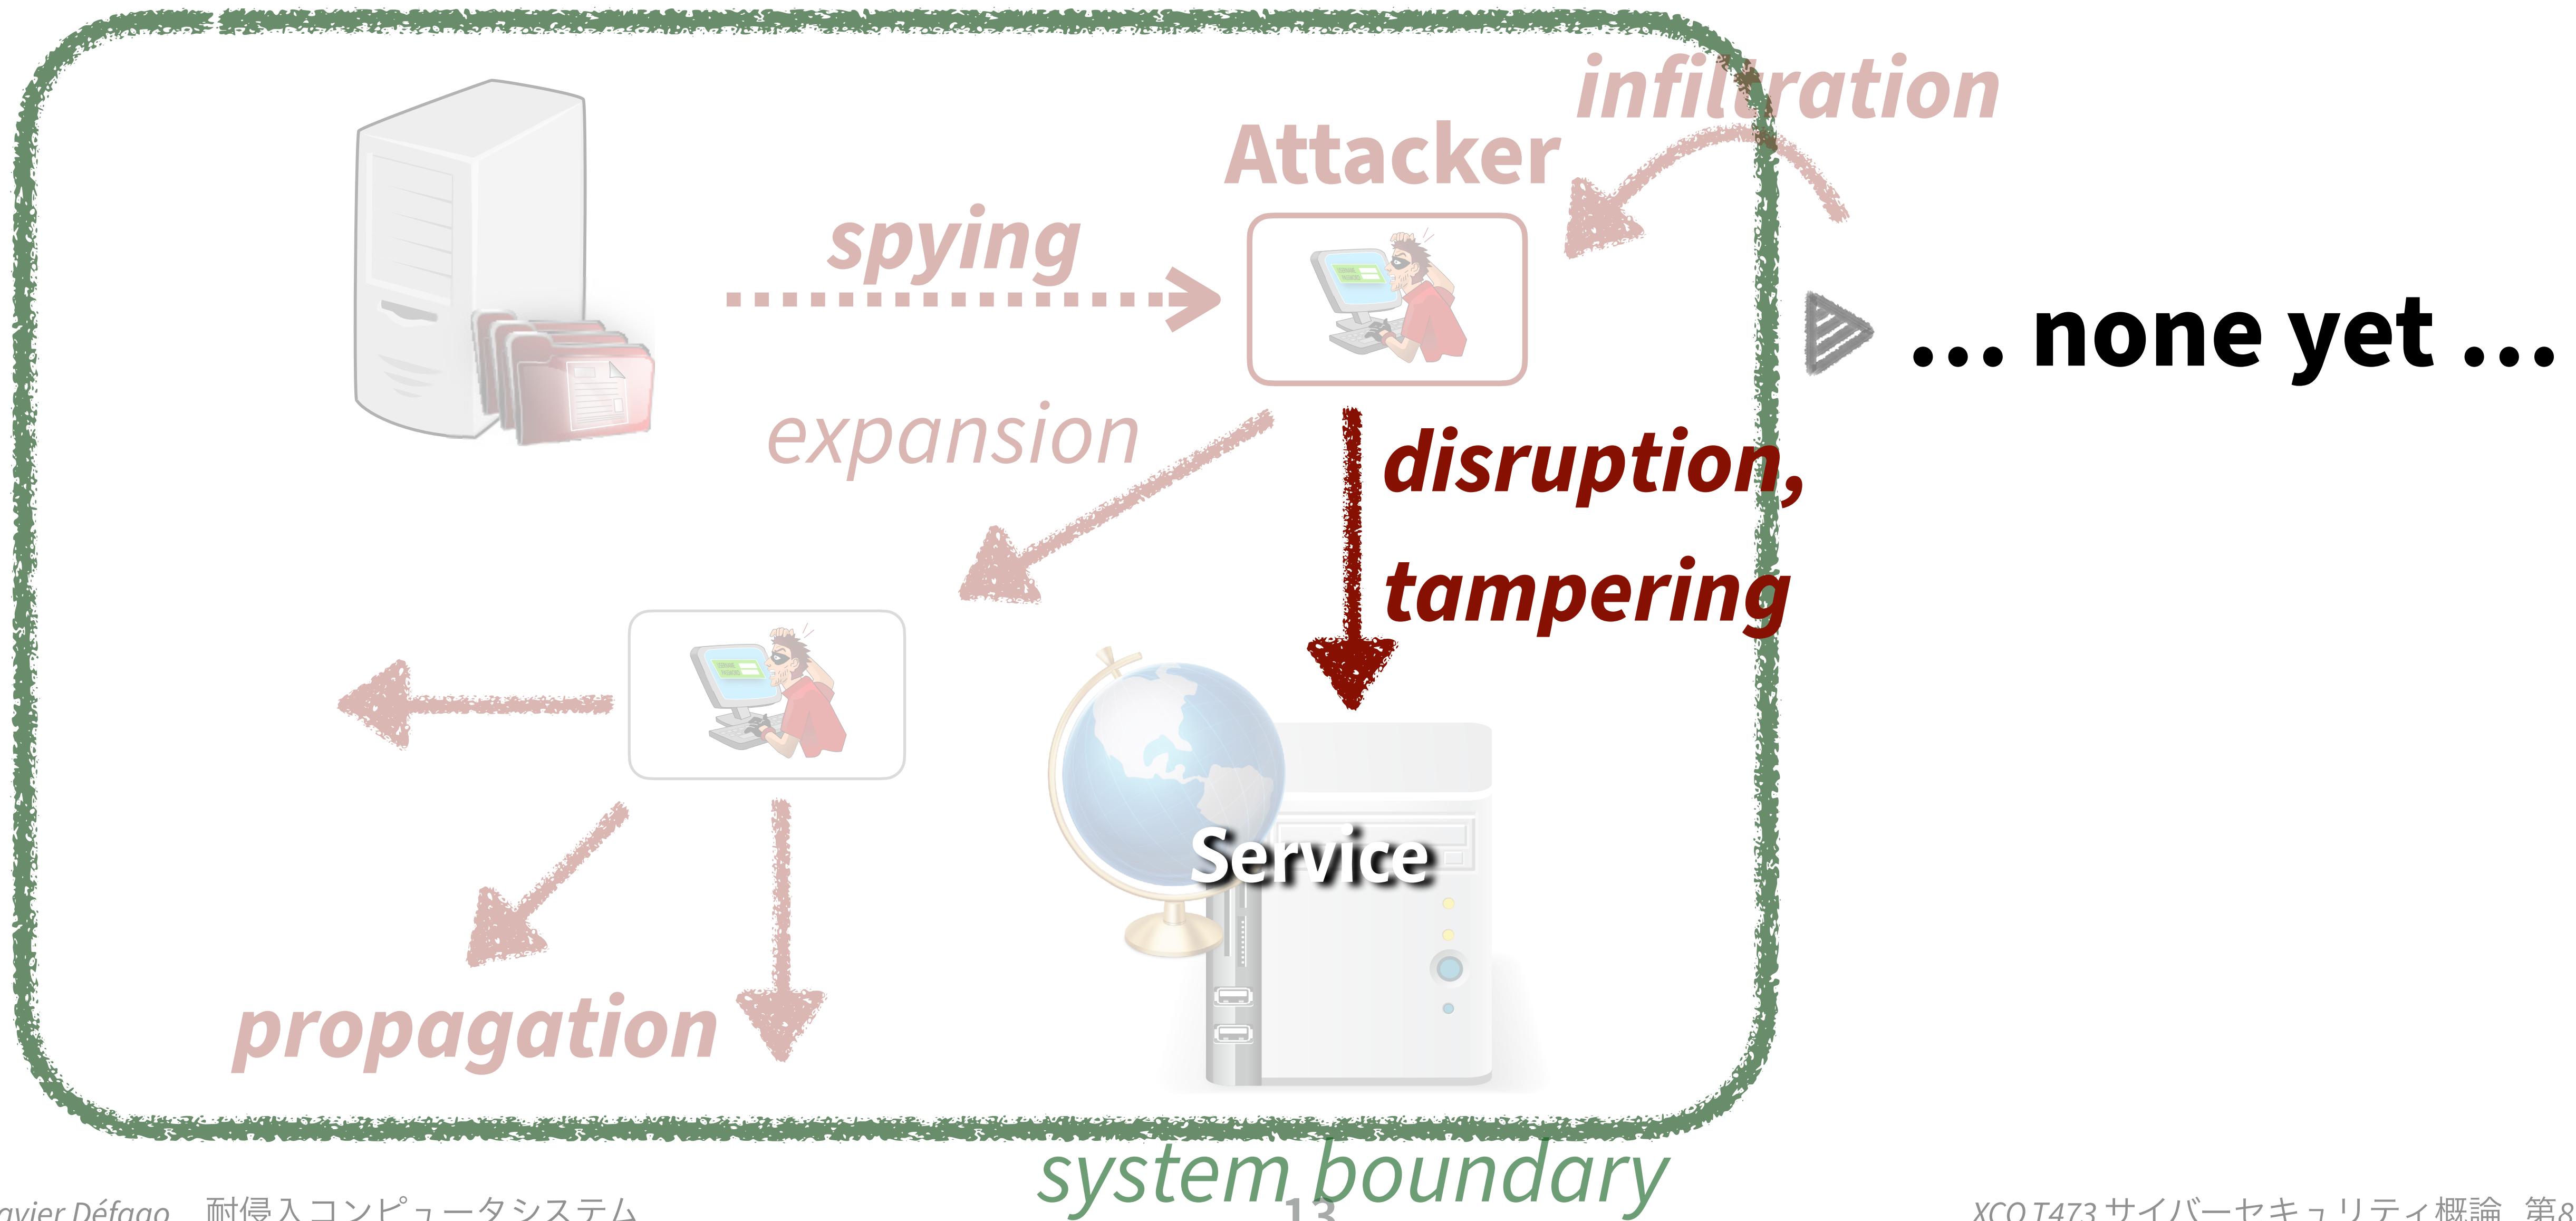
Typical Countermeasures



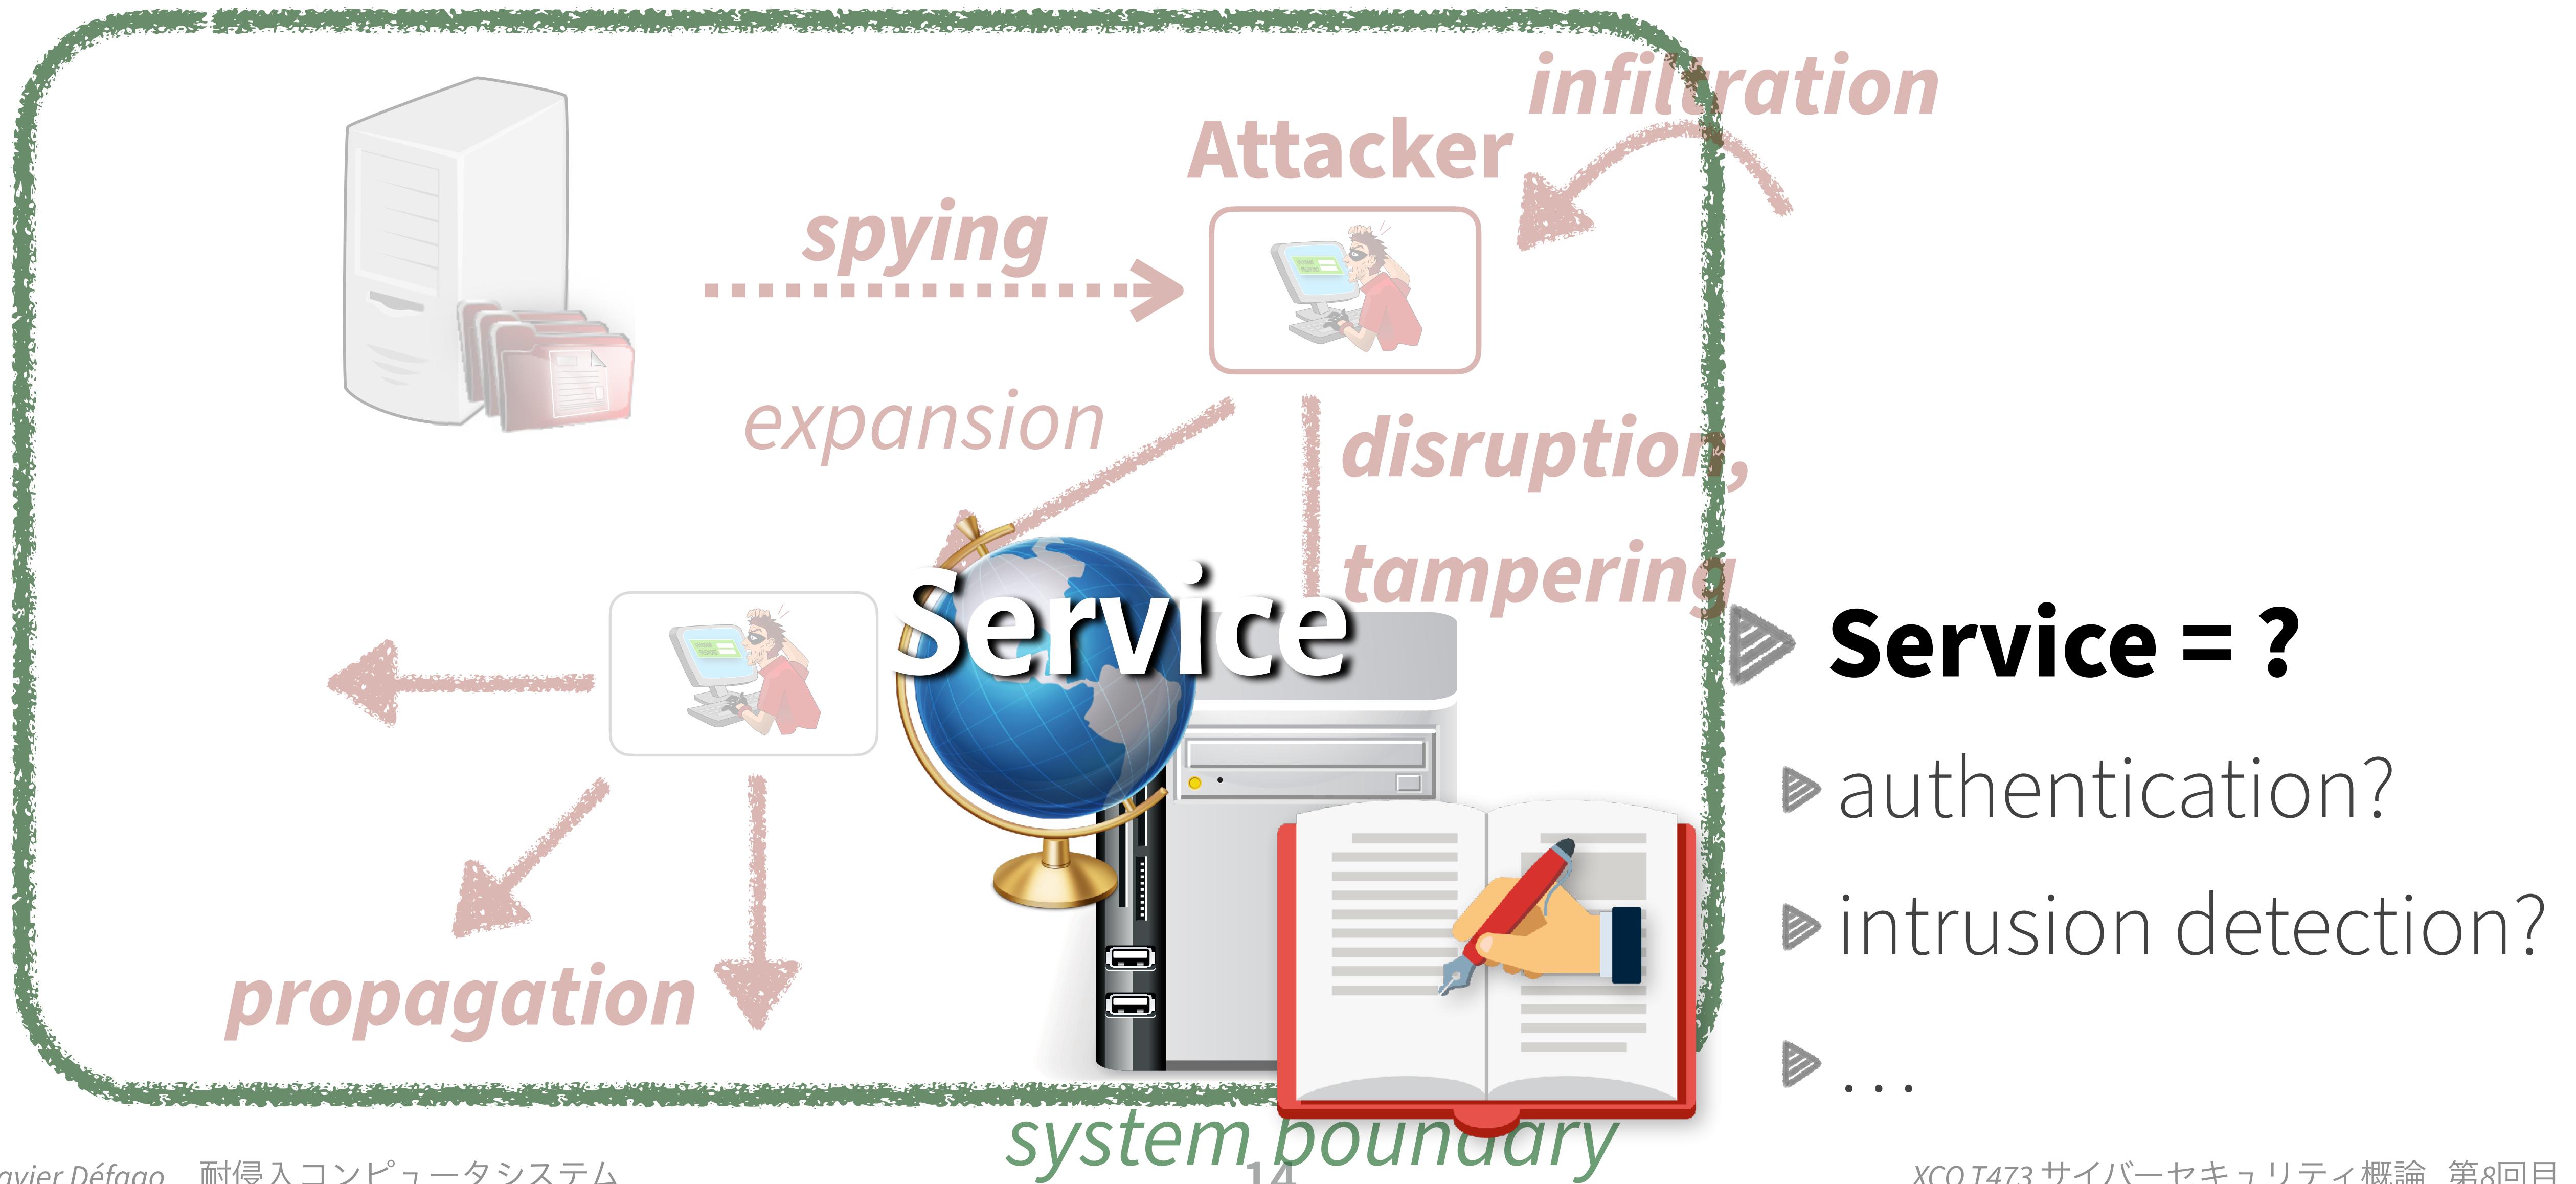
Typical Countermeasures



Typical Countermeasures



Typical Countermeasures



**Keep
System Integrity**

Aspects of Security

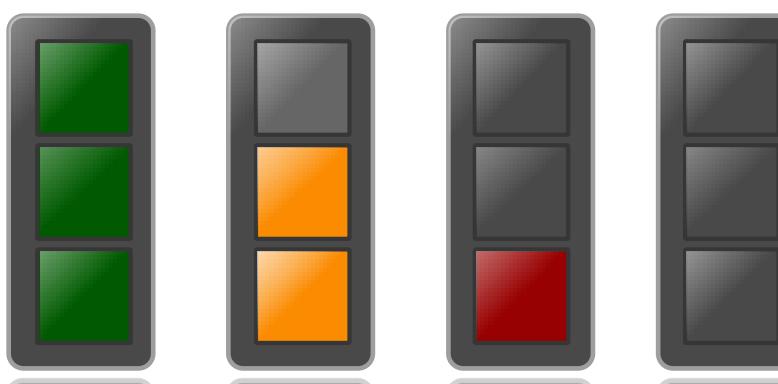
▶ Confidentiality

- ▶ No information leak



▶ Availability

- ▶ Provide operation (no DoS)



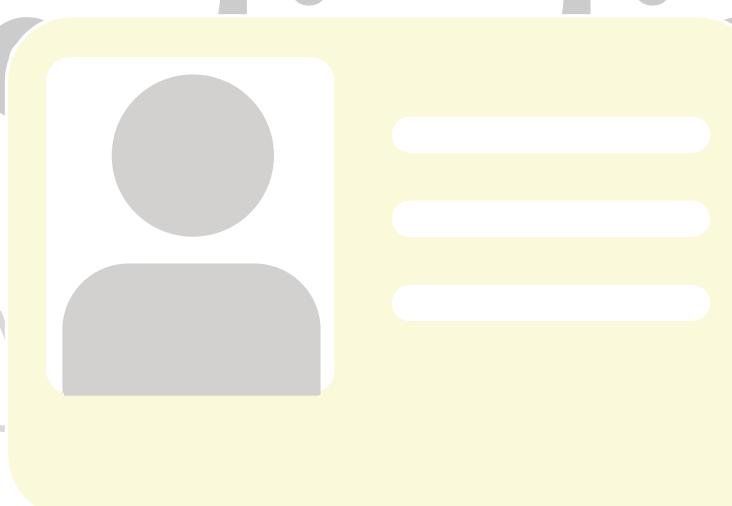
▶ Integrity

- ▶ Present correct state



▶ Authentication

- ▶ “Are you?”



Data / System Integrity

► Requirements

- ▶ No tampering of data / system
- ▶ Distributed system: preserve **consistency**

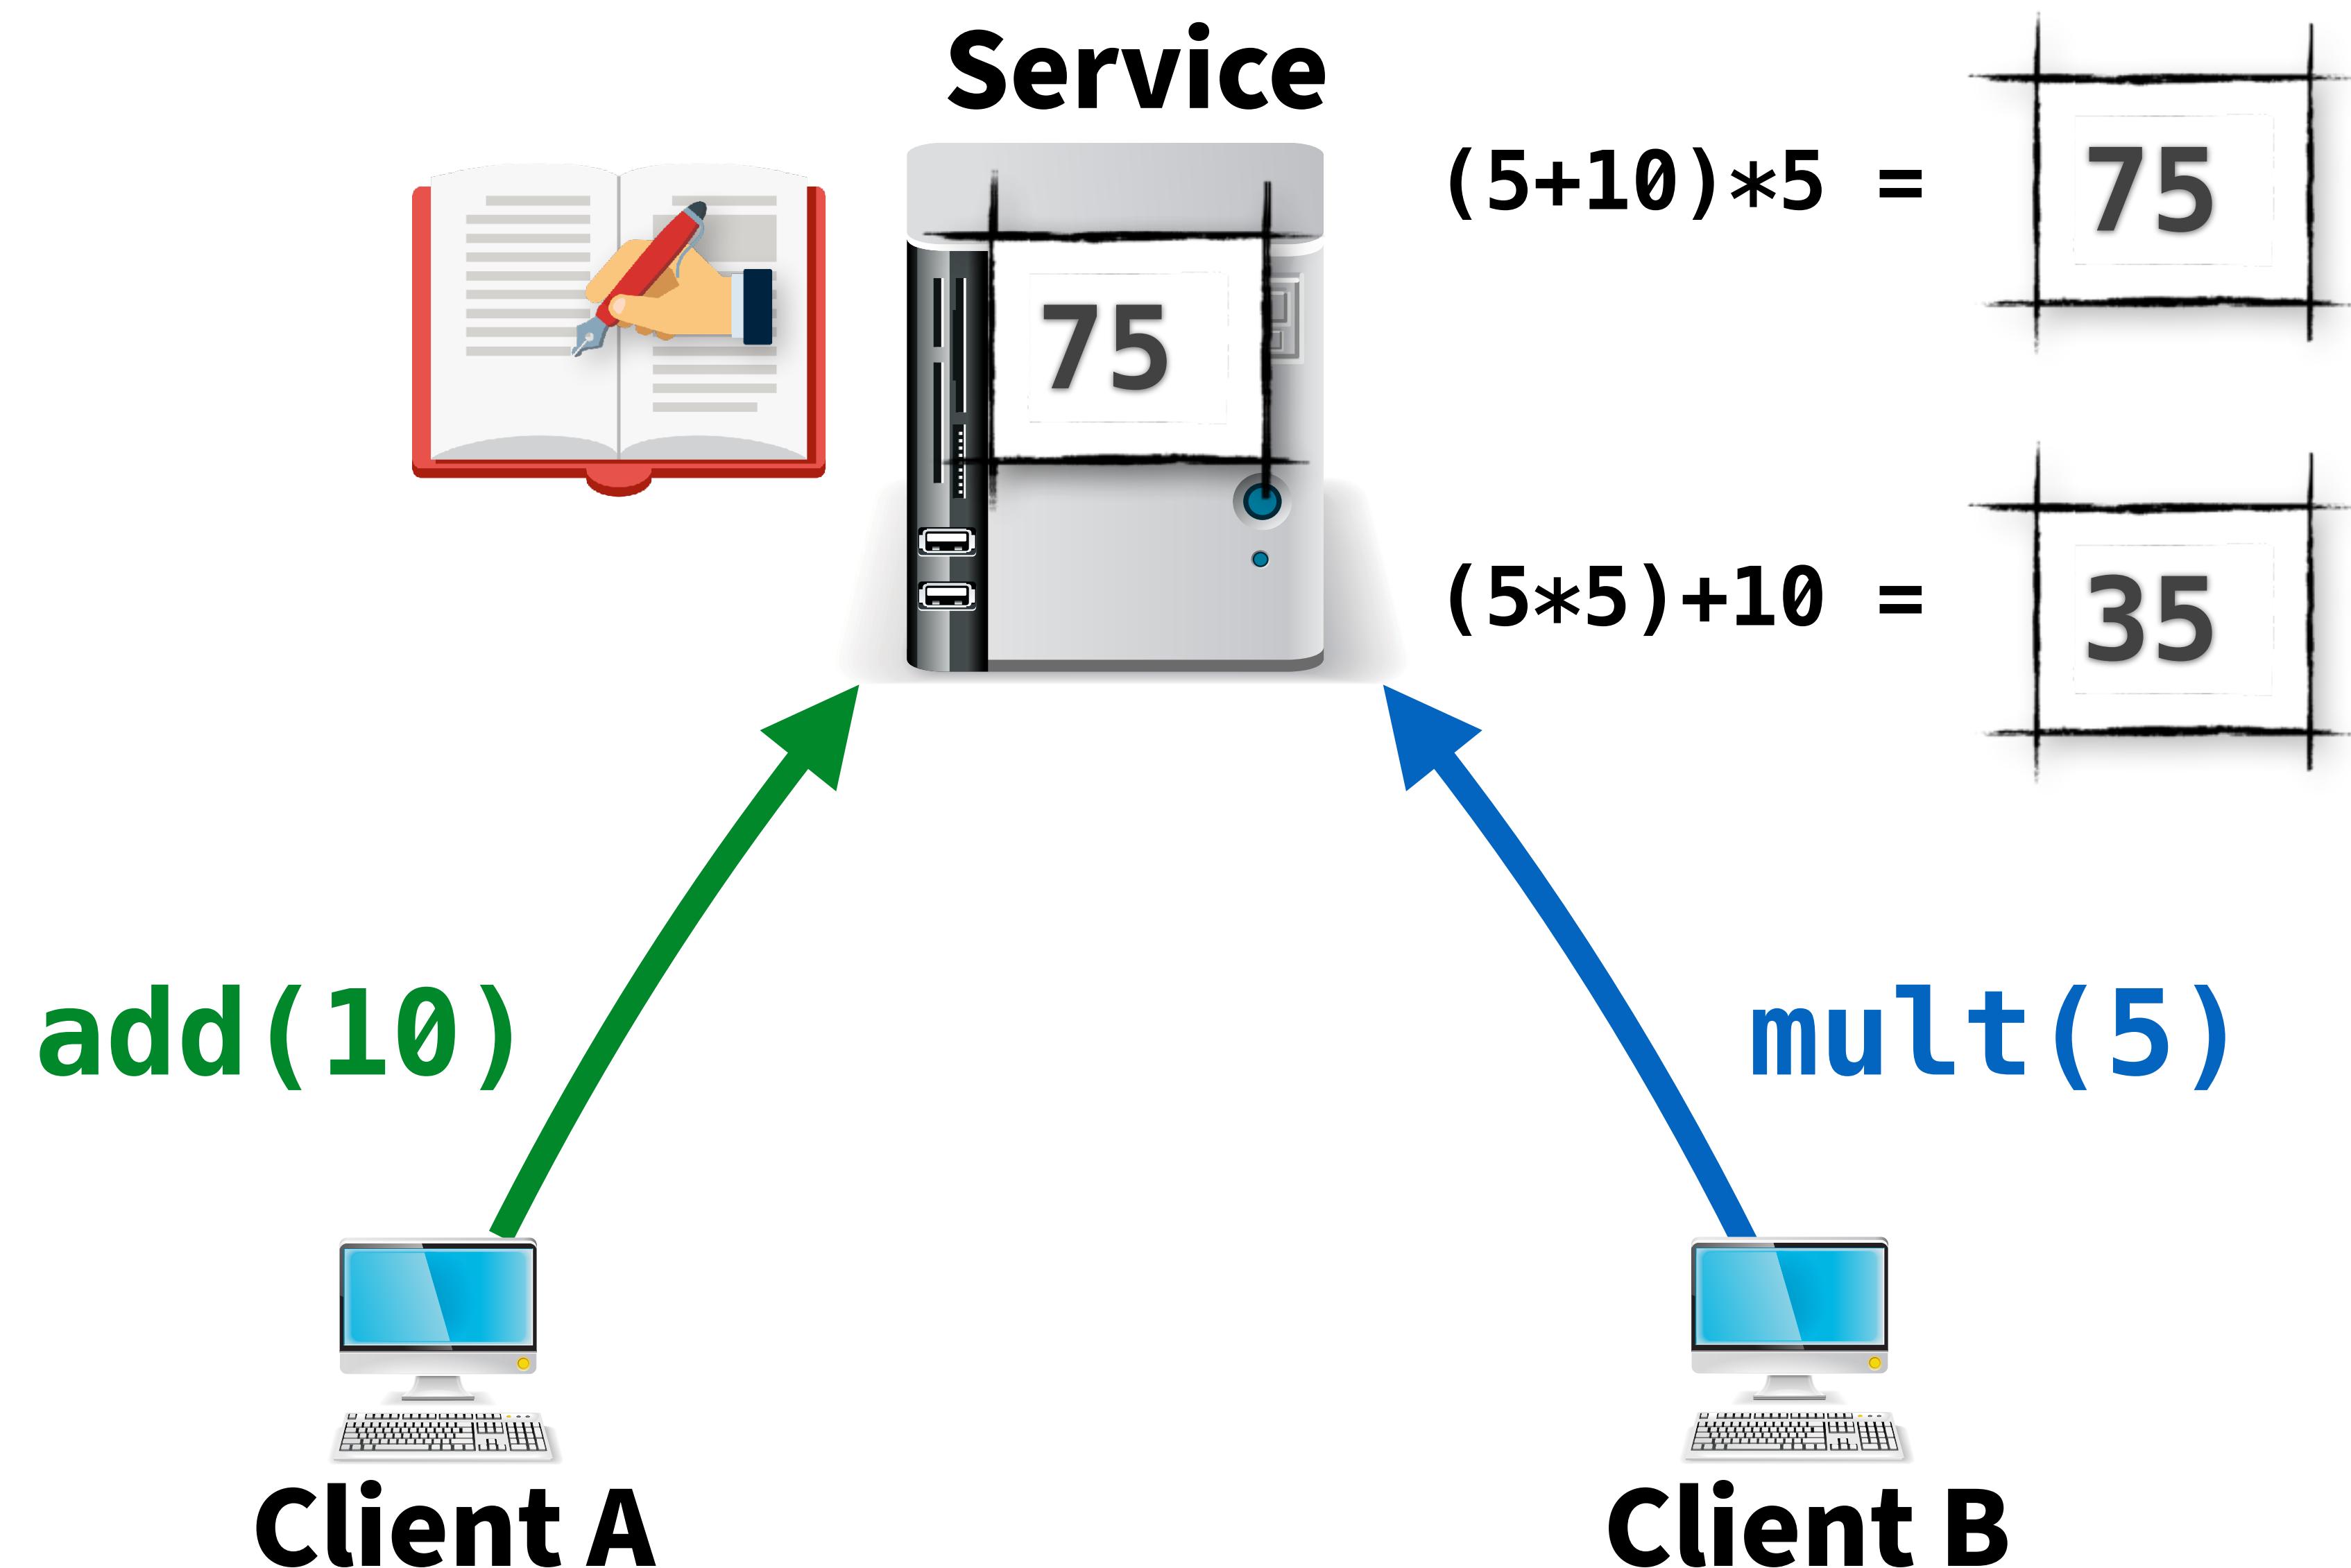
► “Outside” Threat

- ▶ Attacker has **no secret** initially
- ▶ Attacker uses **public** service interface in unexpected ways

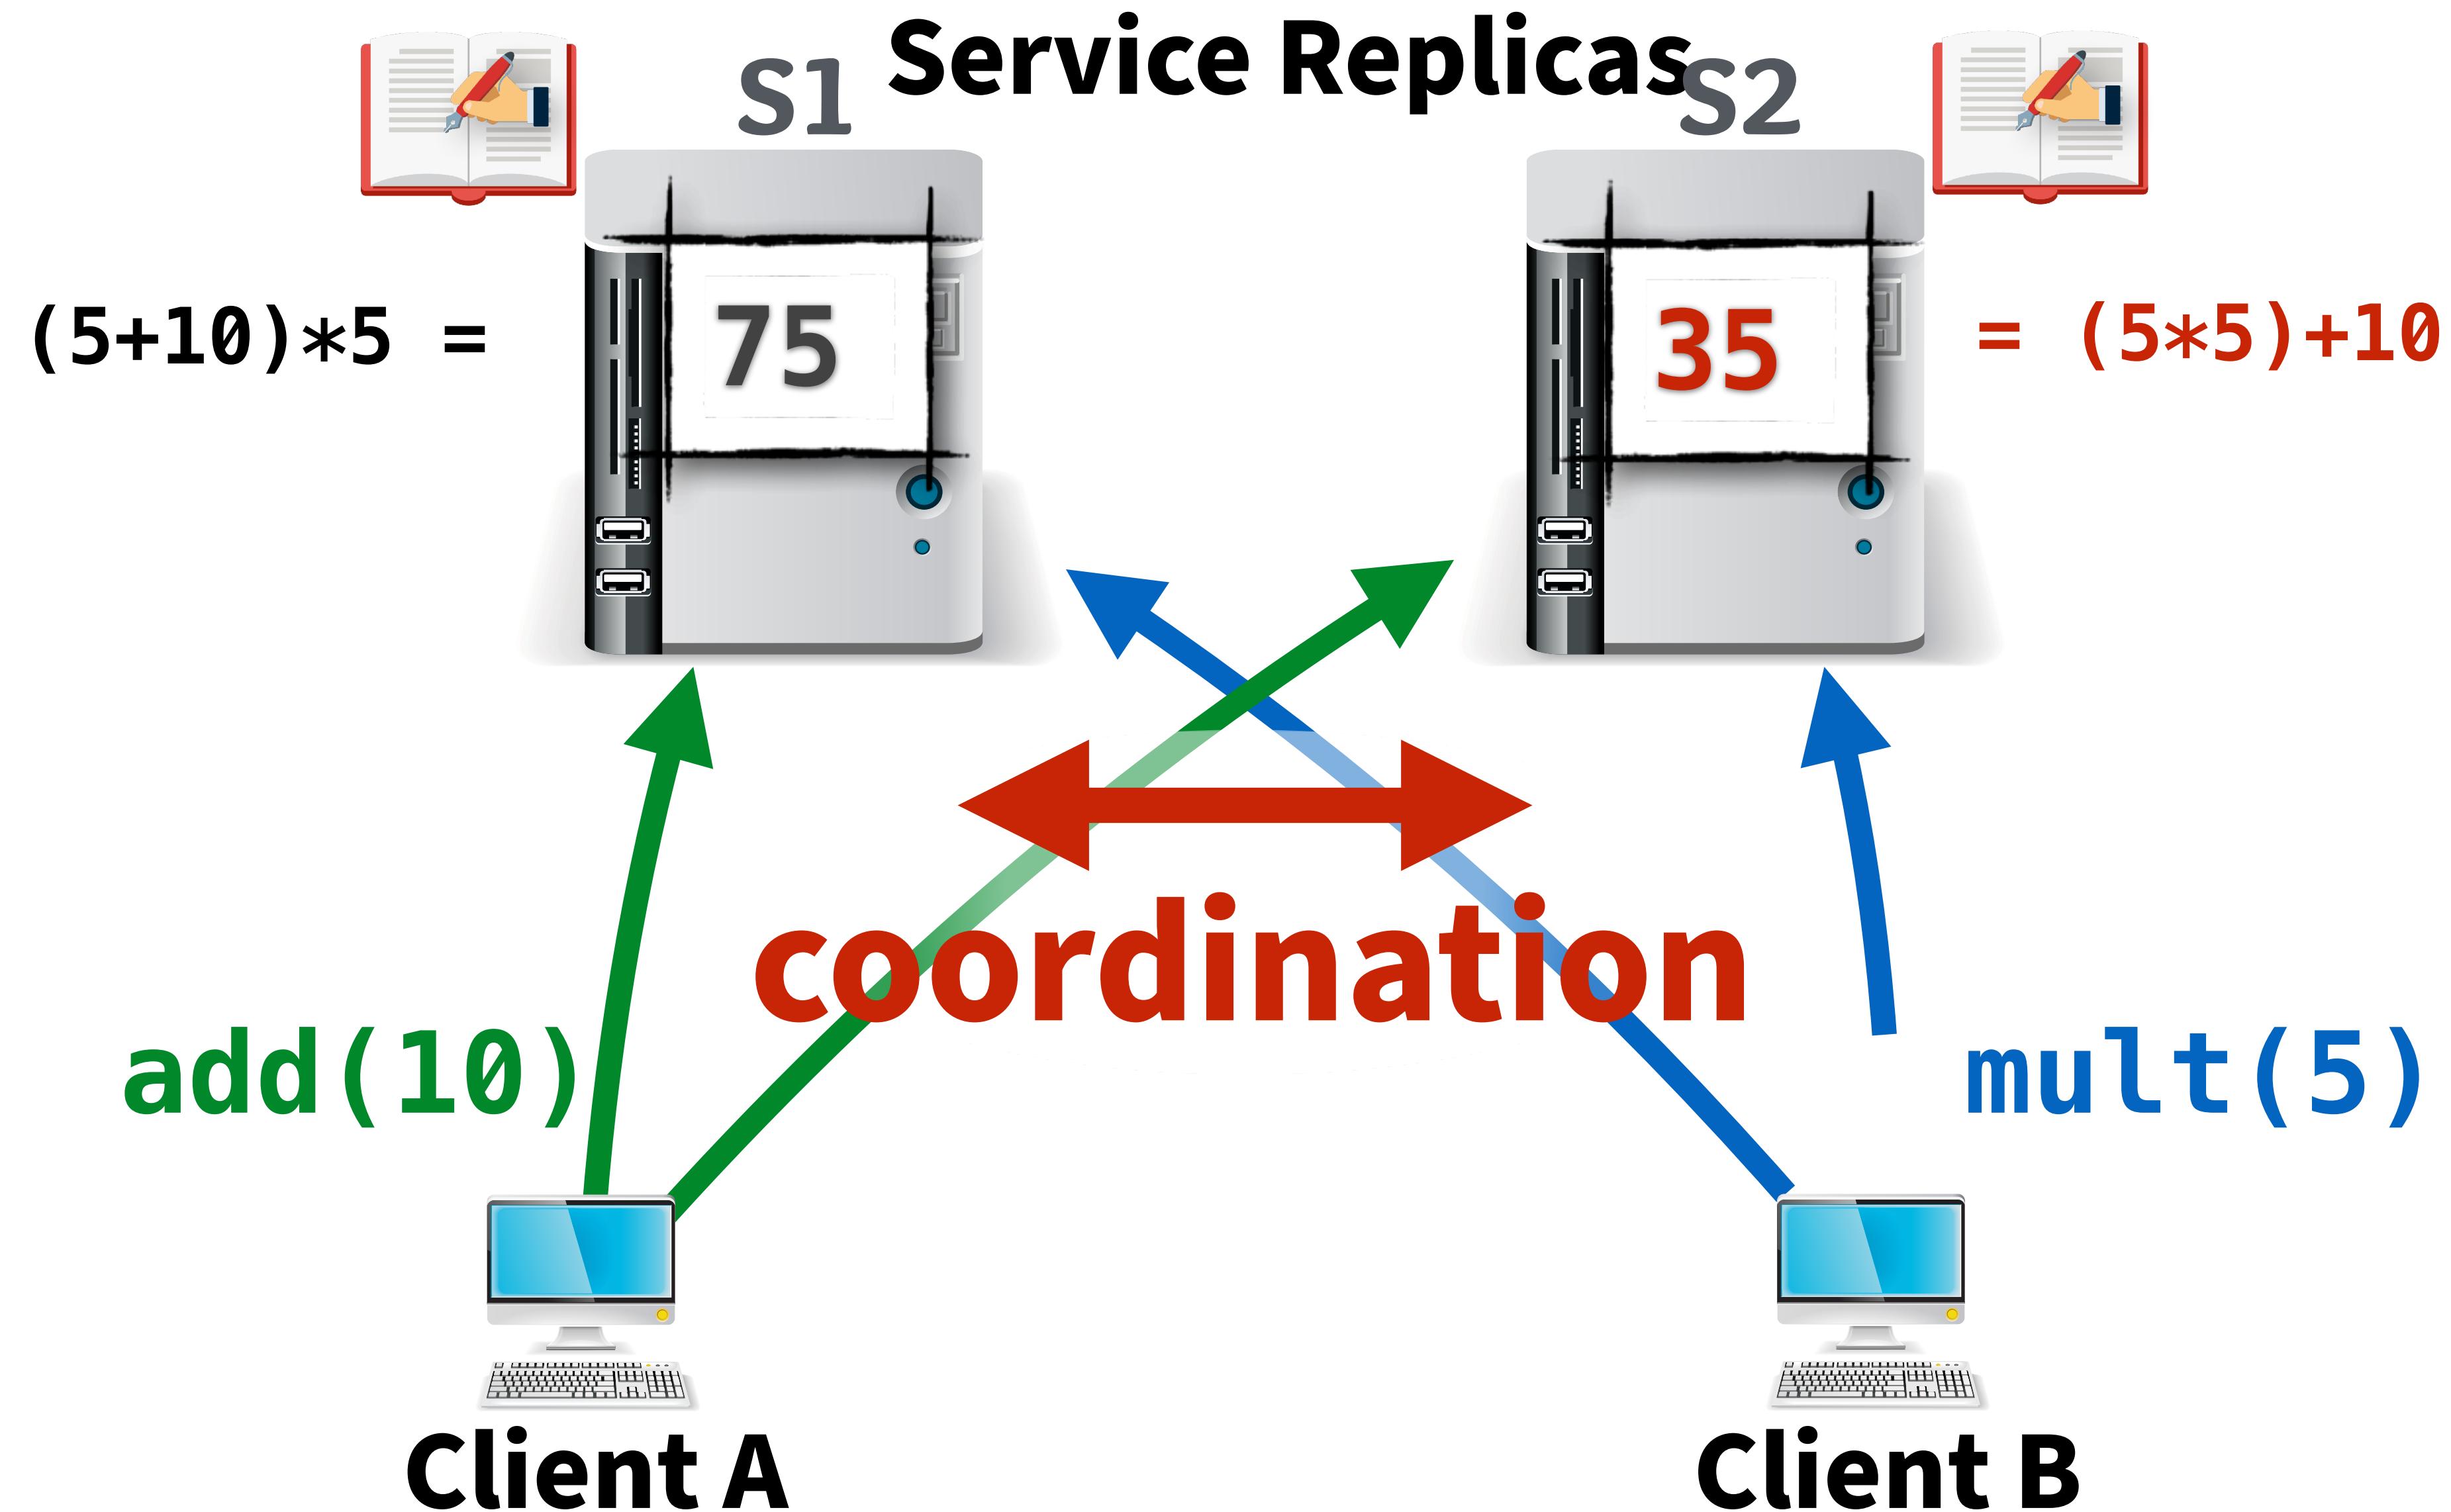
► “Inside” Threat

- ▶ Attacker **passes as legitimate** node / server (eludes detection)
- ▶ Attacker has access to **all information**

The Replicated Register



The Replicated Register

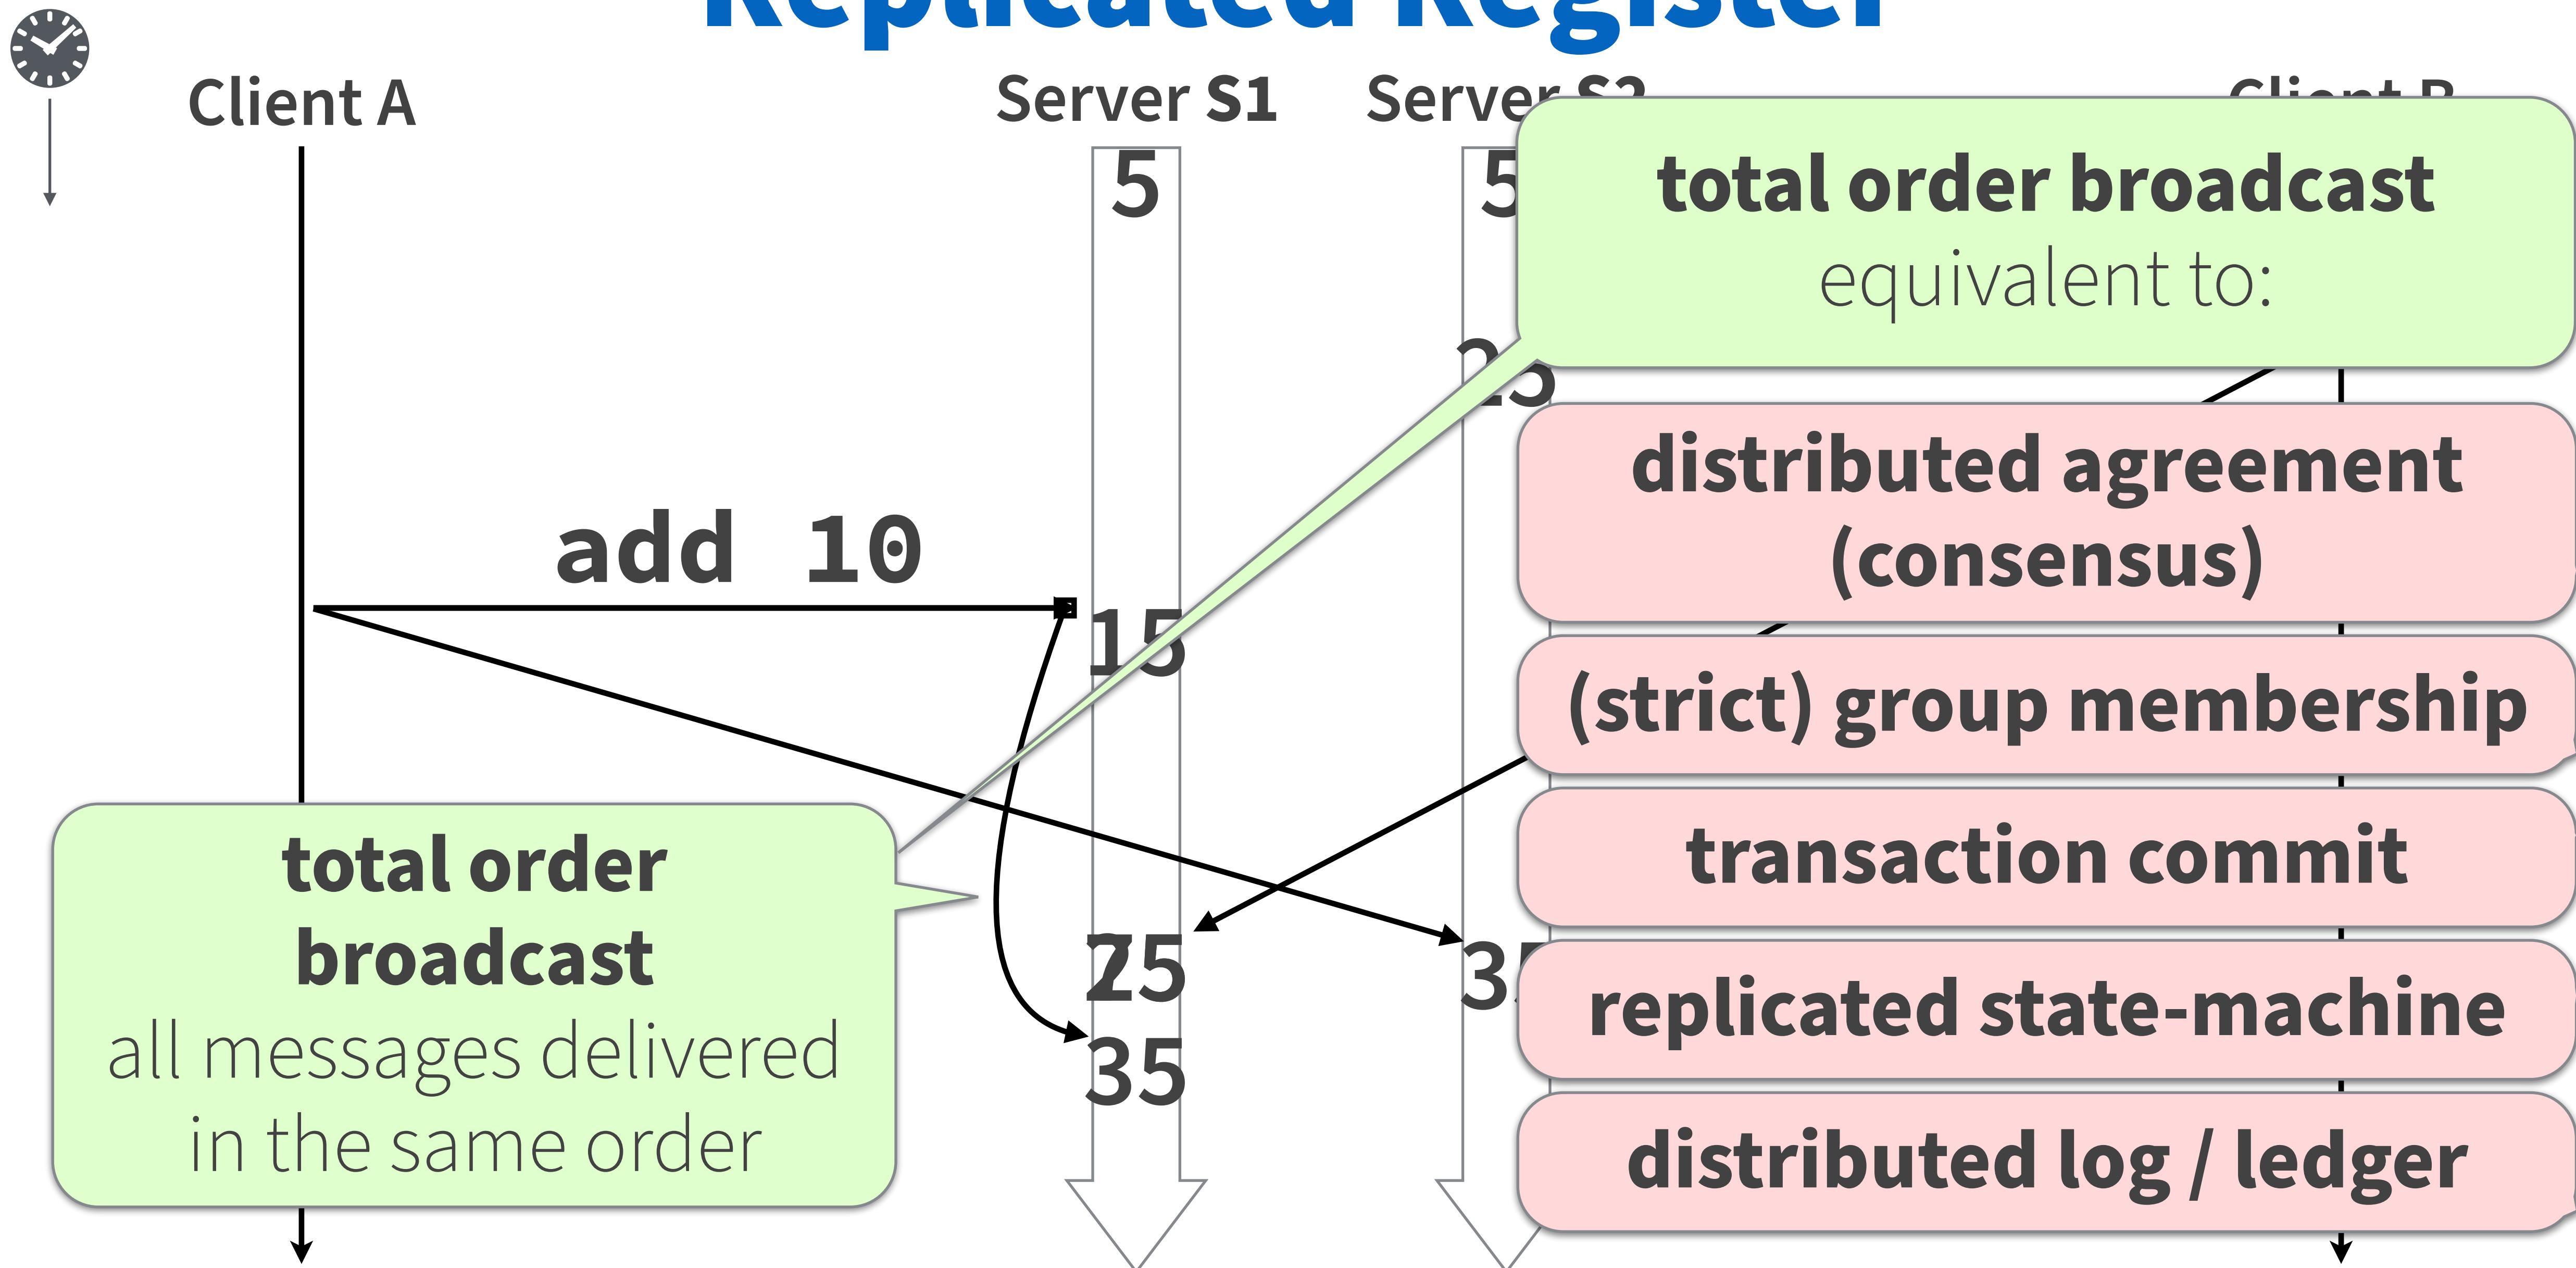


Consistency

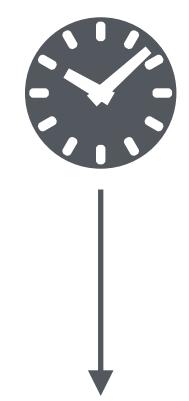
&

Total Order

Replicated Register



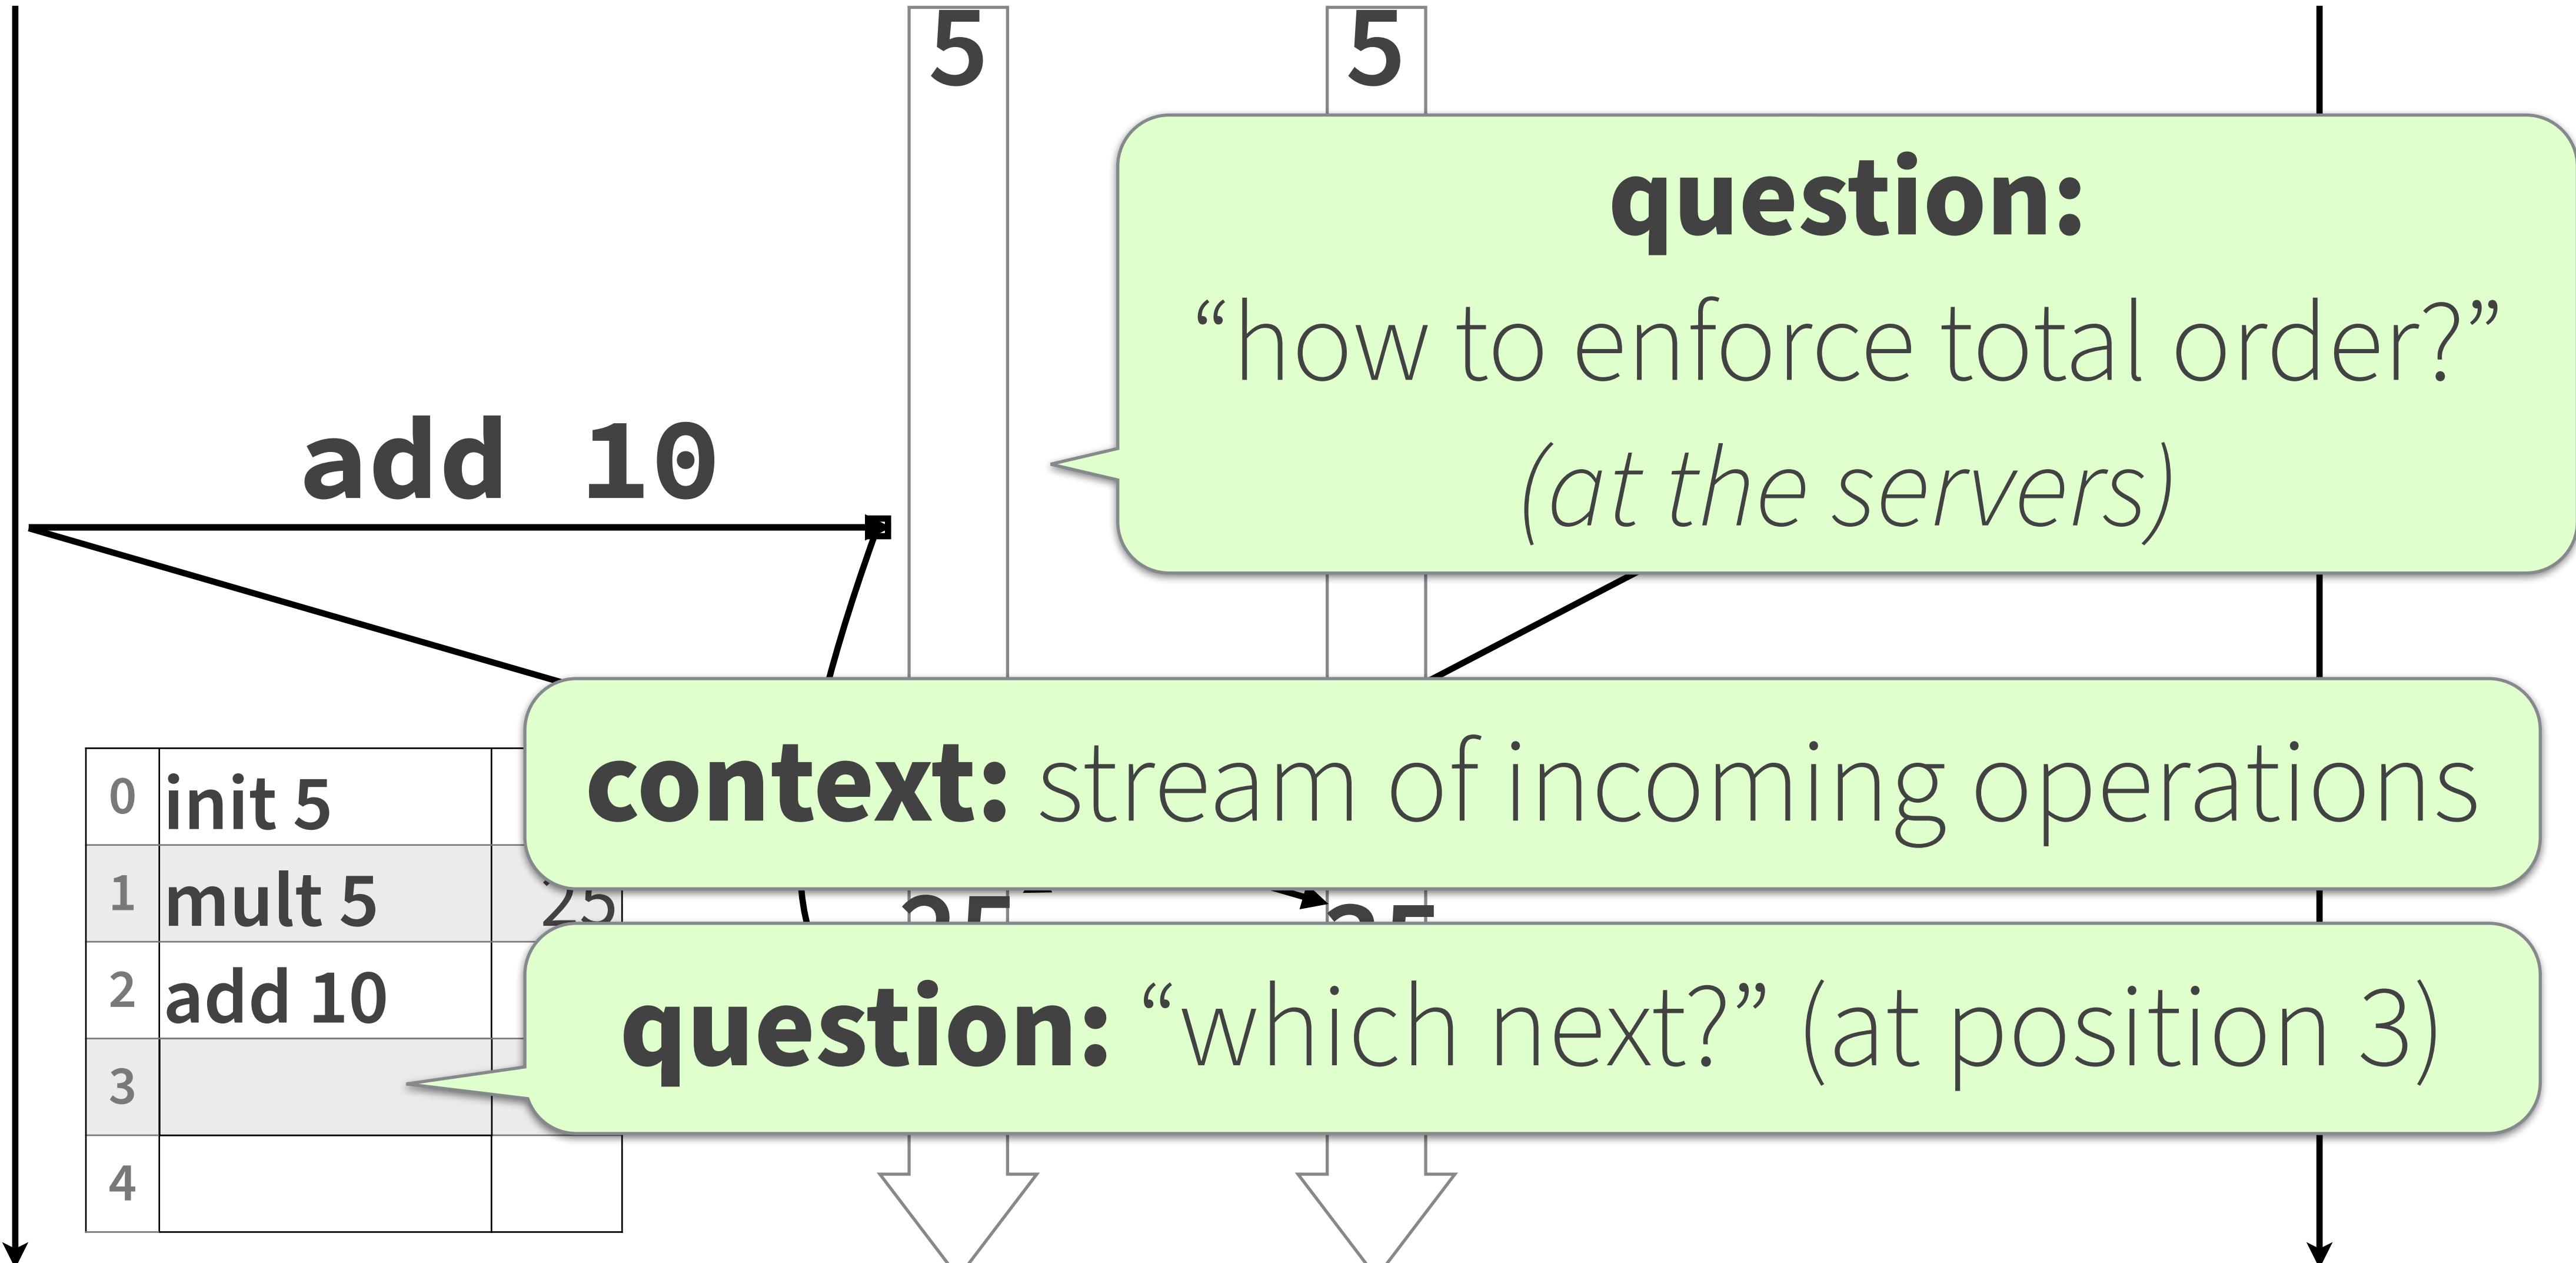
Replicated Register



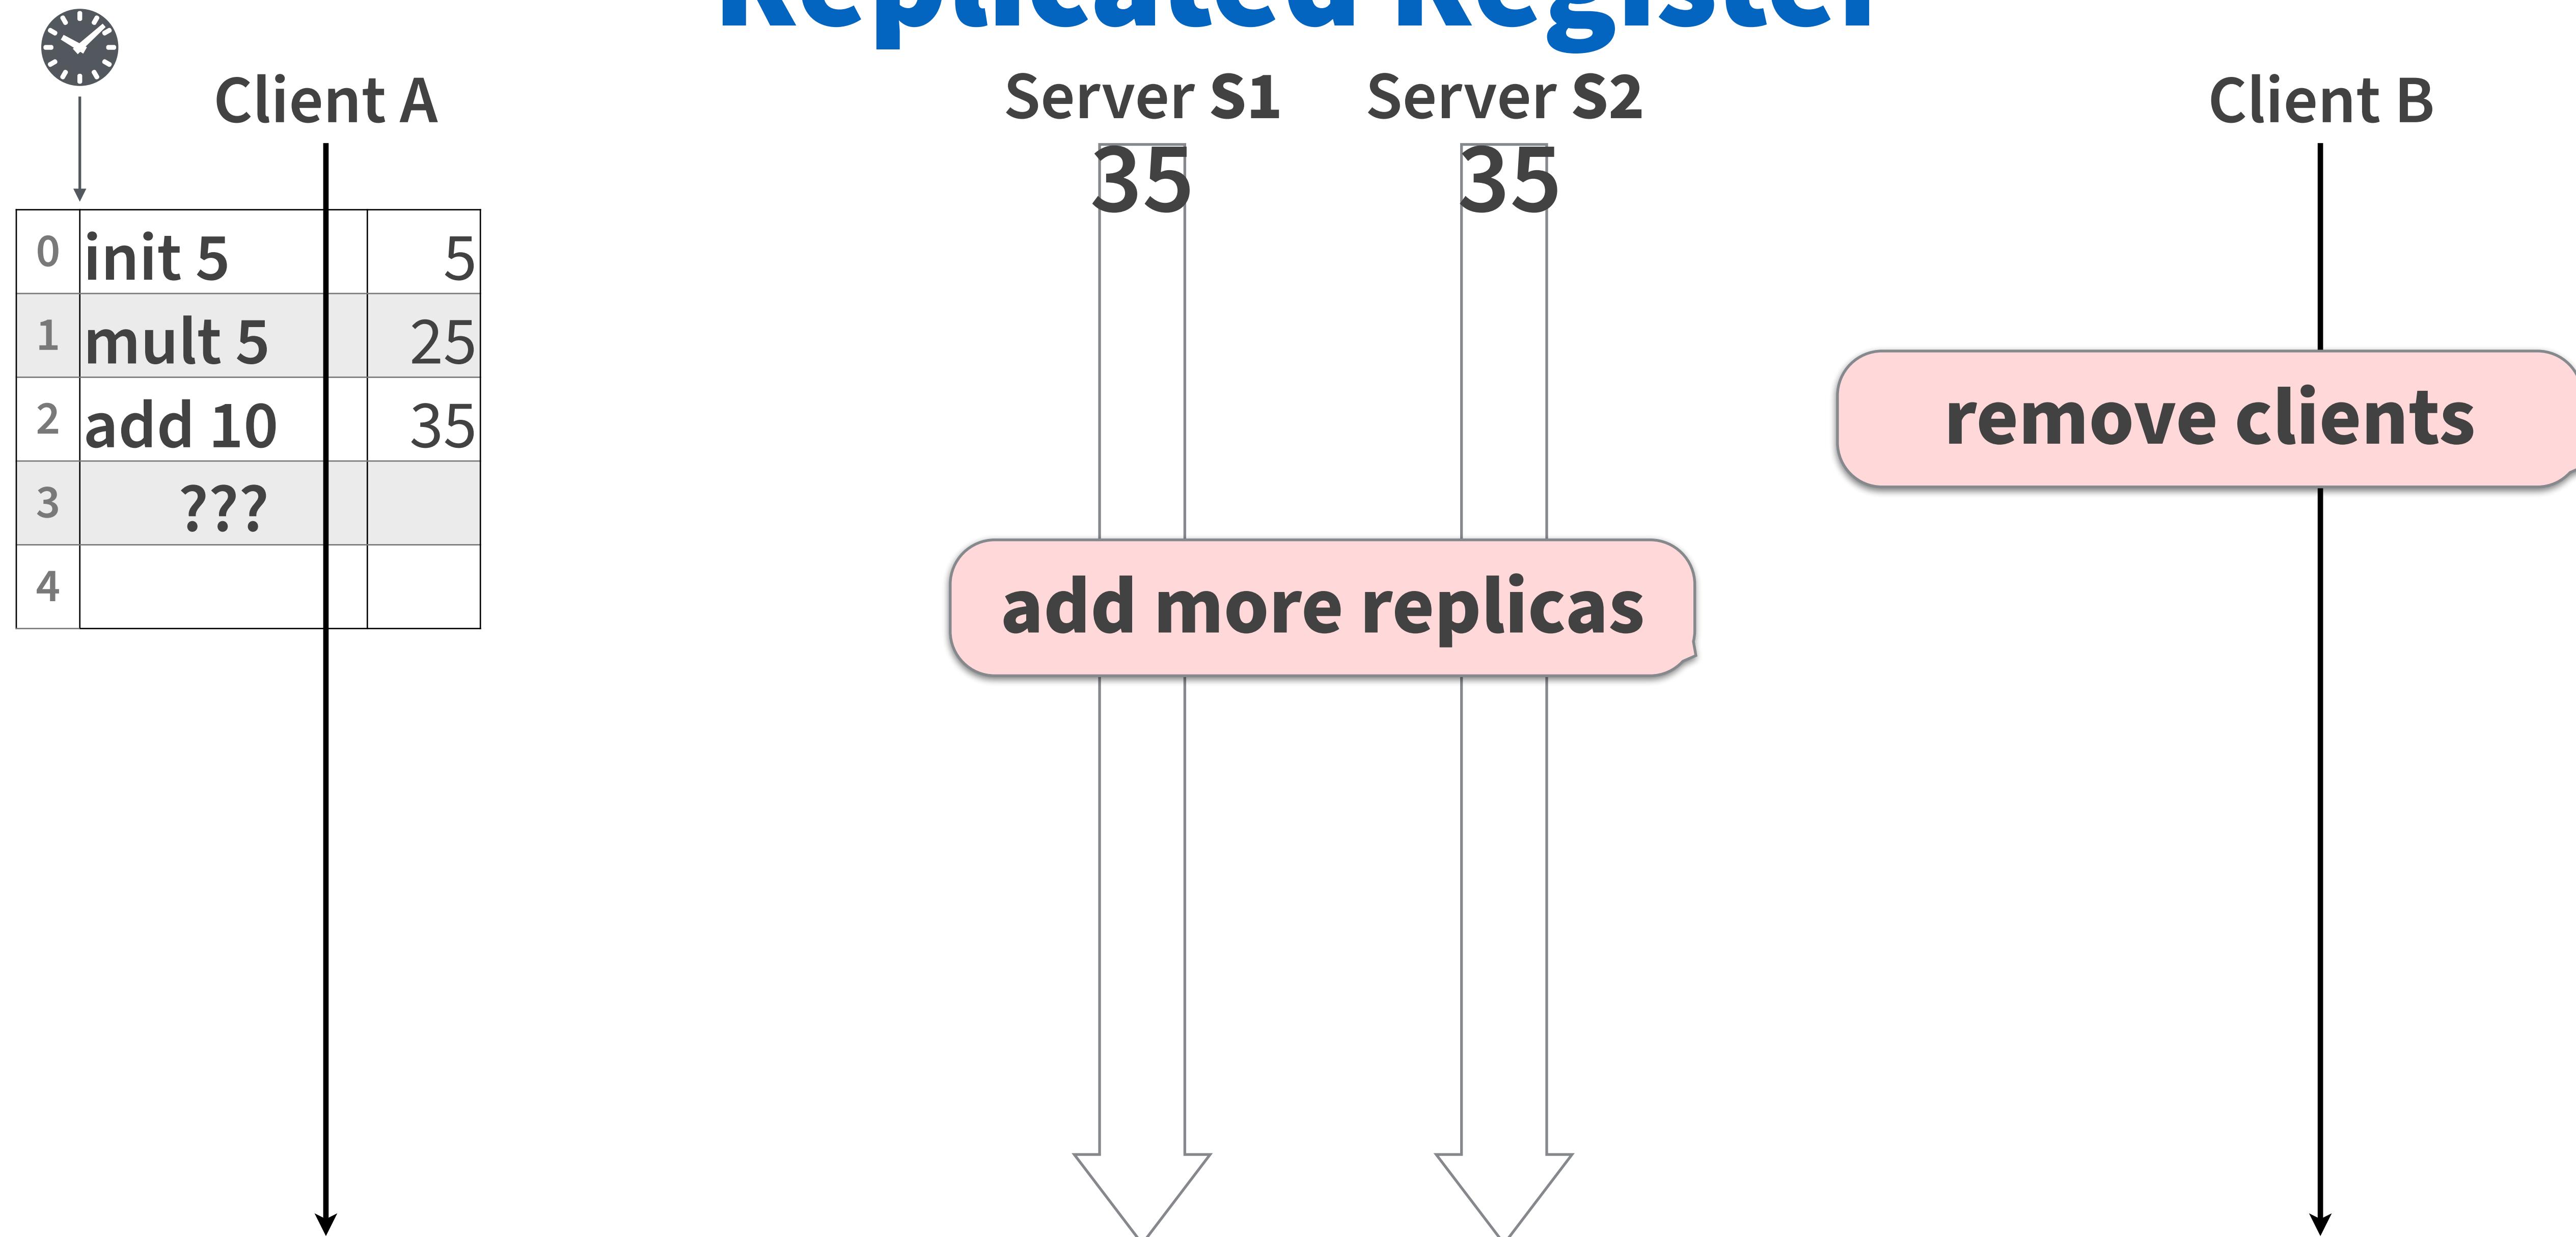
Client A

Server S1 Server S2

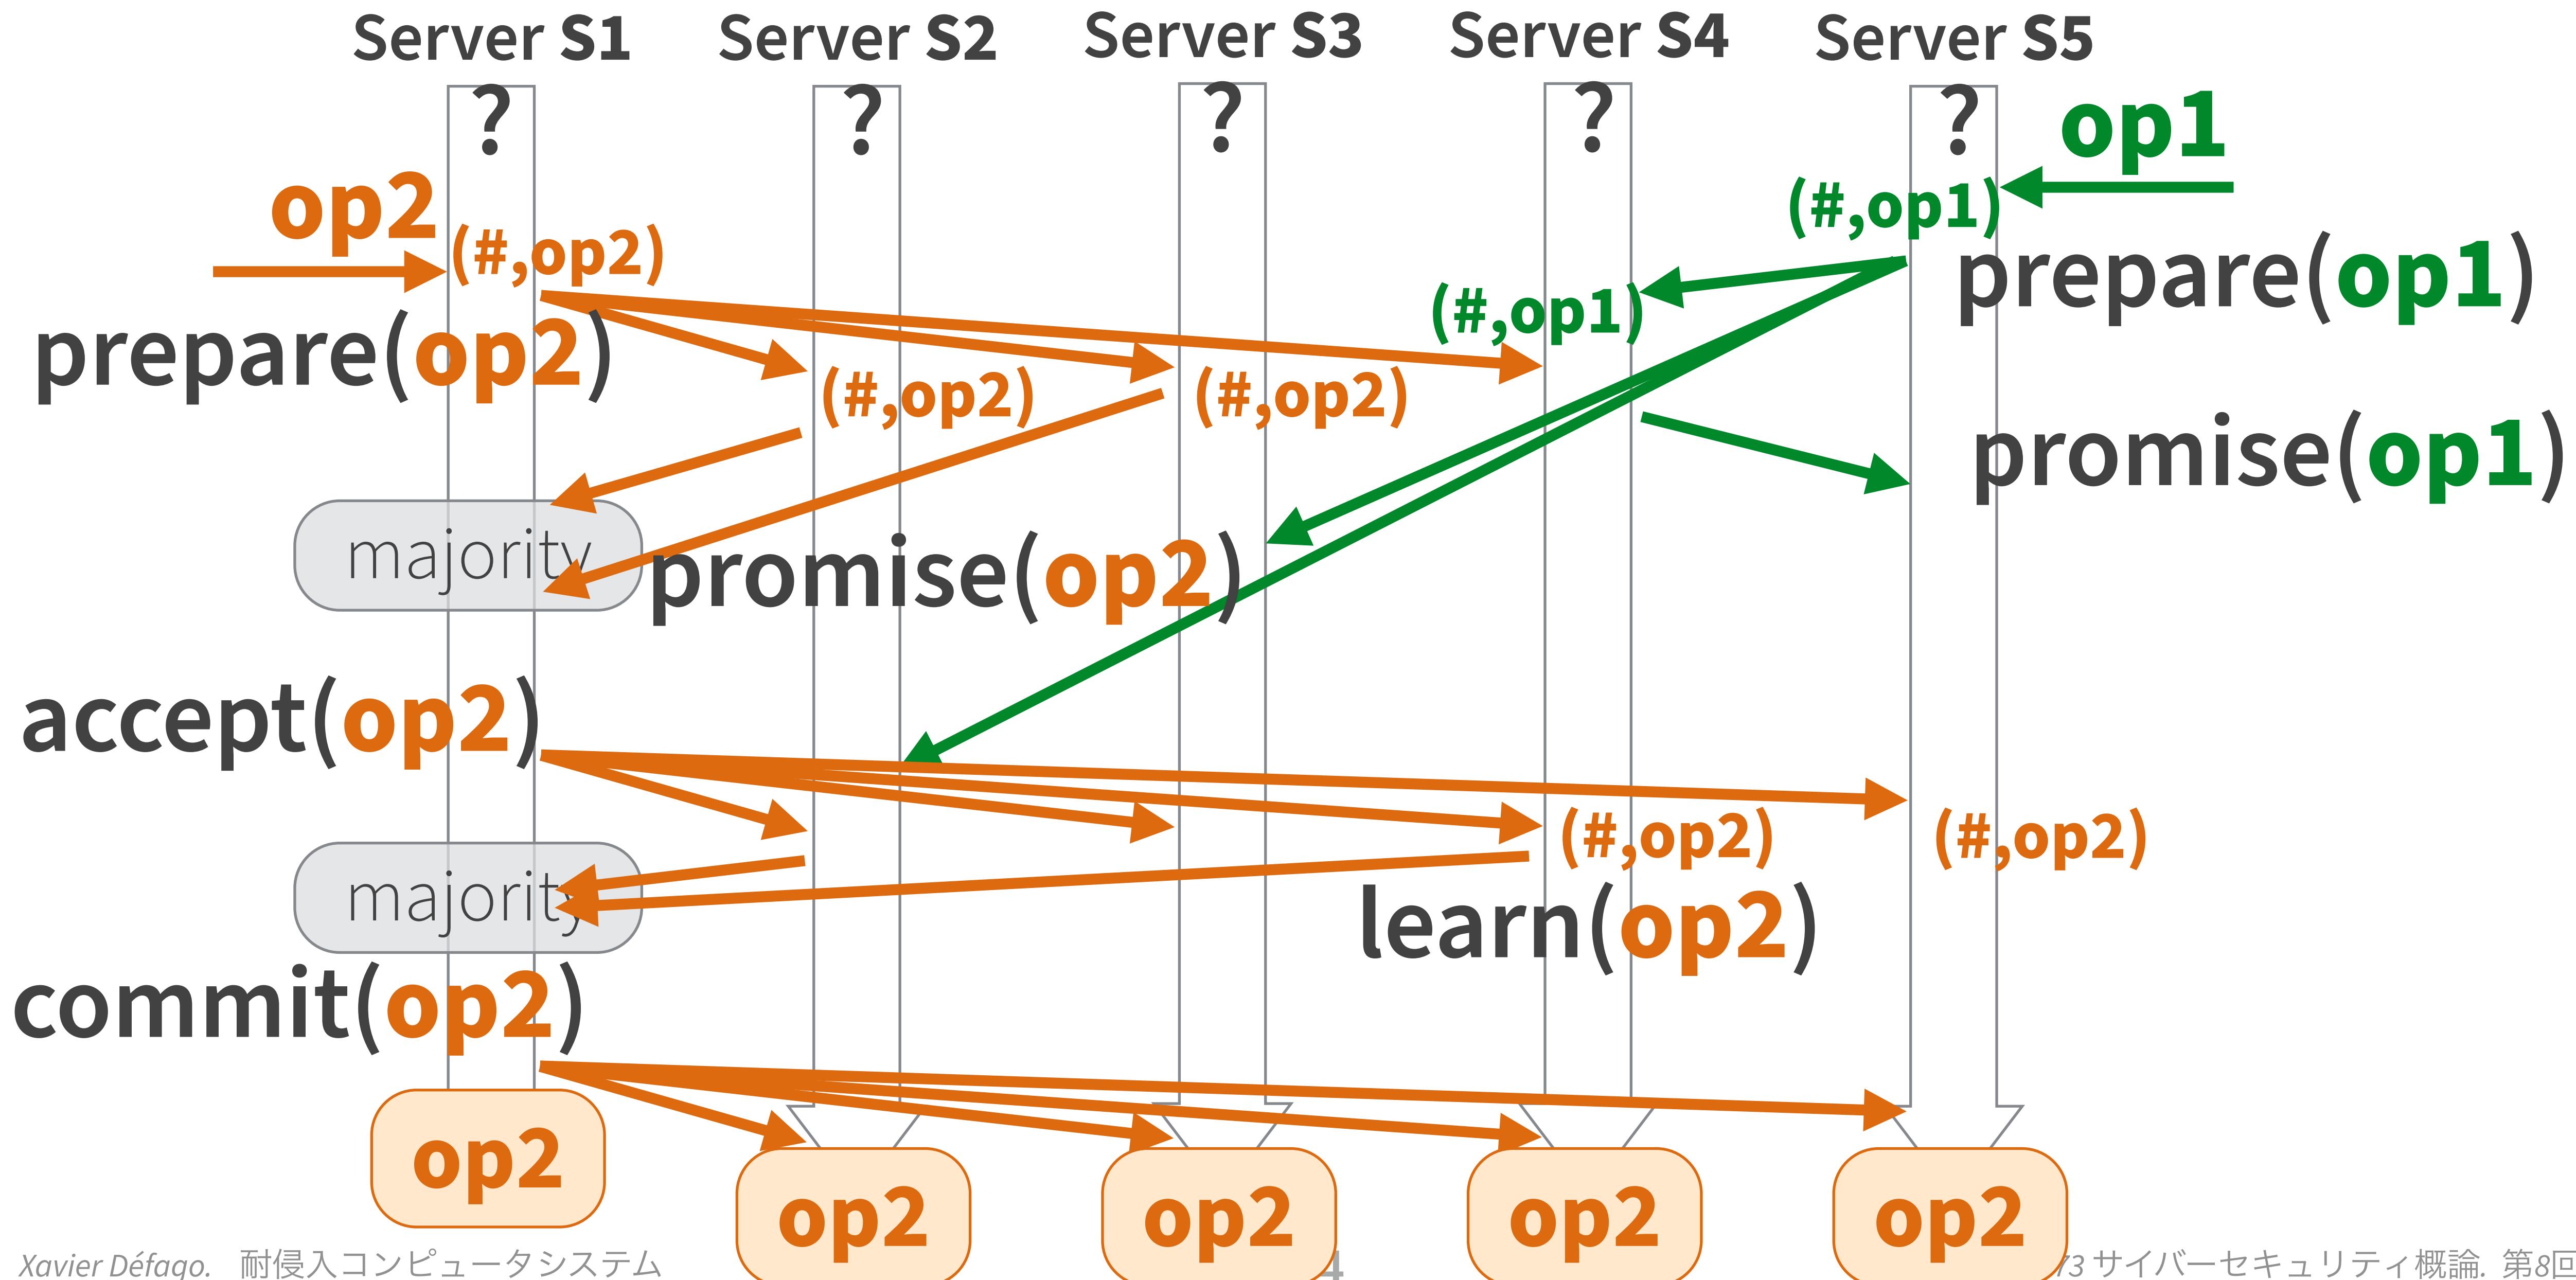
Client B



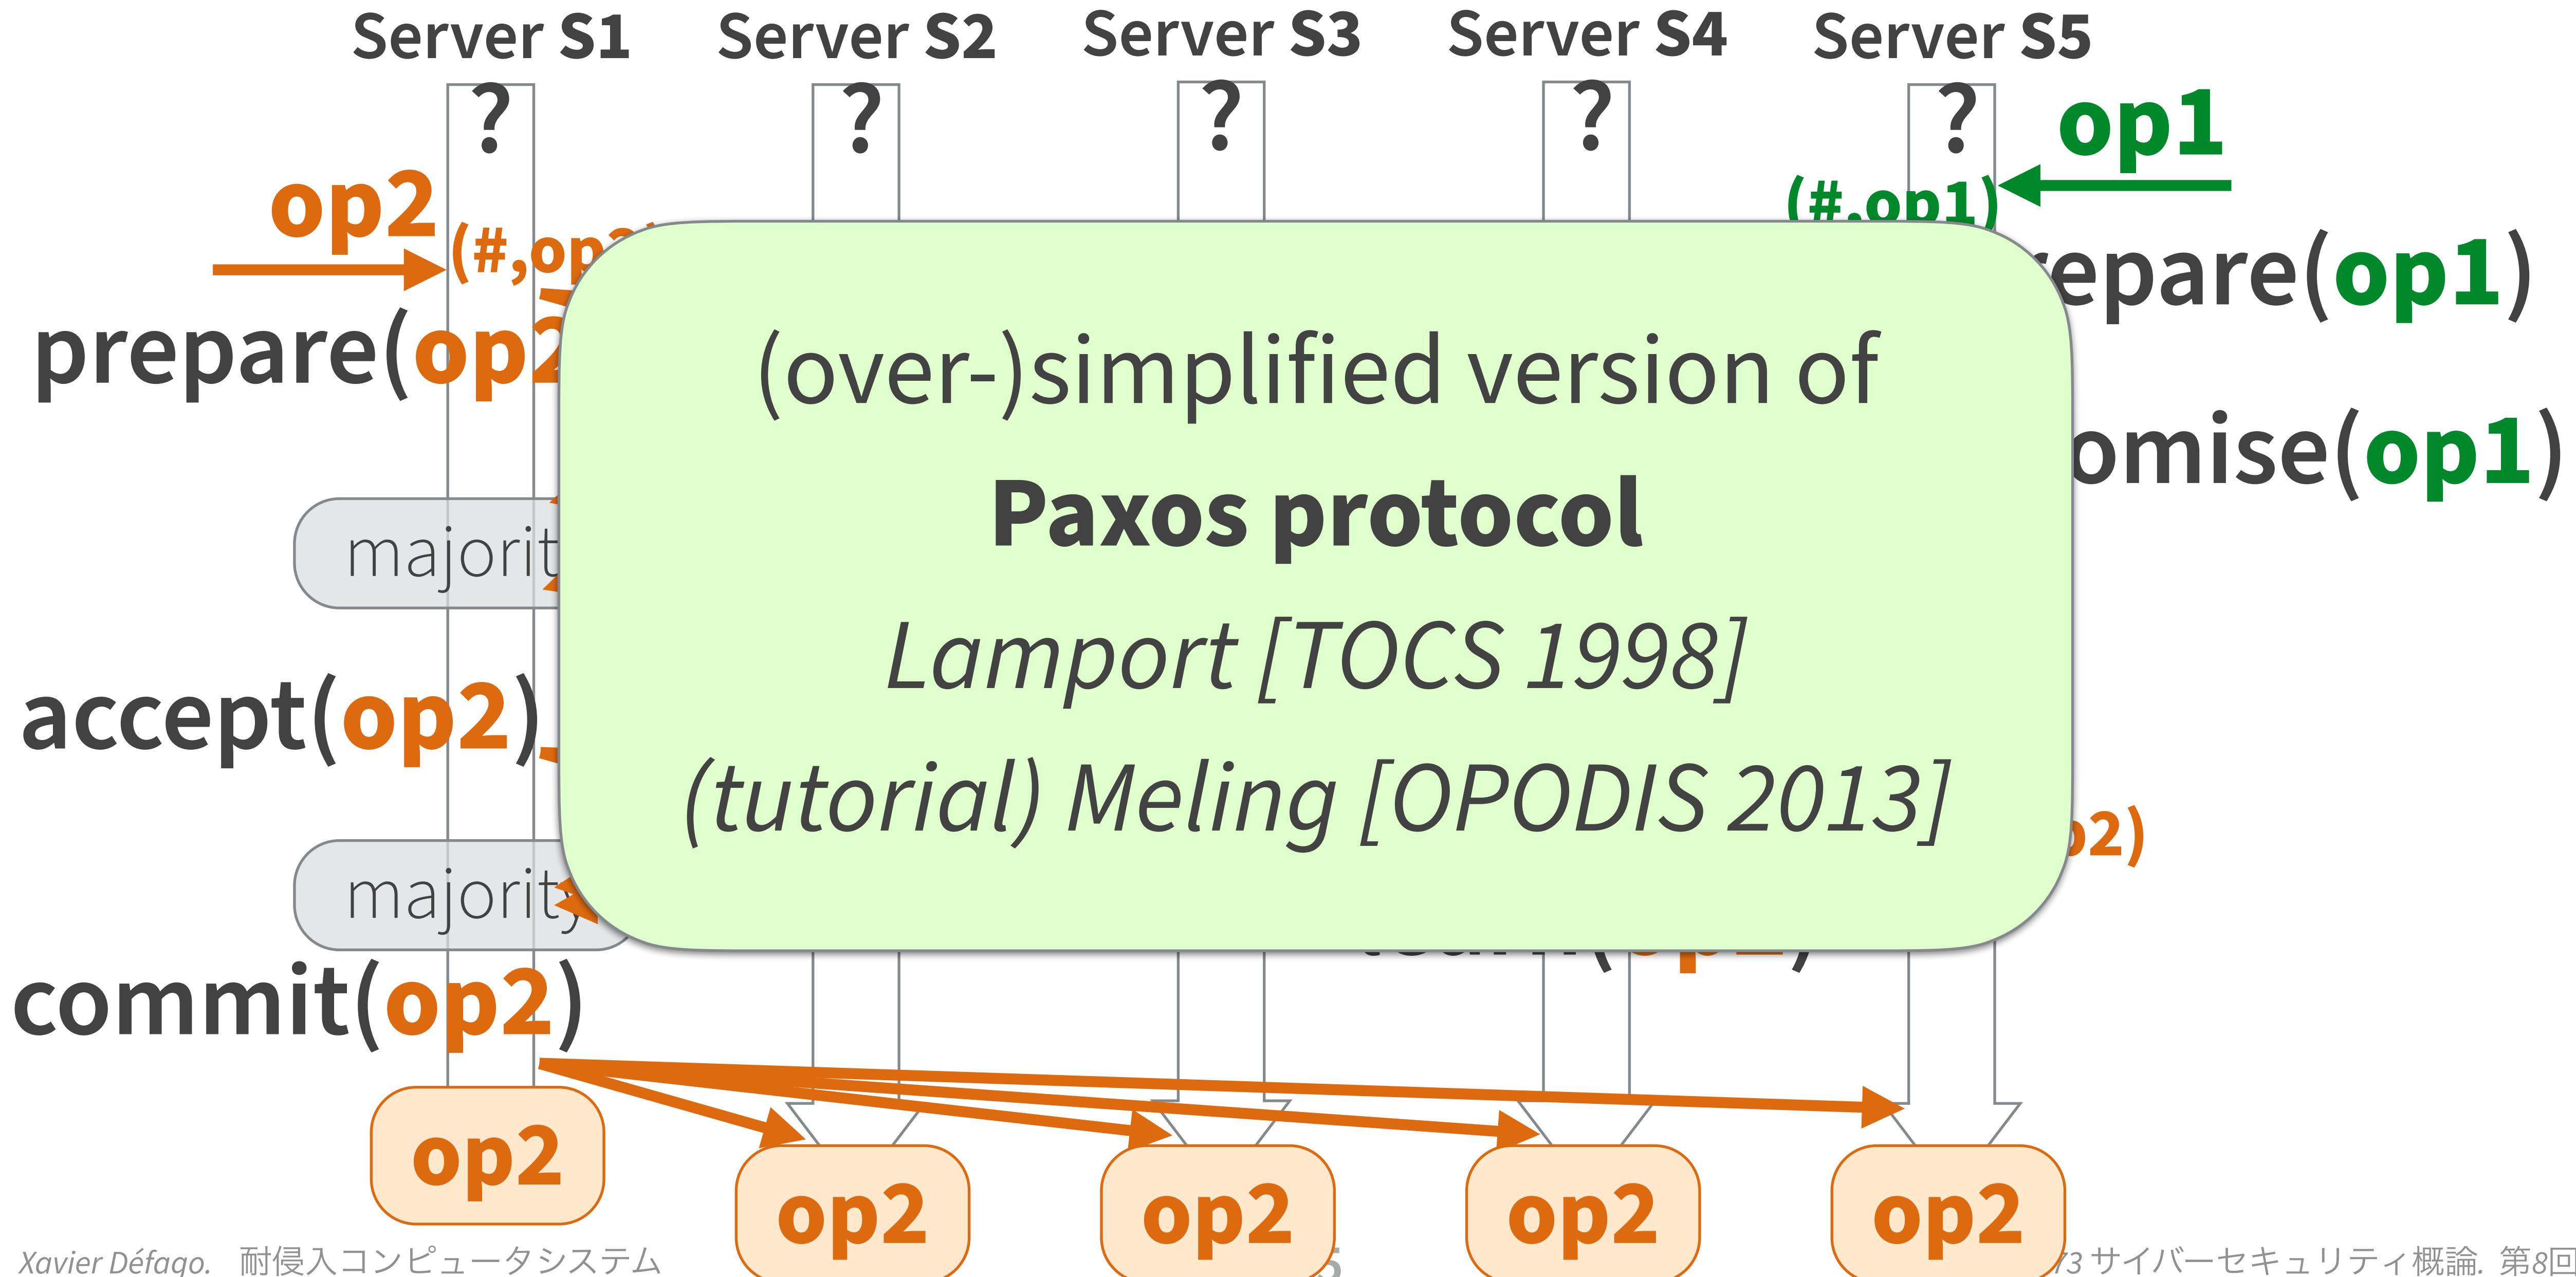
Replicated Register



Replicated Register



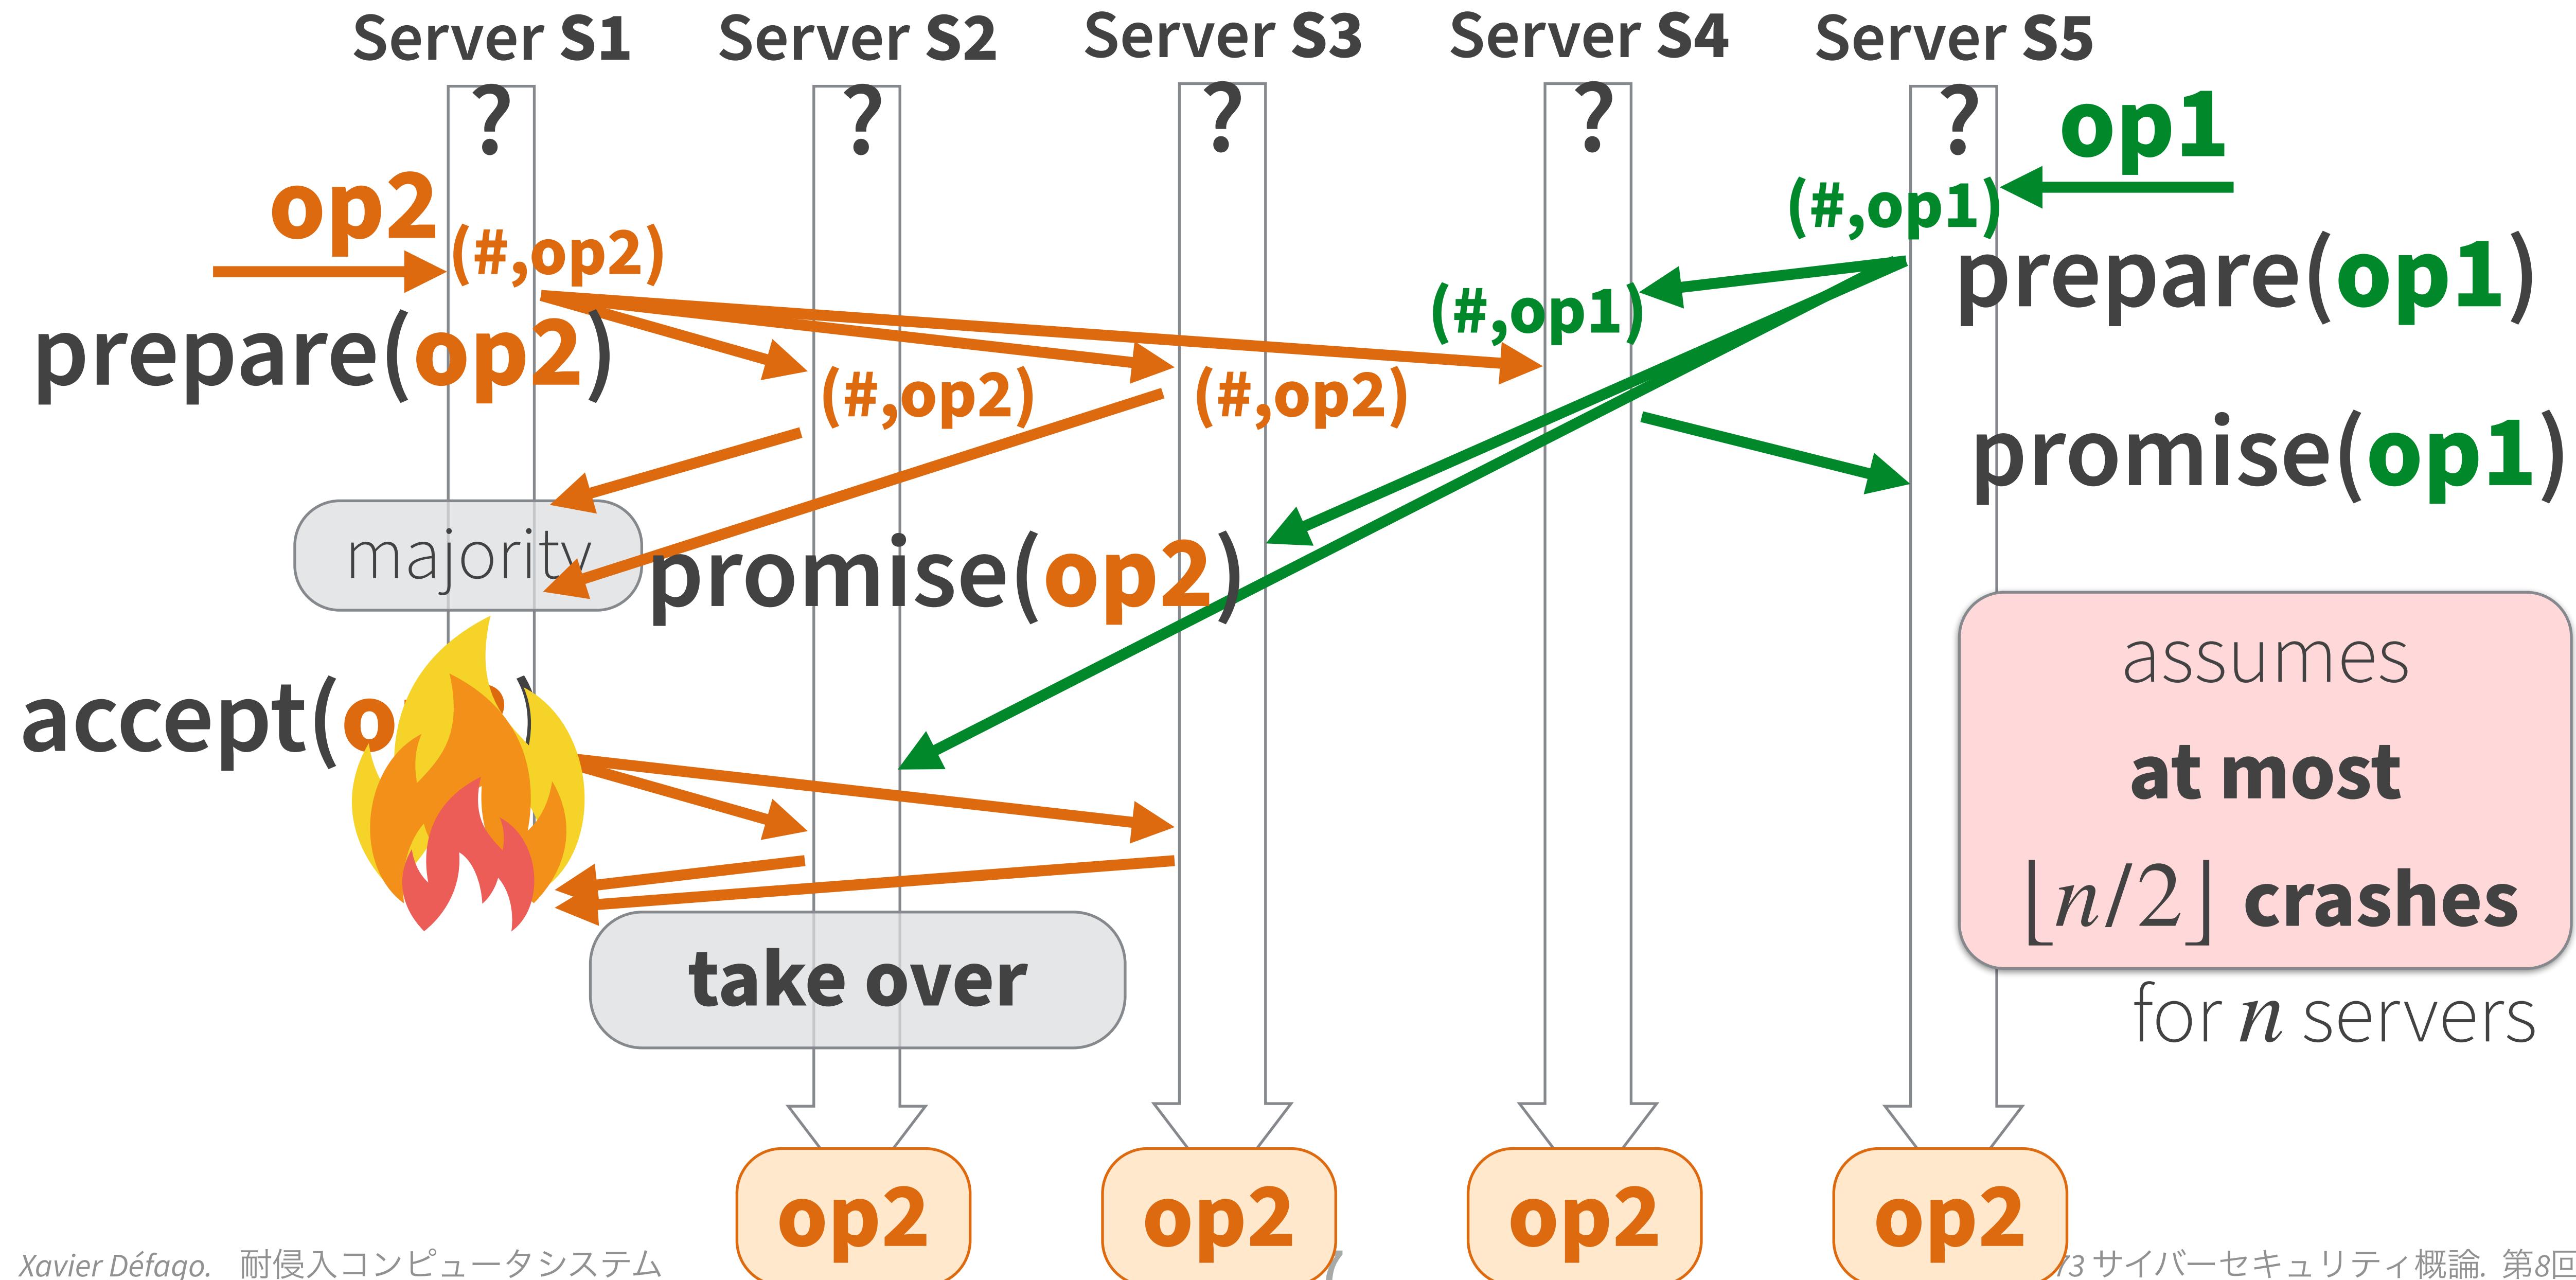
Replicated Register



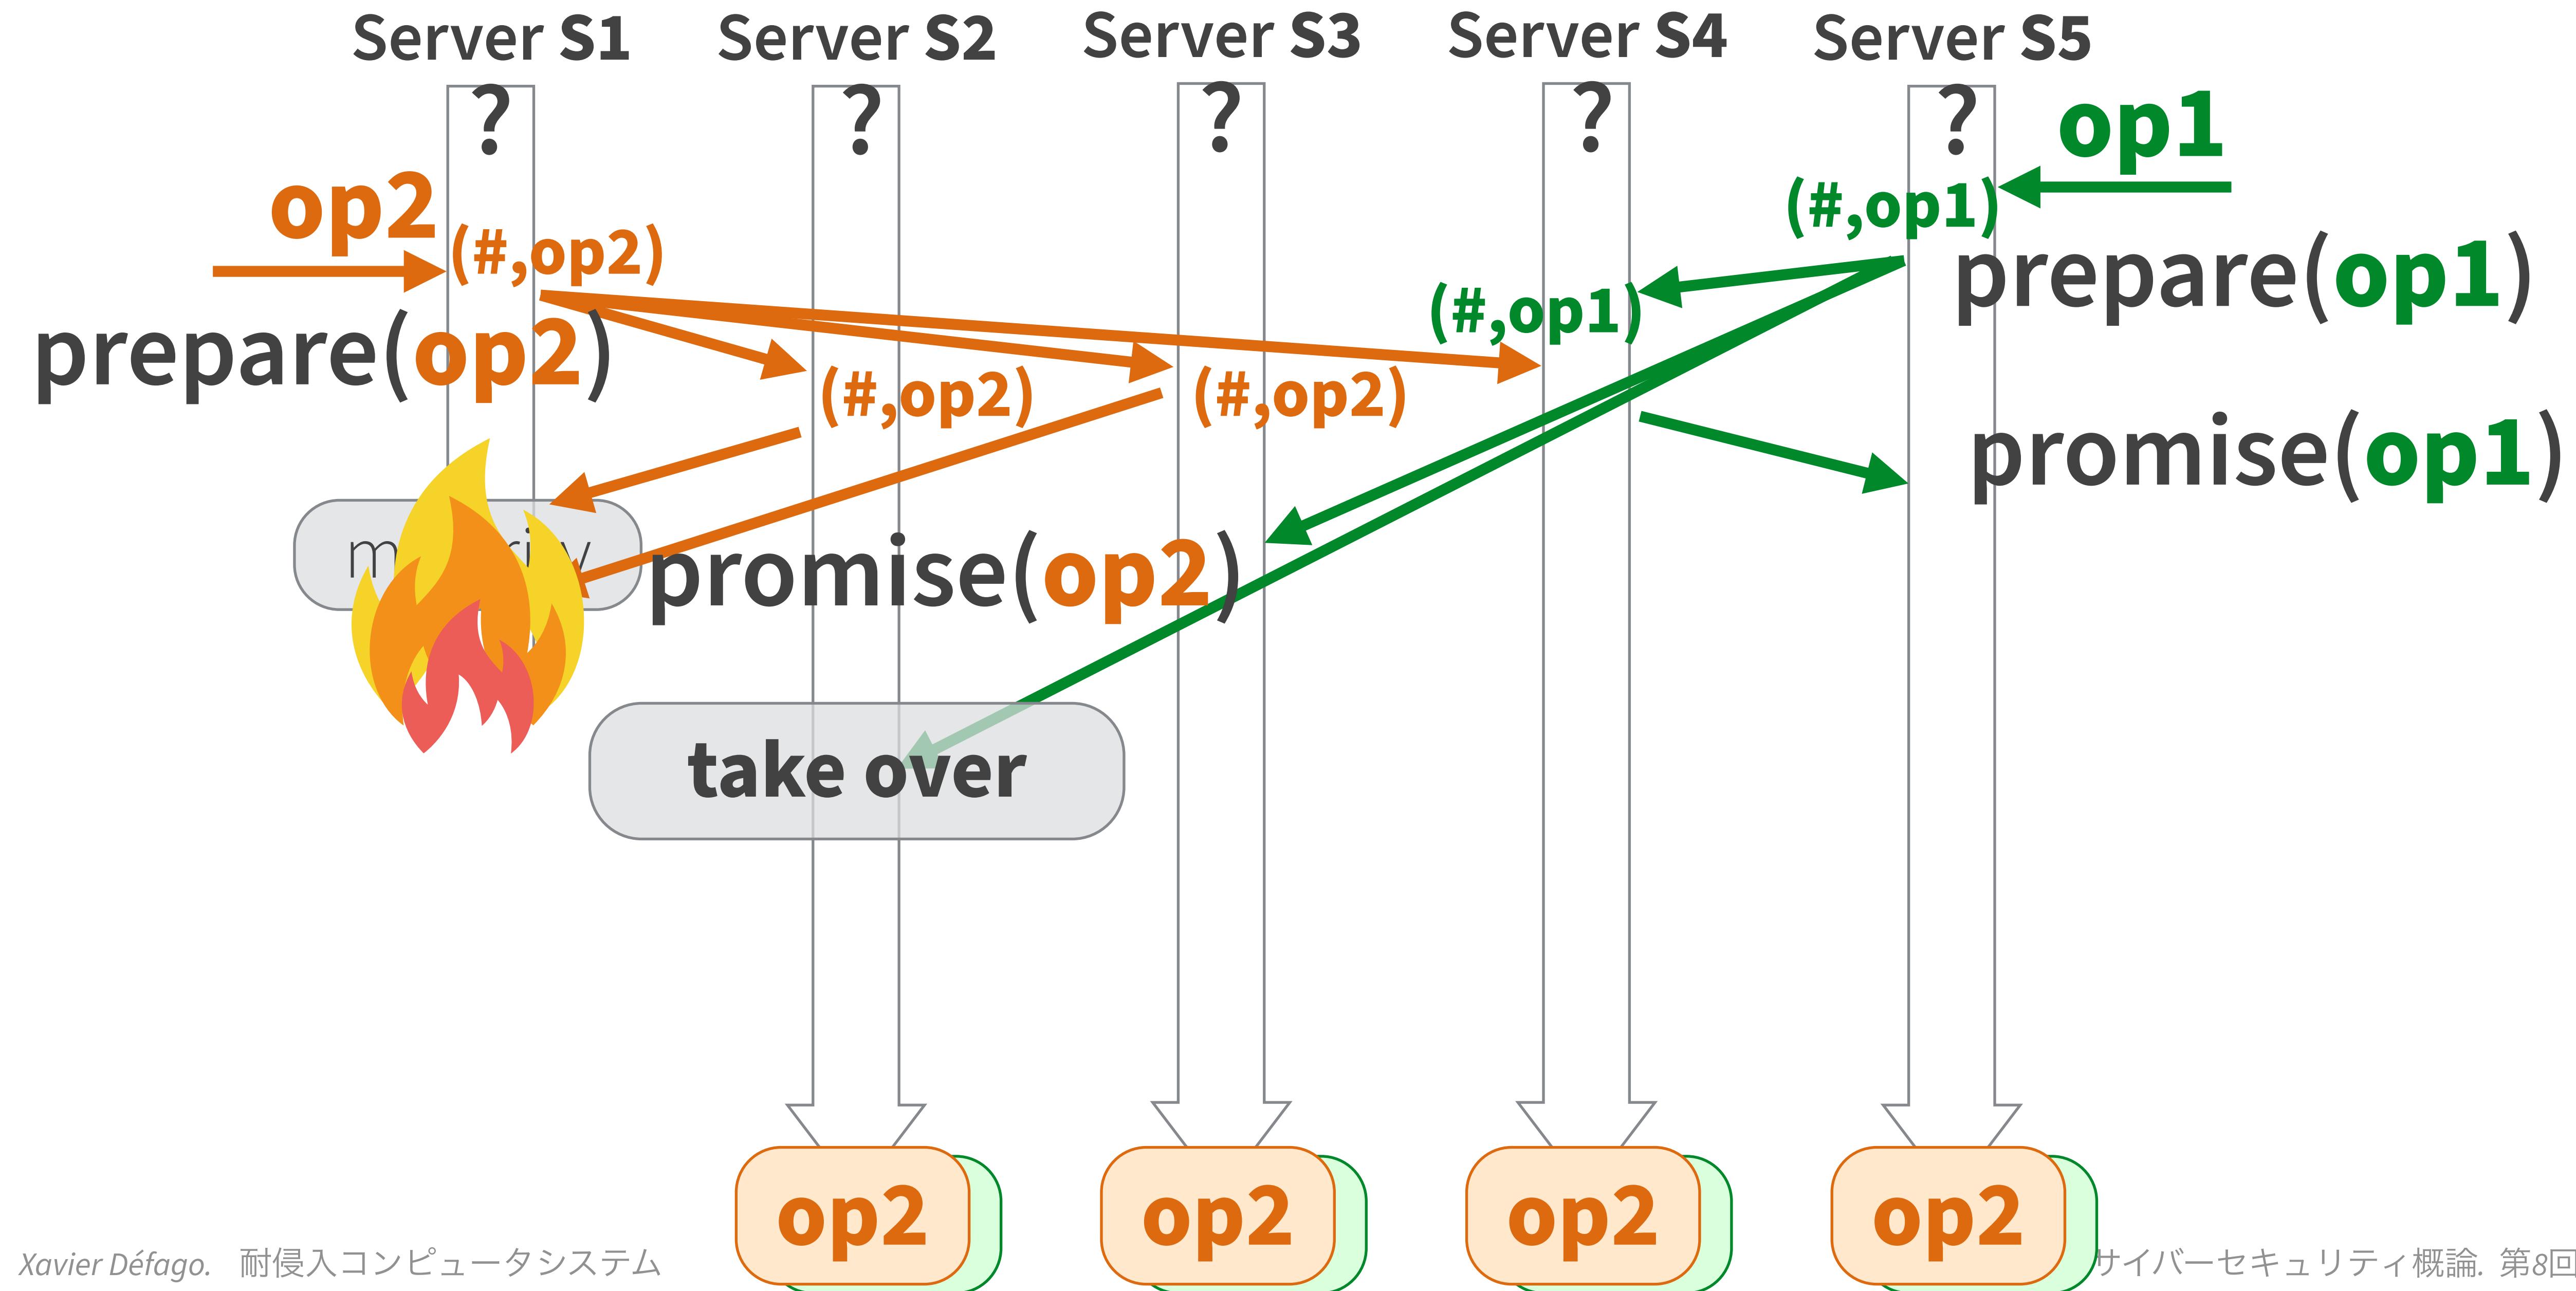
Quorums

Why majority?

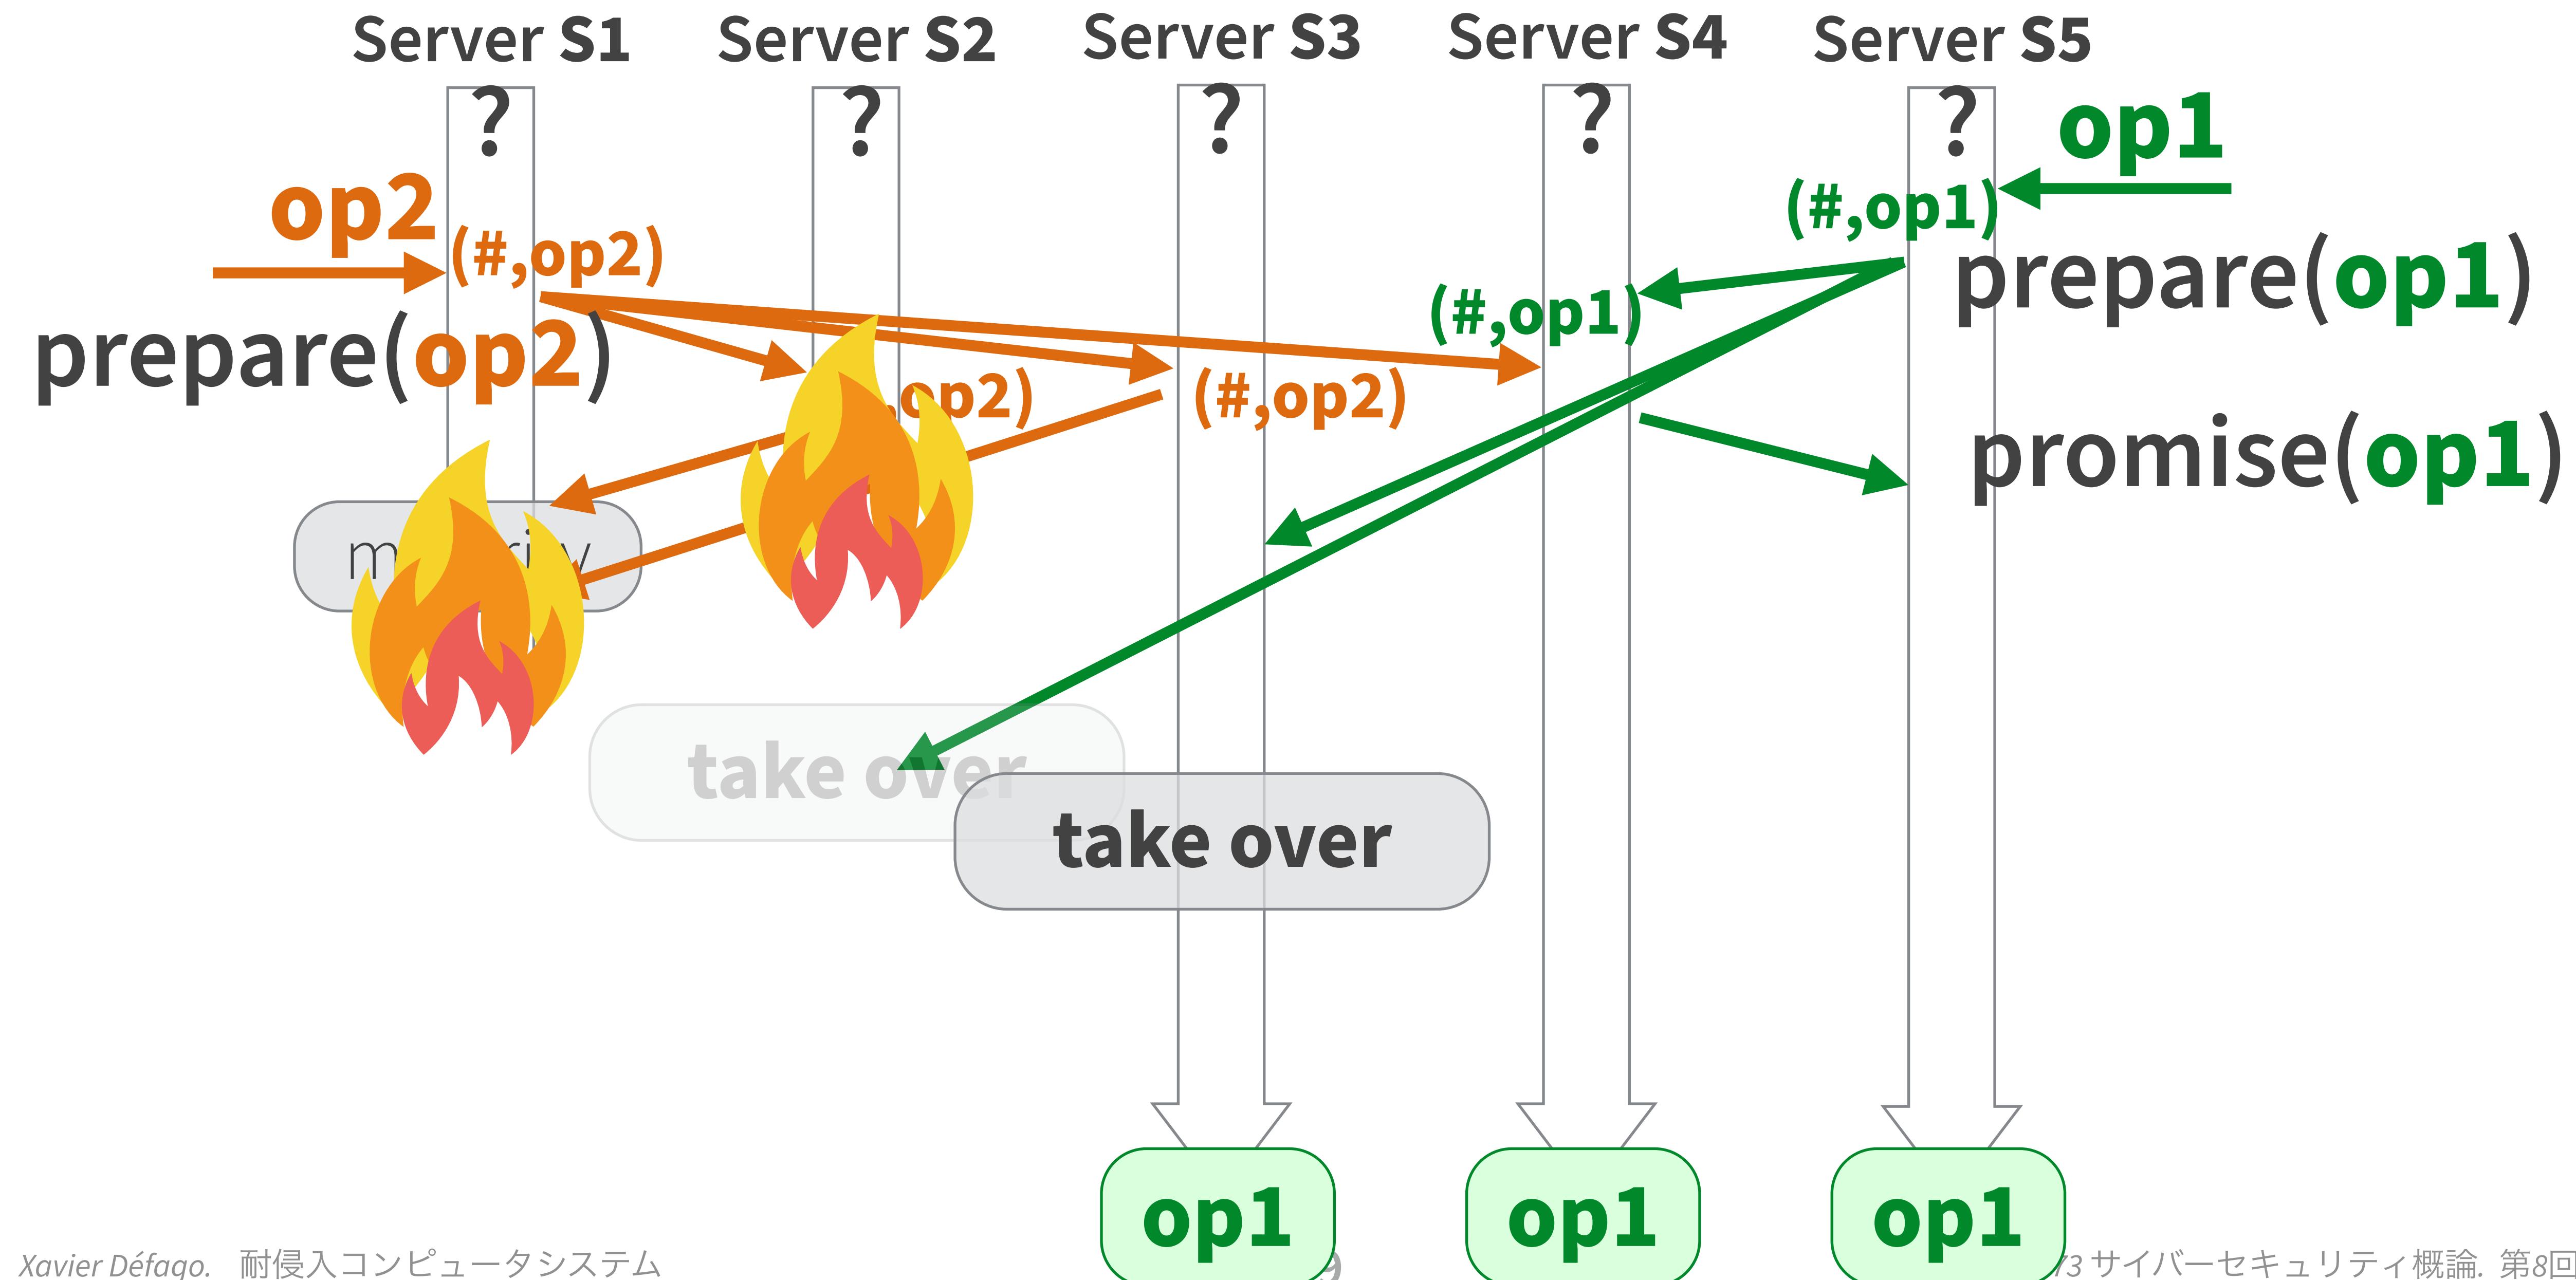
Replicated Register



Replicated Register



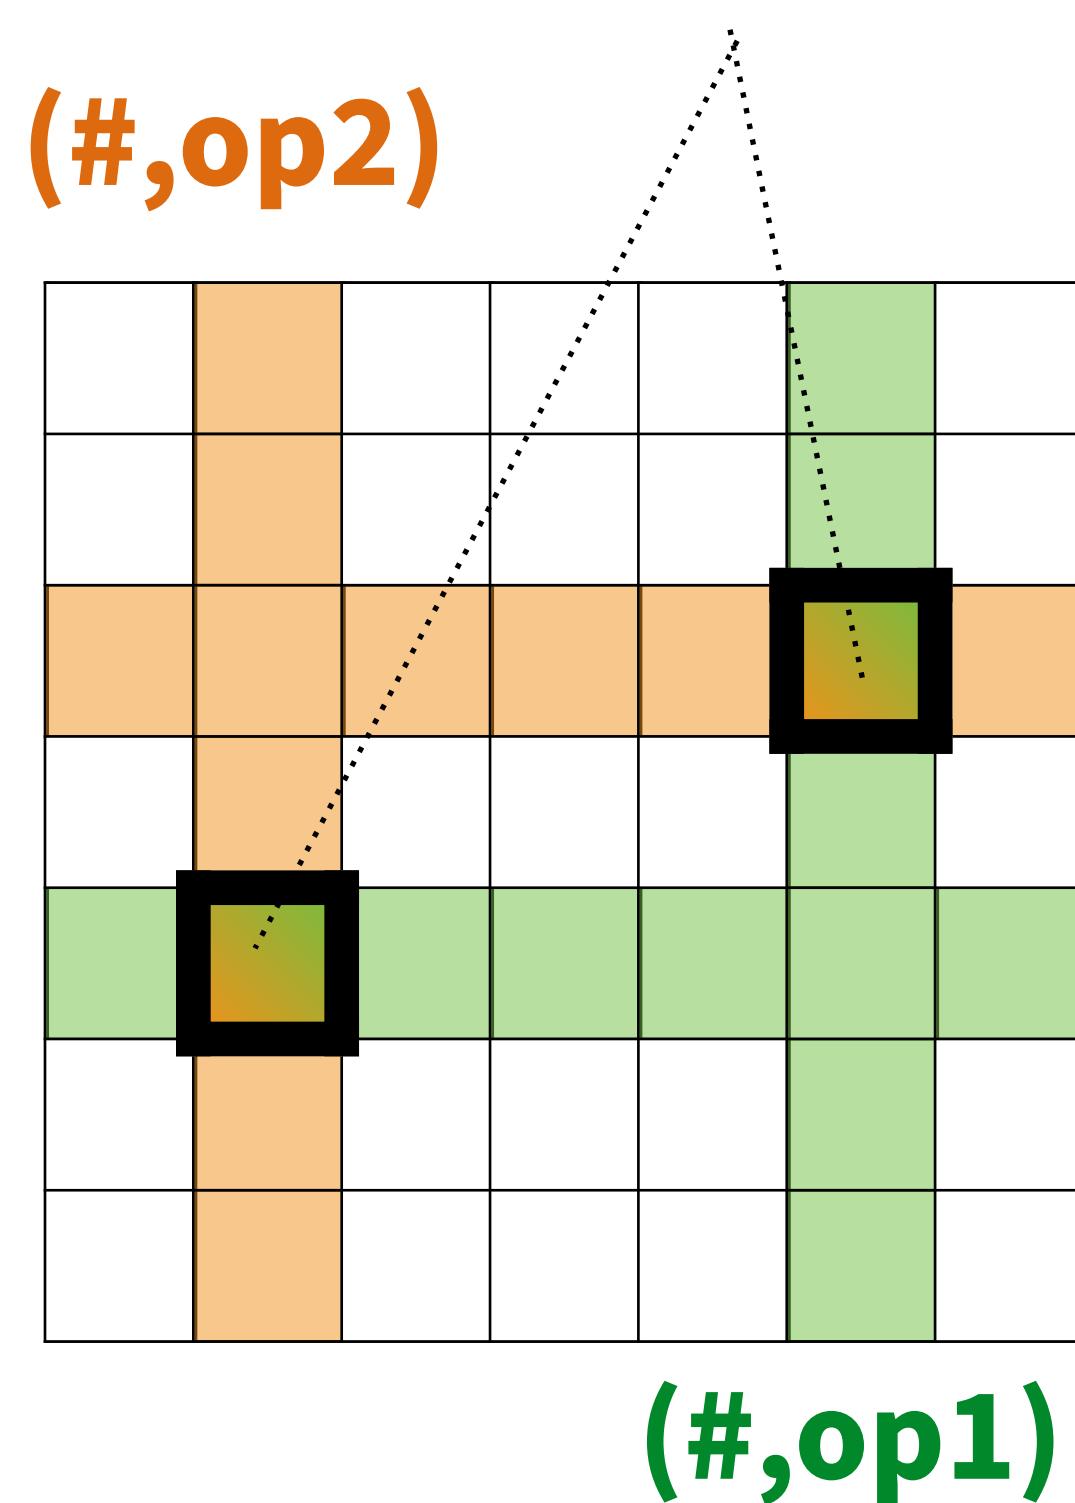
Quorums



Quorums

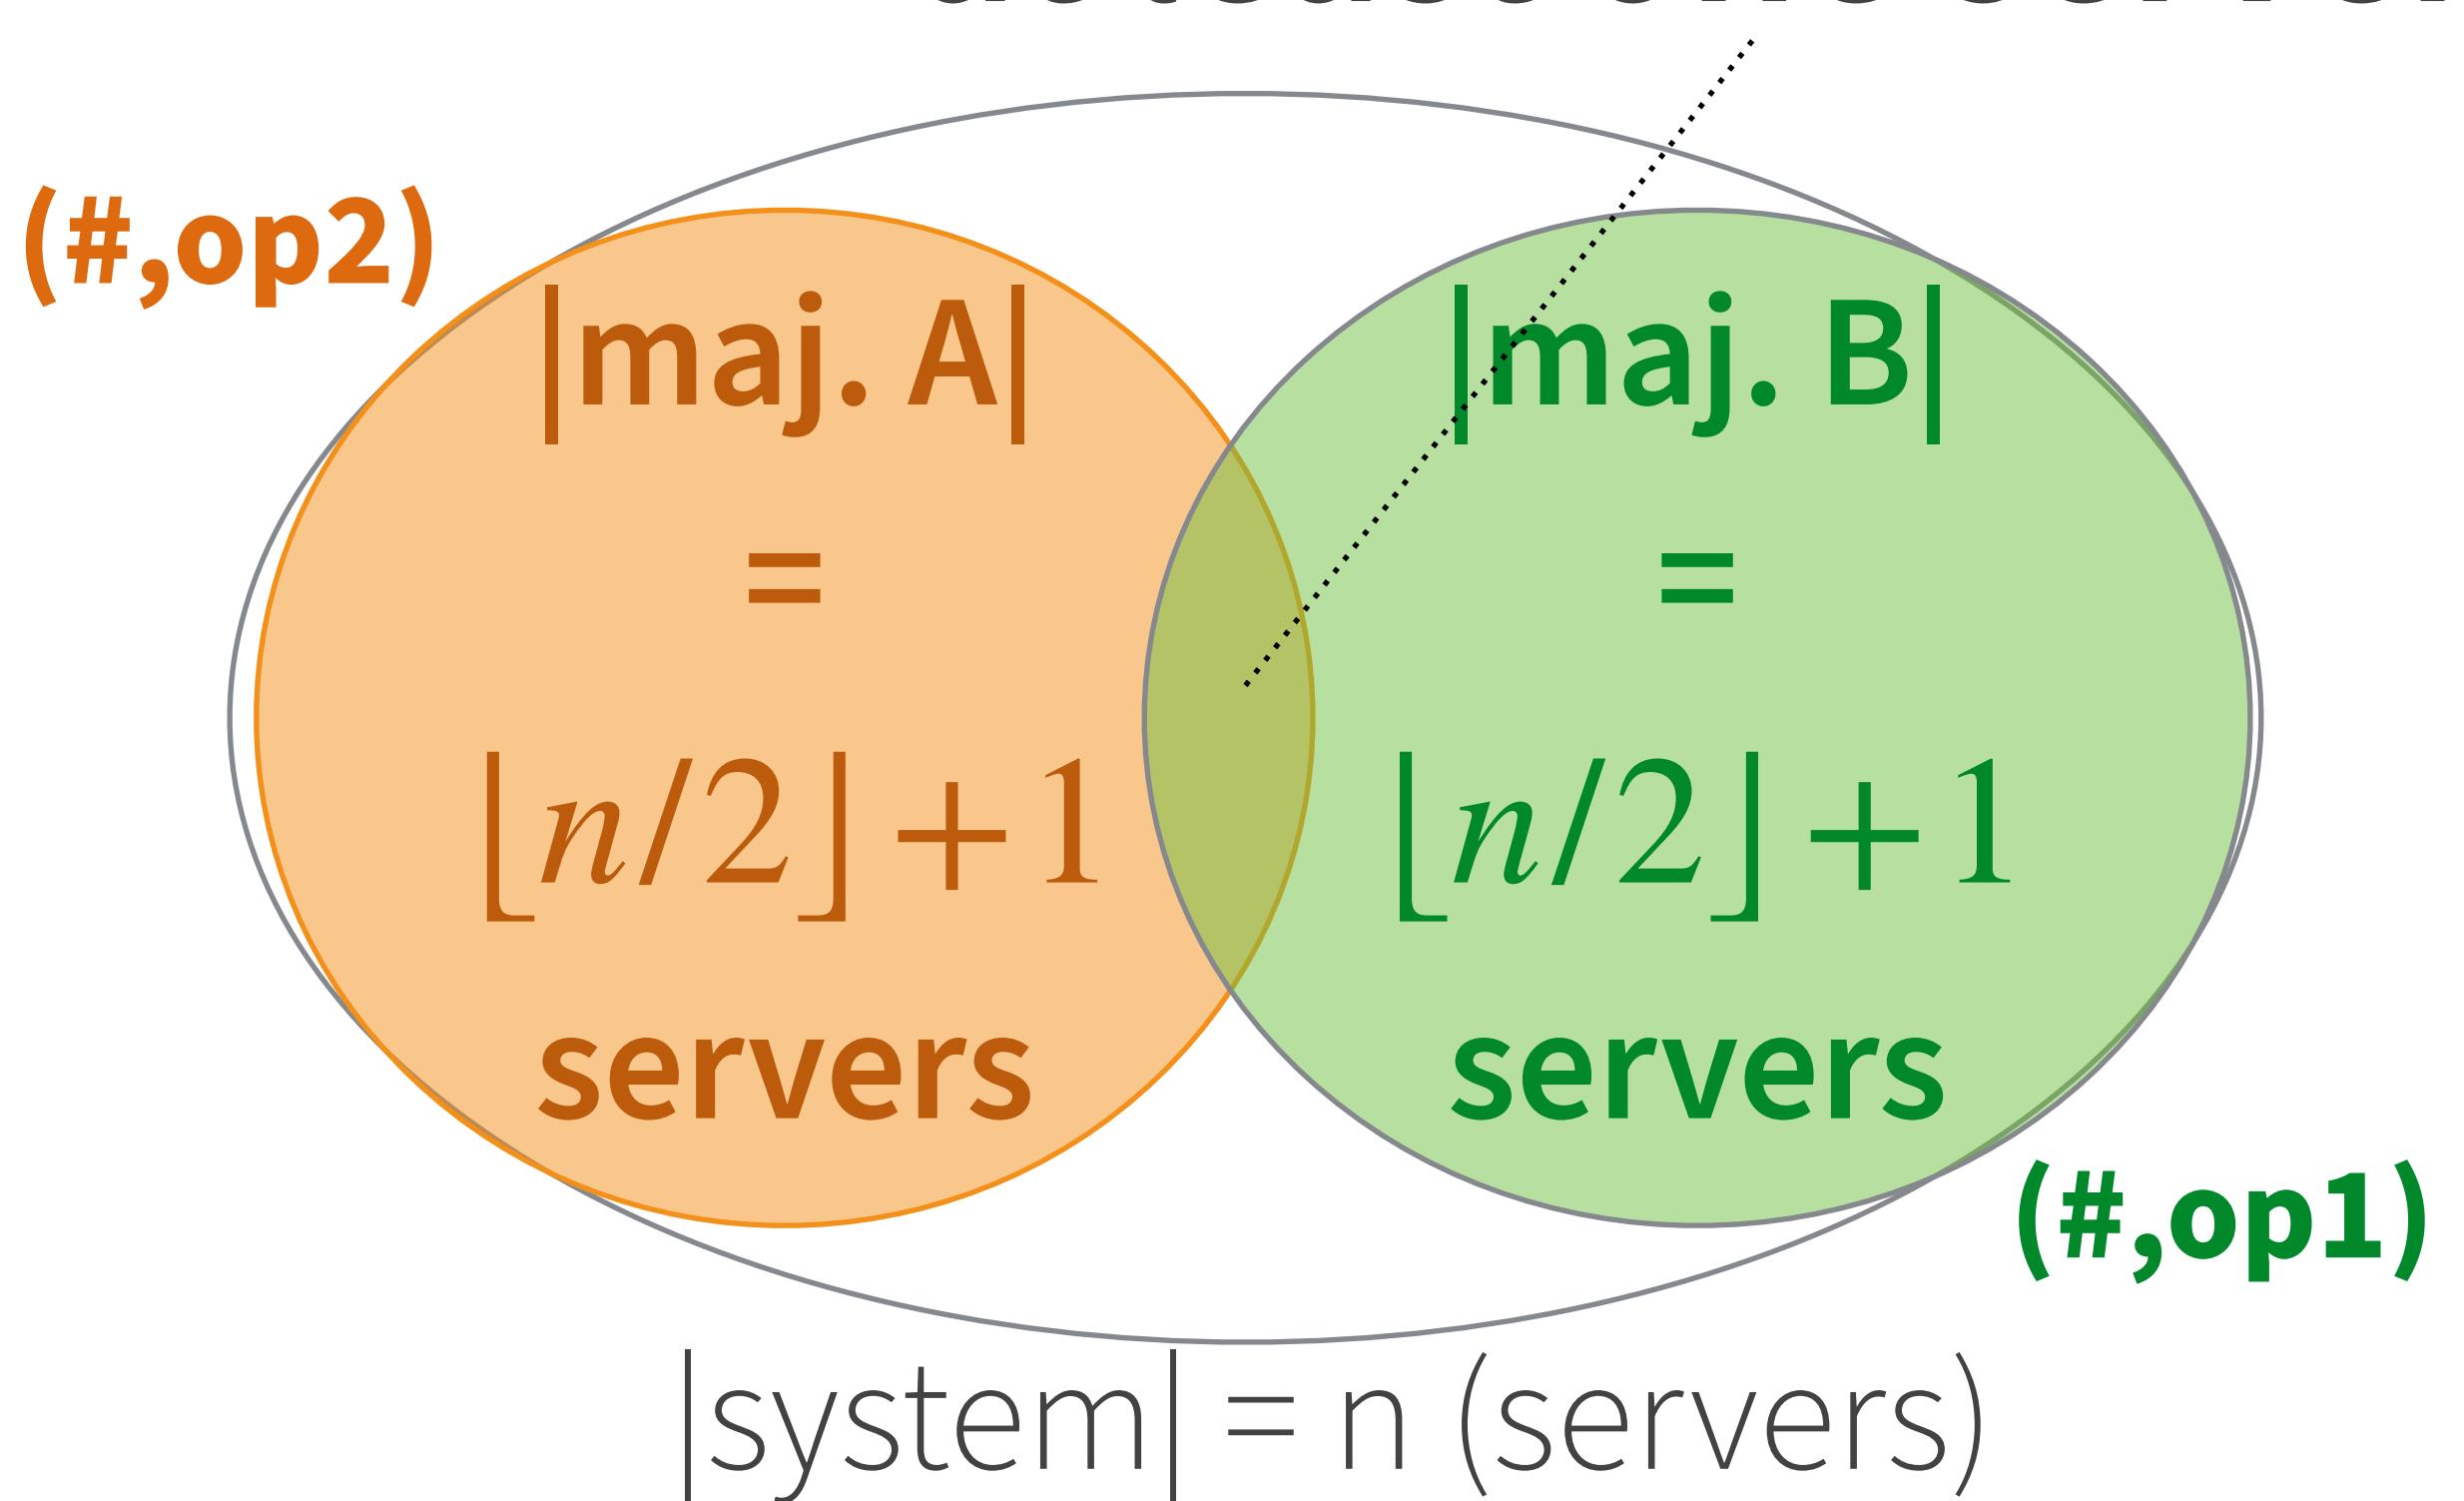
grid quorum

at least two servers



majority quorum

at least one server



Recovery (take over)

▶ Terminology

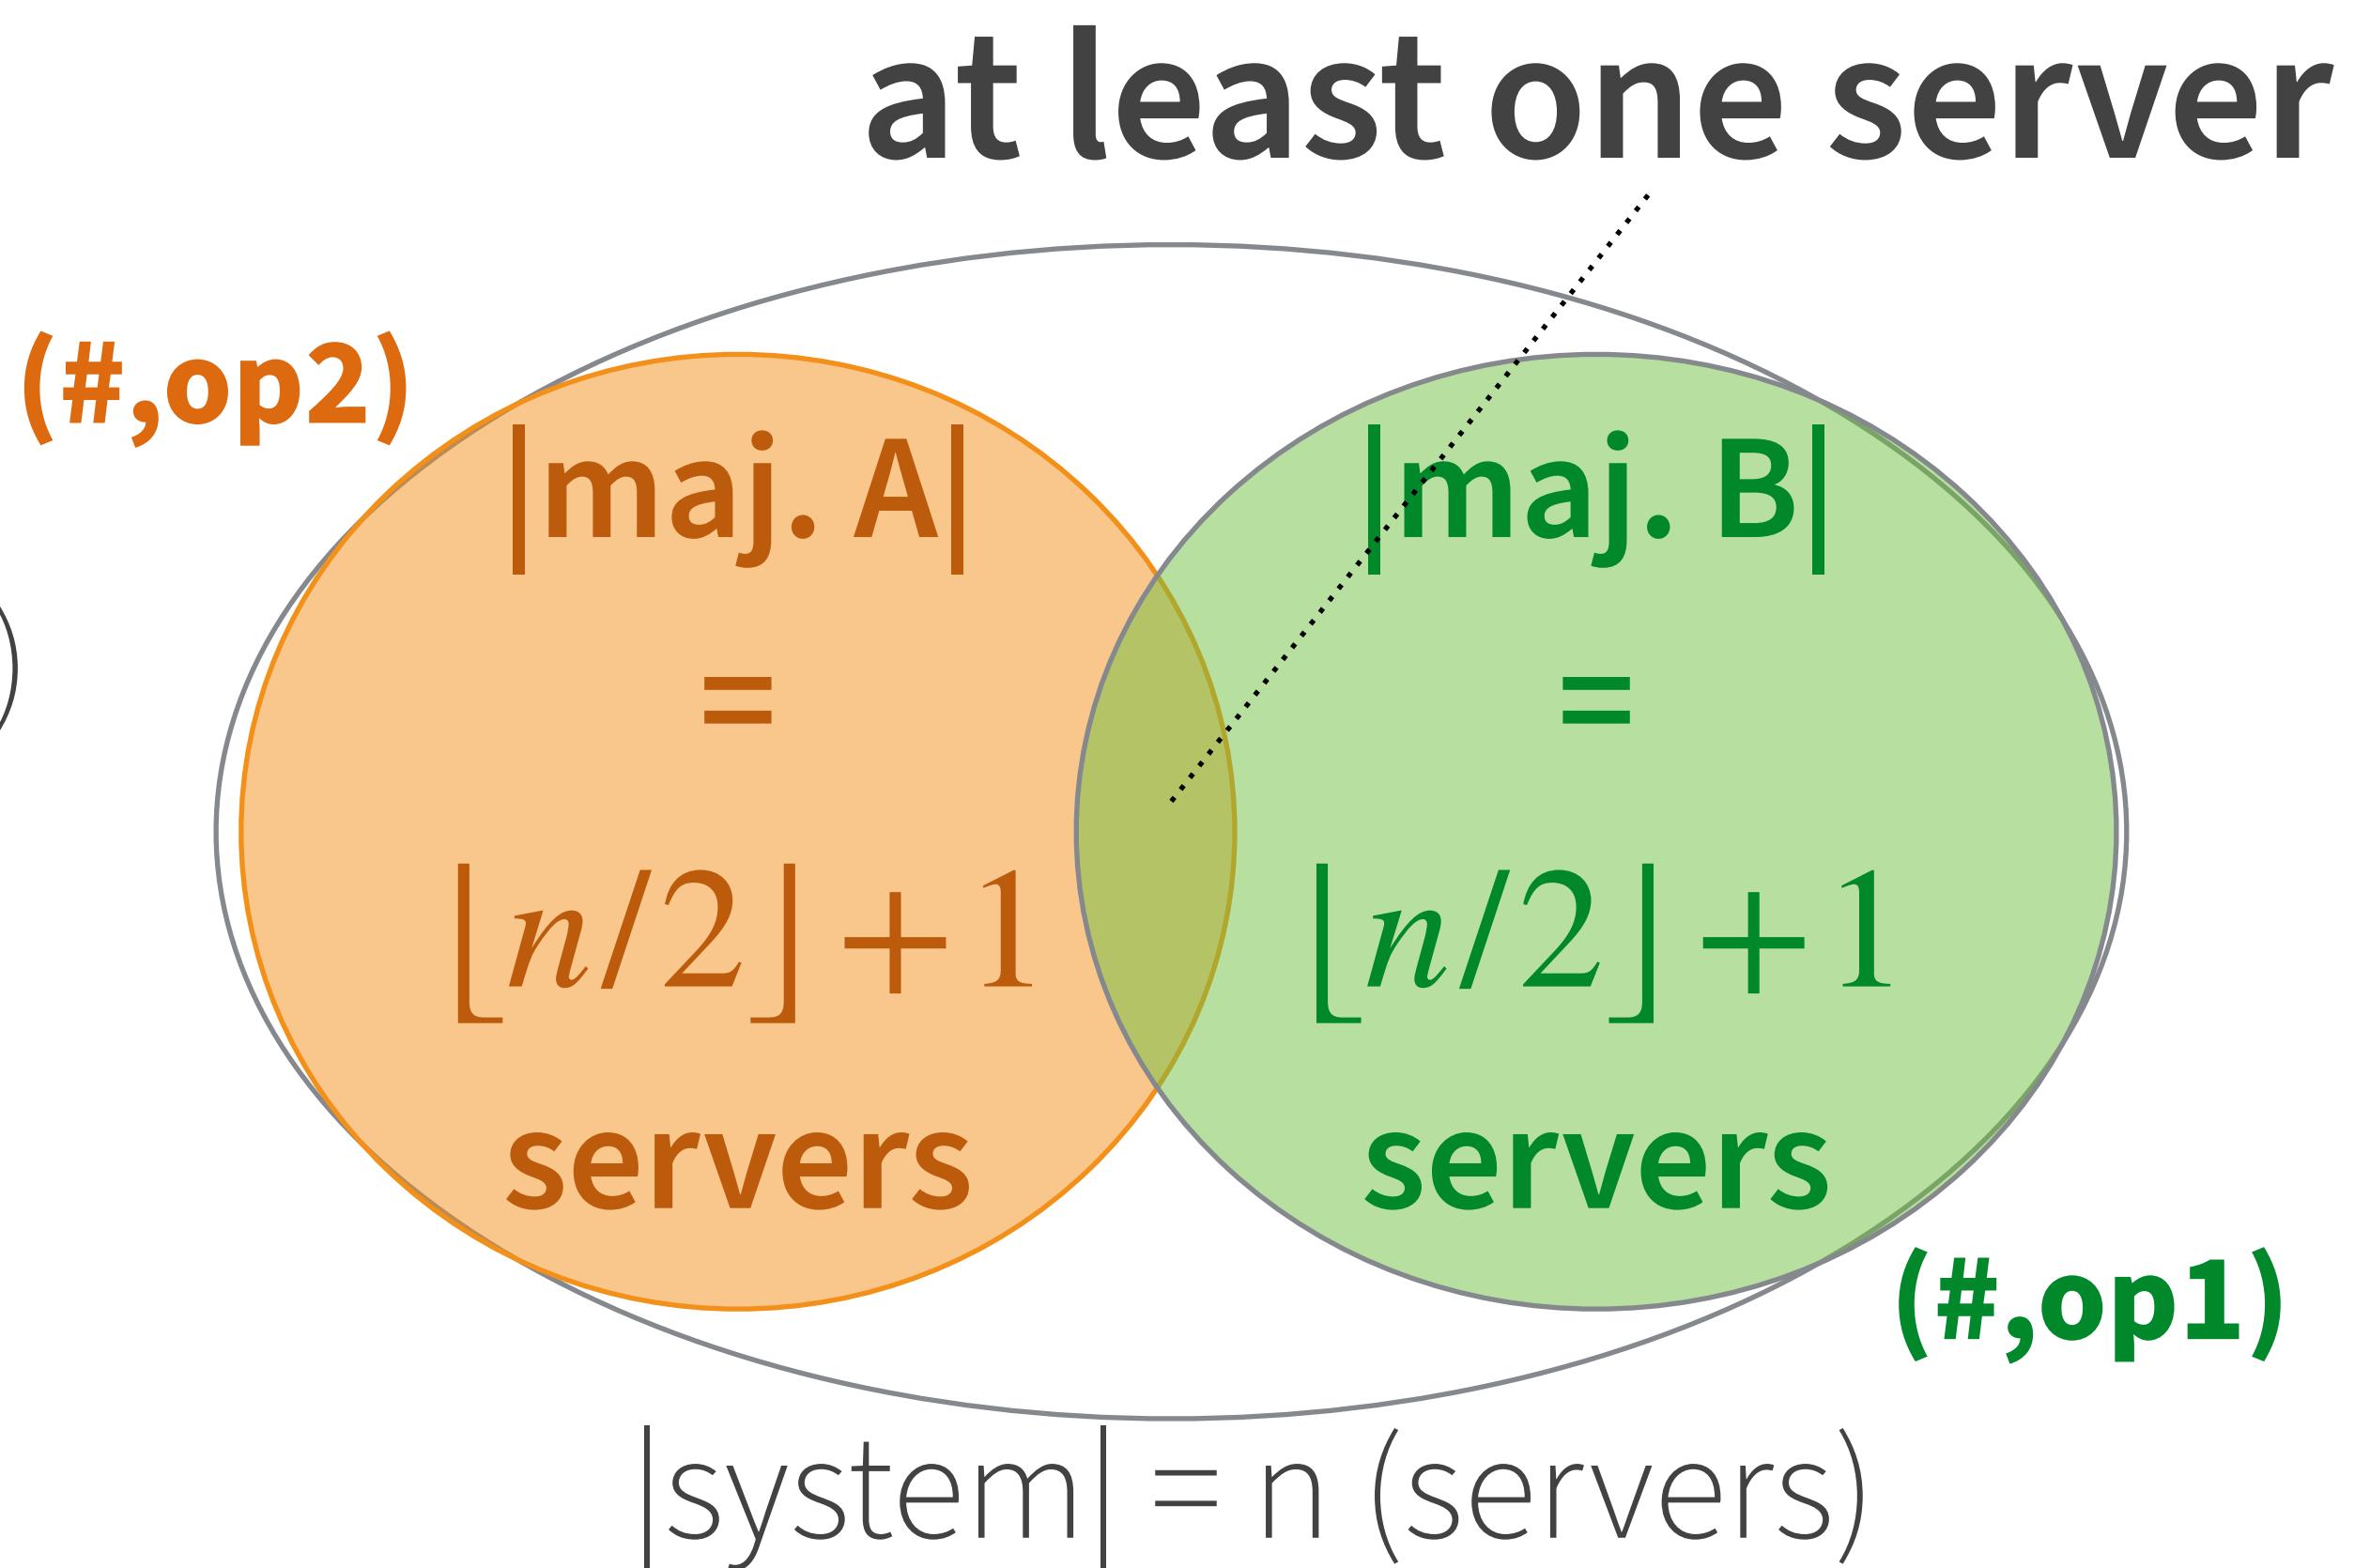
- ▶ n servers; f faulty (crash); $n \geq 2f+1$

▶ Recovery (take over)

- ▶ accept (#,op2) by majority (A)
- ▶ recovery with info. from majority (B)
- ▶ at least one process from A

▶ Non-blocking

- ▶ broadcast message
- ▶ wait for at most $n-f$ replies



Byzantine Fault-Tolerance

Data / System Integrity

► Requirements

- ▶ No tampering of data / system
- ▶ Distributed system: preserve **consistency**

► “Outside” Threat

- ▶ Attacker has **no secret** initially
- ▶ Attacker uses **public** service interface in unexpected ways

► “Inside” Threat

- ▶ Attacker **passes as legitimate** node / server (eludes detection)
- ▶ Attacker has access to **all information**

Byzantine Faults

▶ Fault

- ▶ behavior deviates from specification

▶ Fault Modes

▶ Crash

stops working permanently

▶ Omission

misses sending/receiving of messages

▶ Rational (selfish)

(mis)behavior that maximizes local benefit

▶ Byzantine (accidental / malicious)

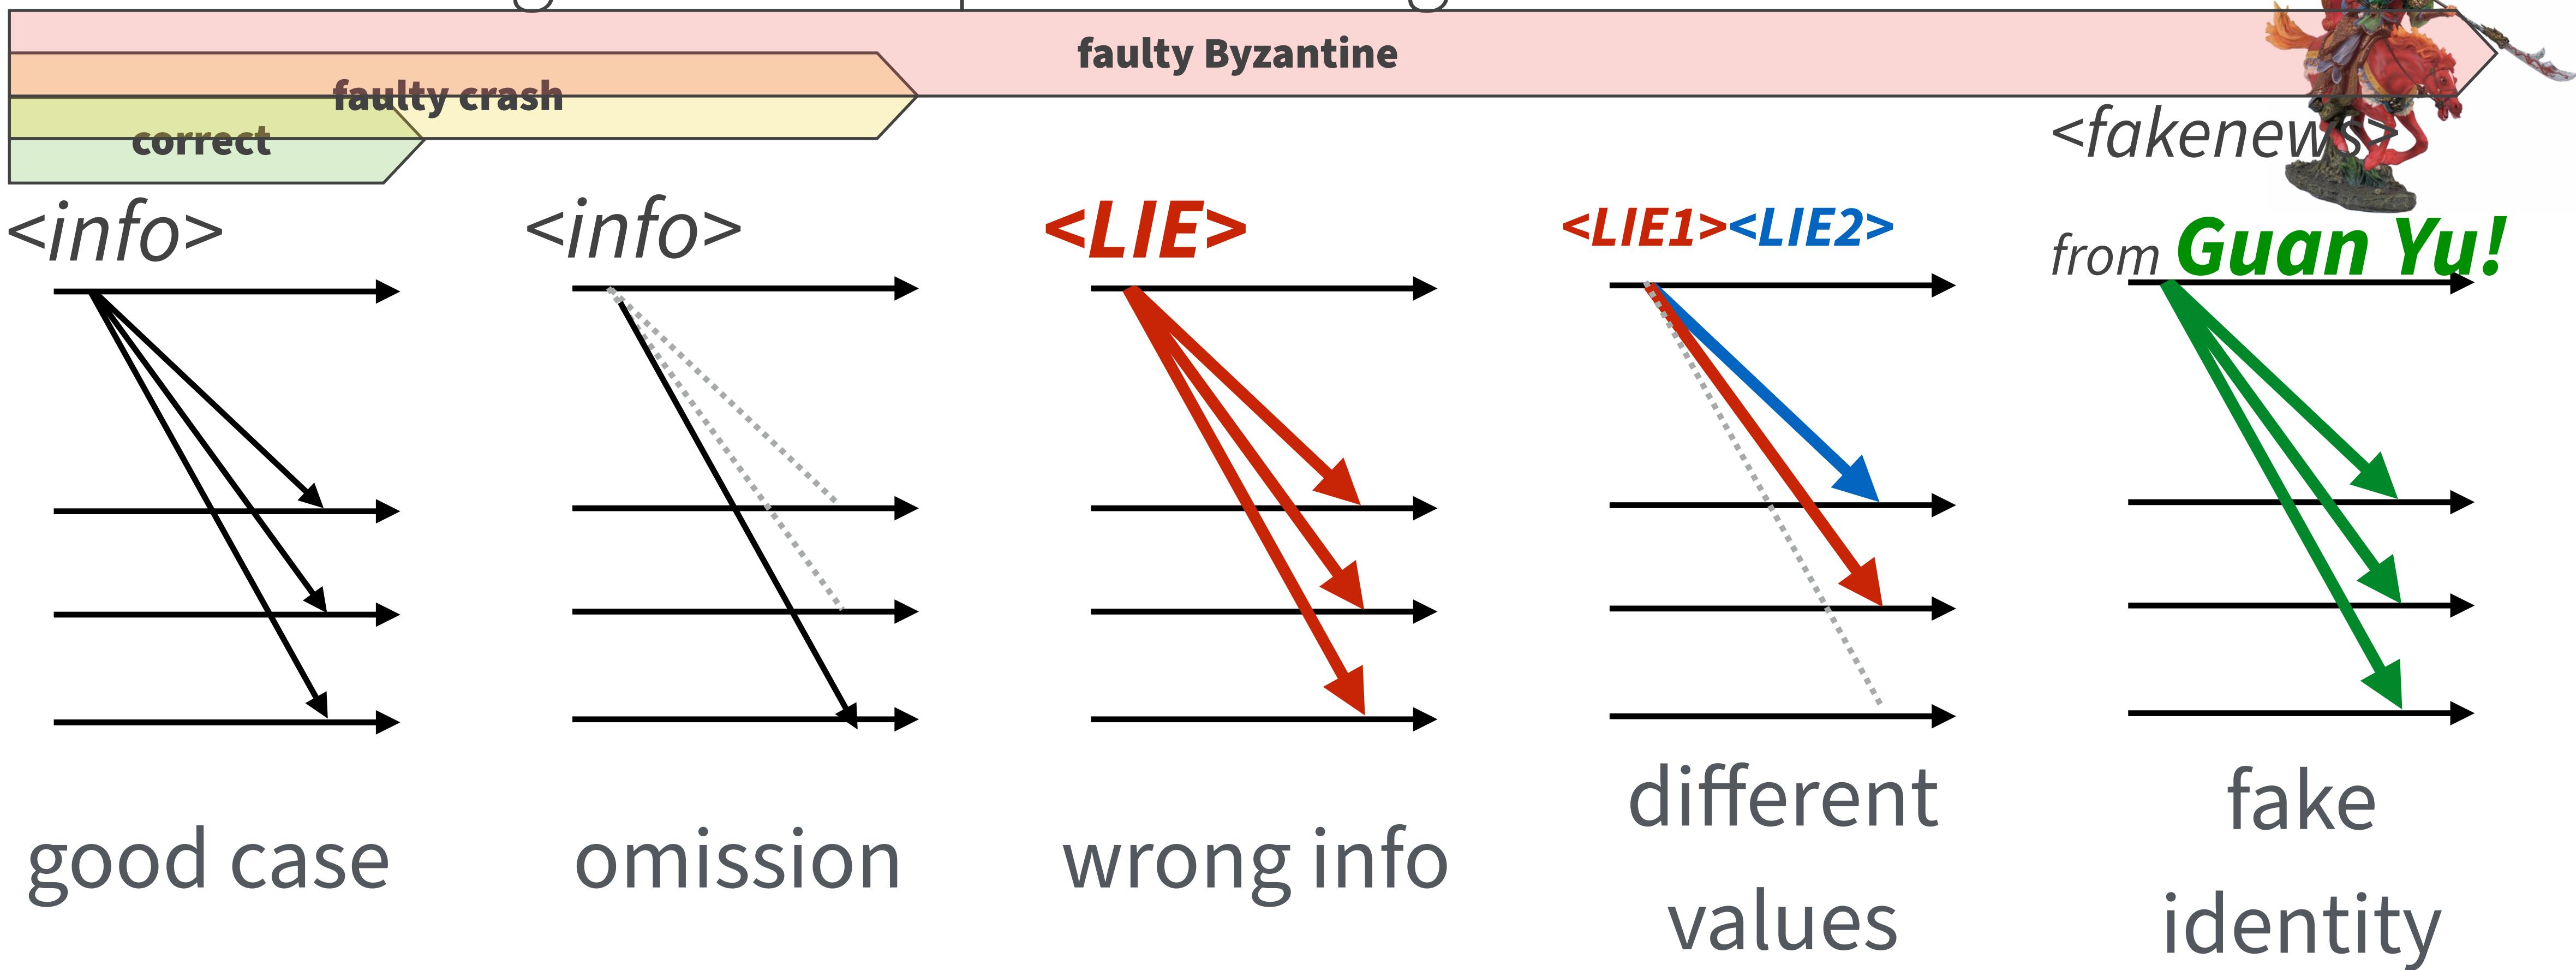
any type of misbehavior



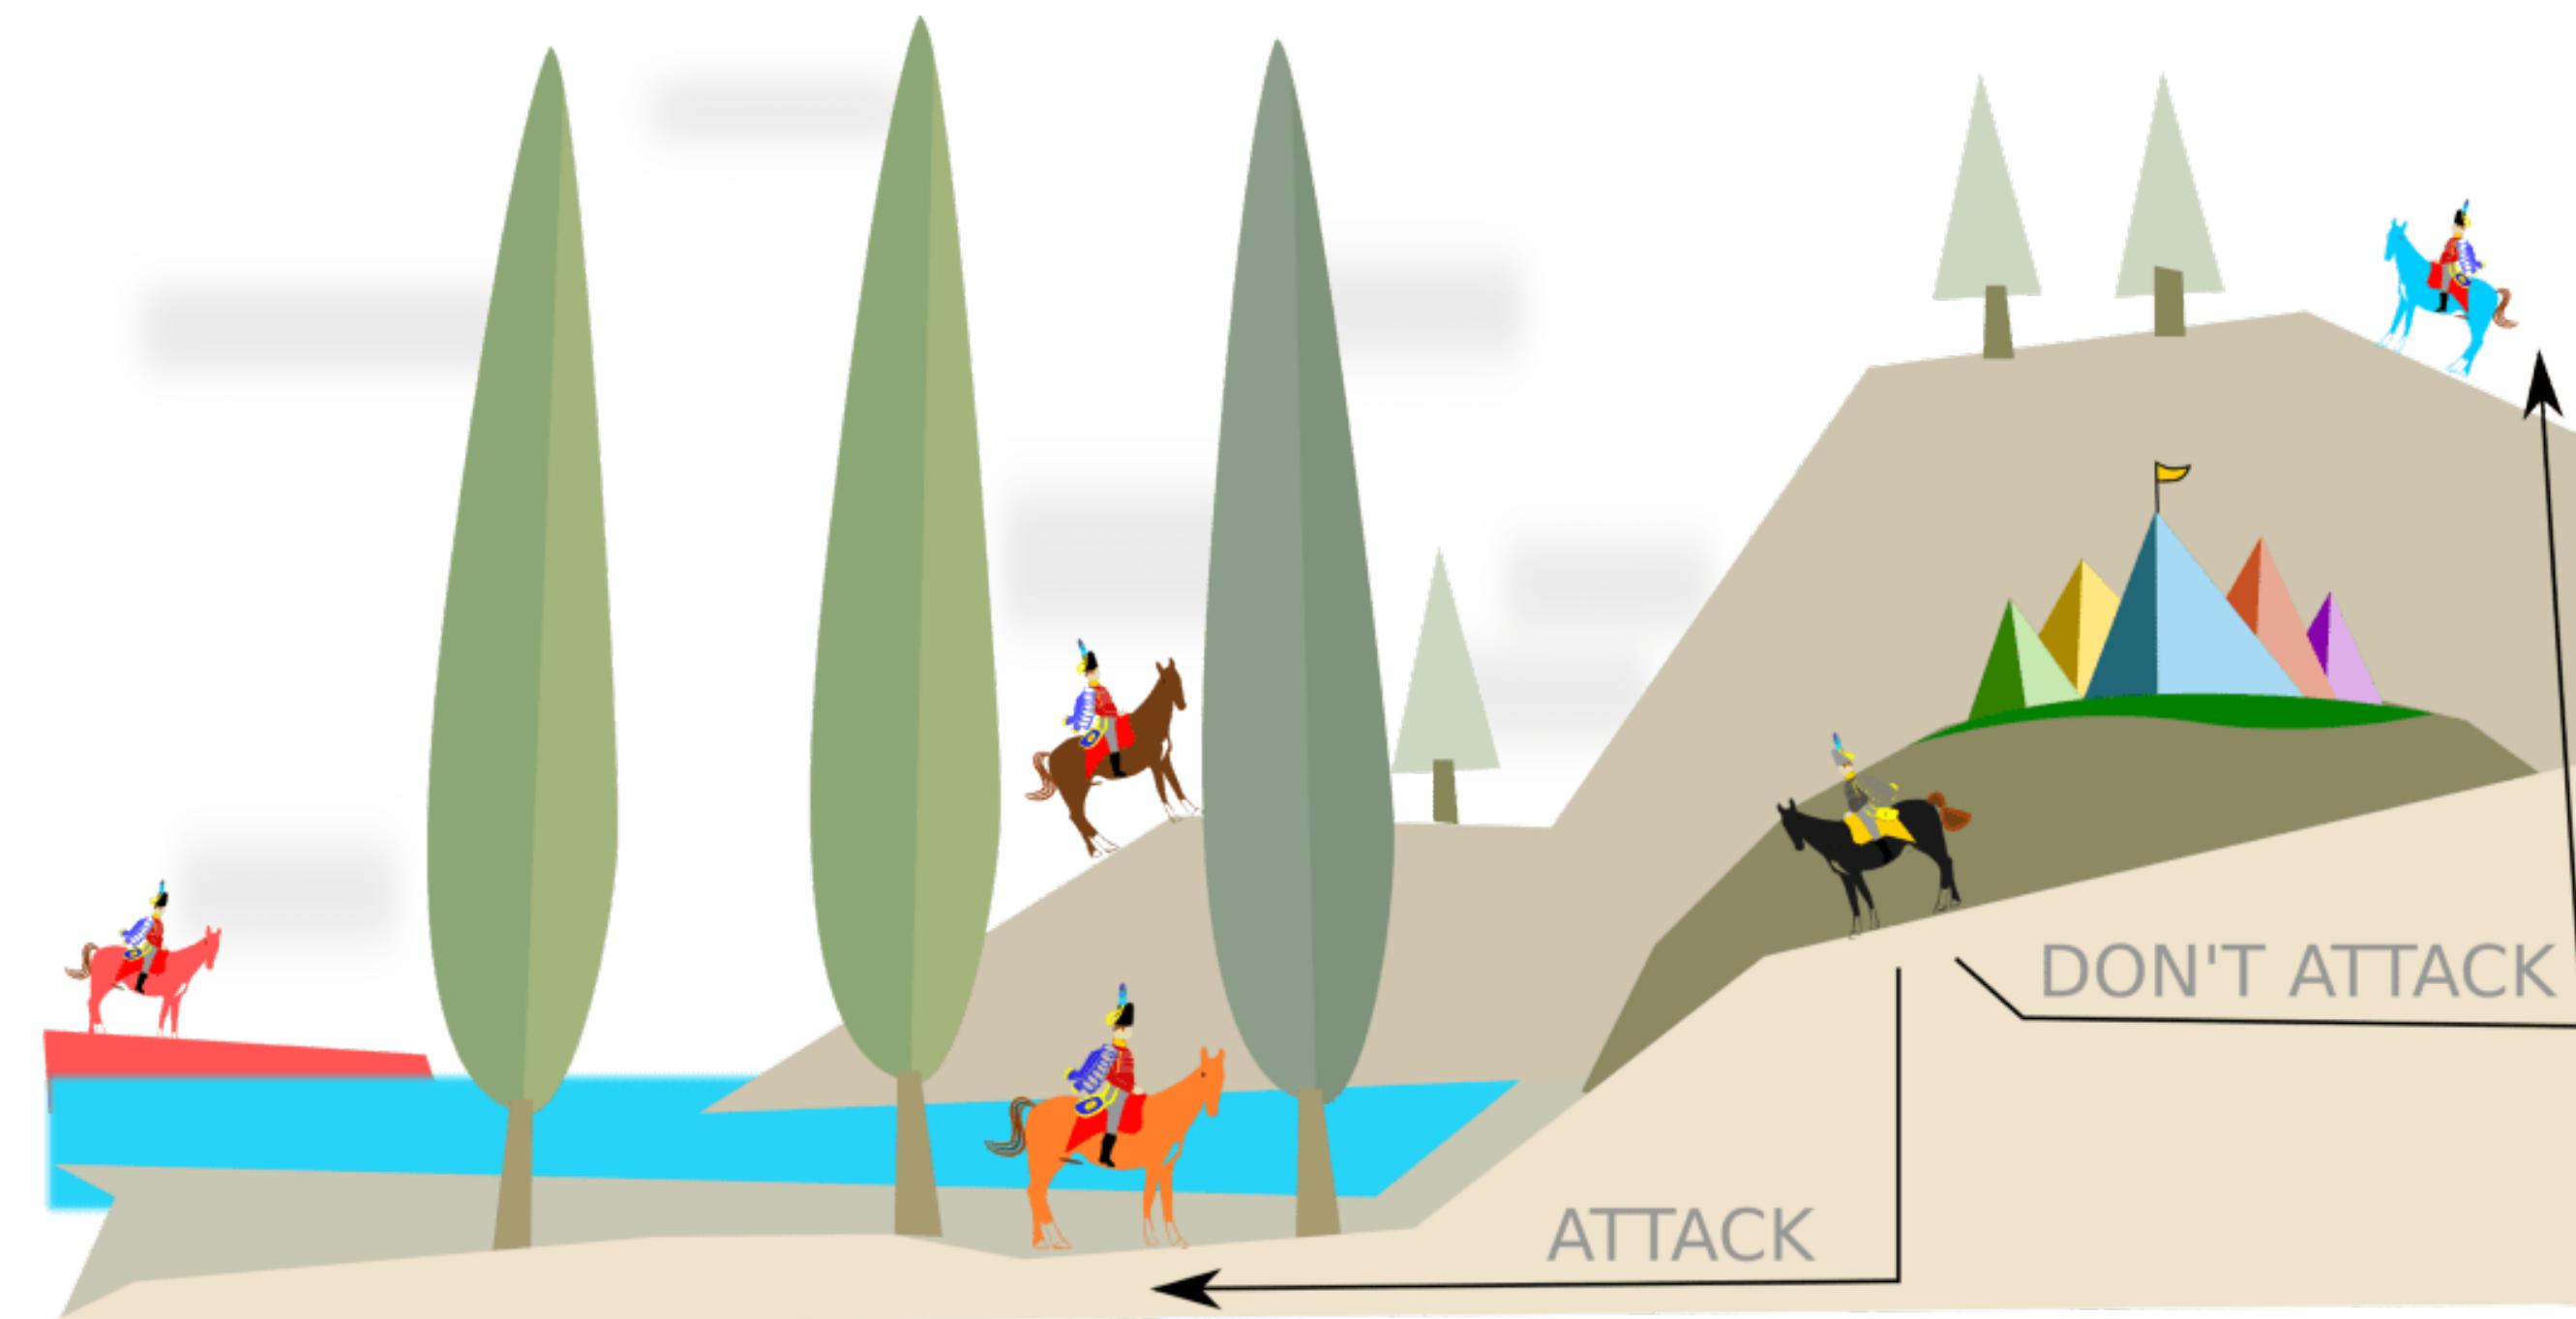
Byzantine Faults

▶ Concretely

- ▶ broadcasting **<info>** as part of an algorithm



Byzantine Generals

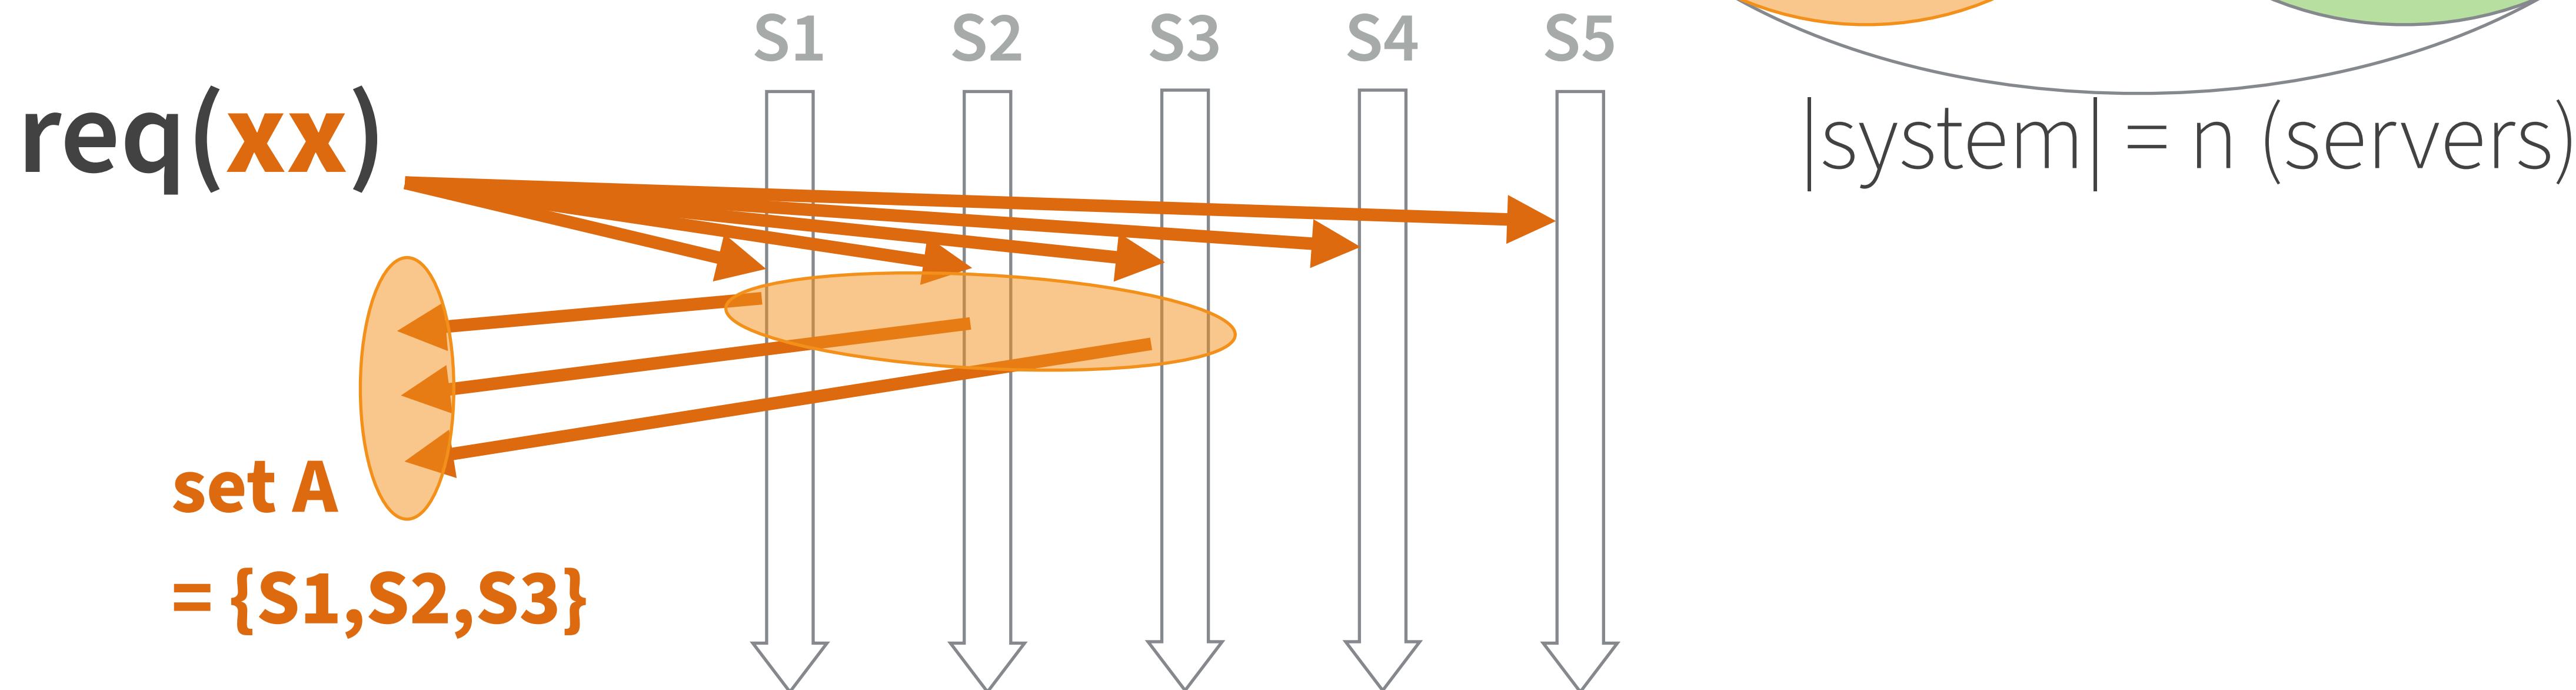


[LSP82] Lamport, Shostak, Pease: **The Byzantine Generals Problem.** *ACM Trans. Program. Lang. Syst.* 4(3): 382-401 (1982)

Byzantine Quorum

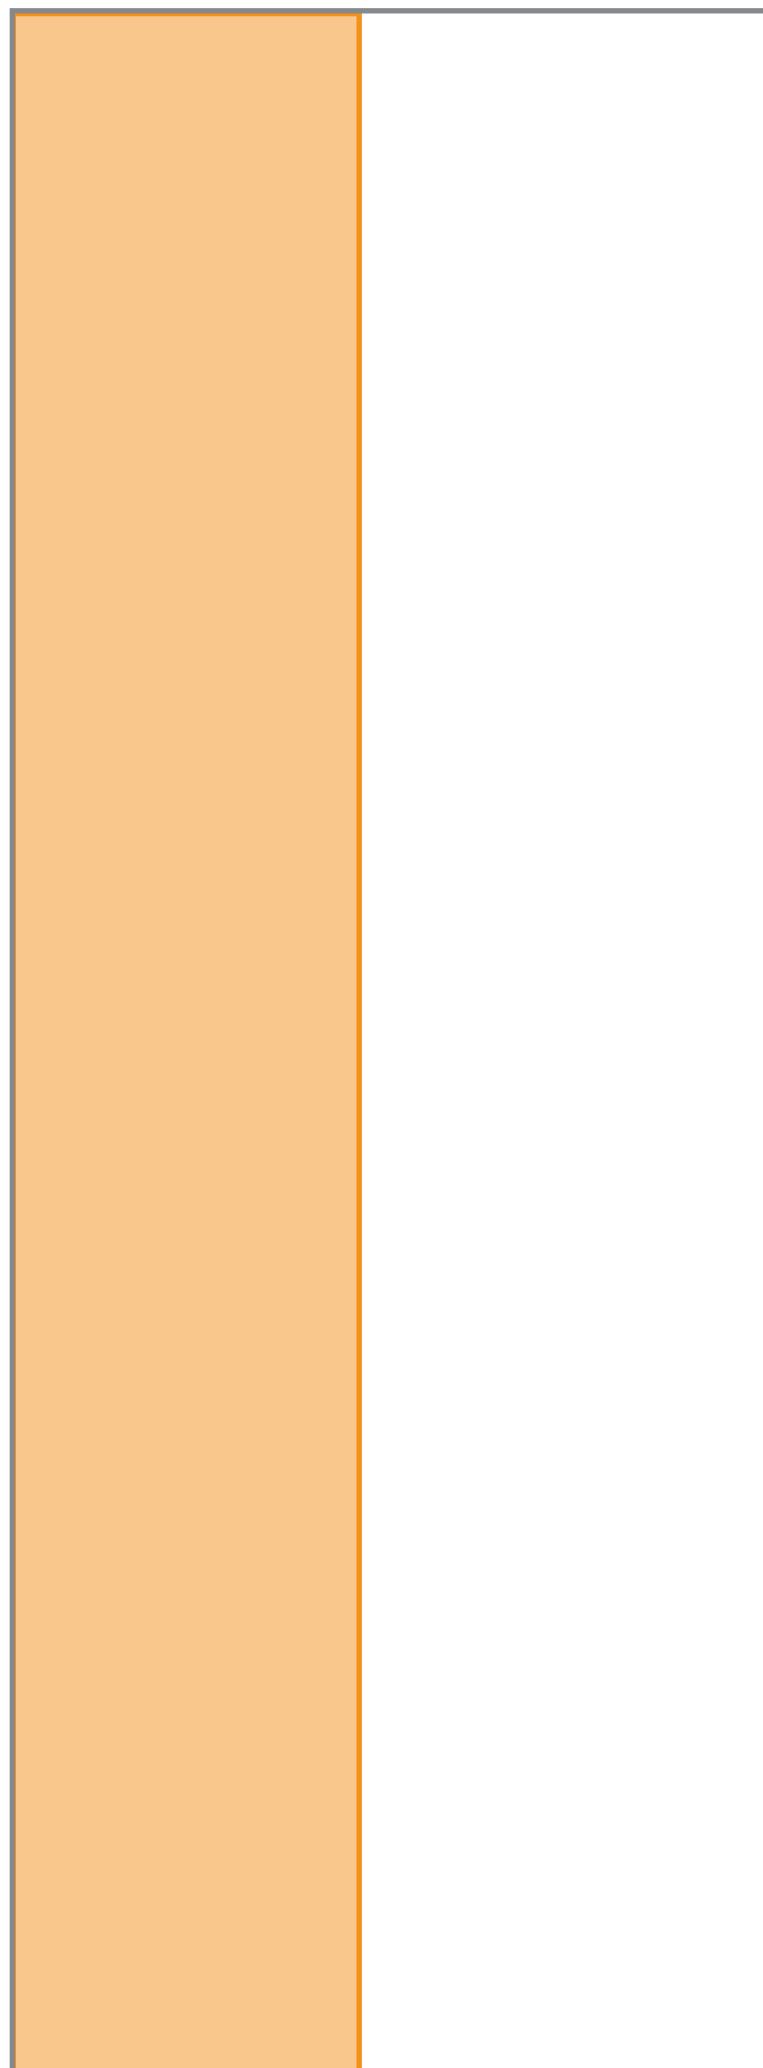
▶ Terminology

- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults



Byzantine Quorum

set A



$|\text{system}| = n$ (servers)

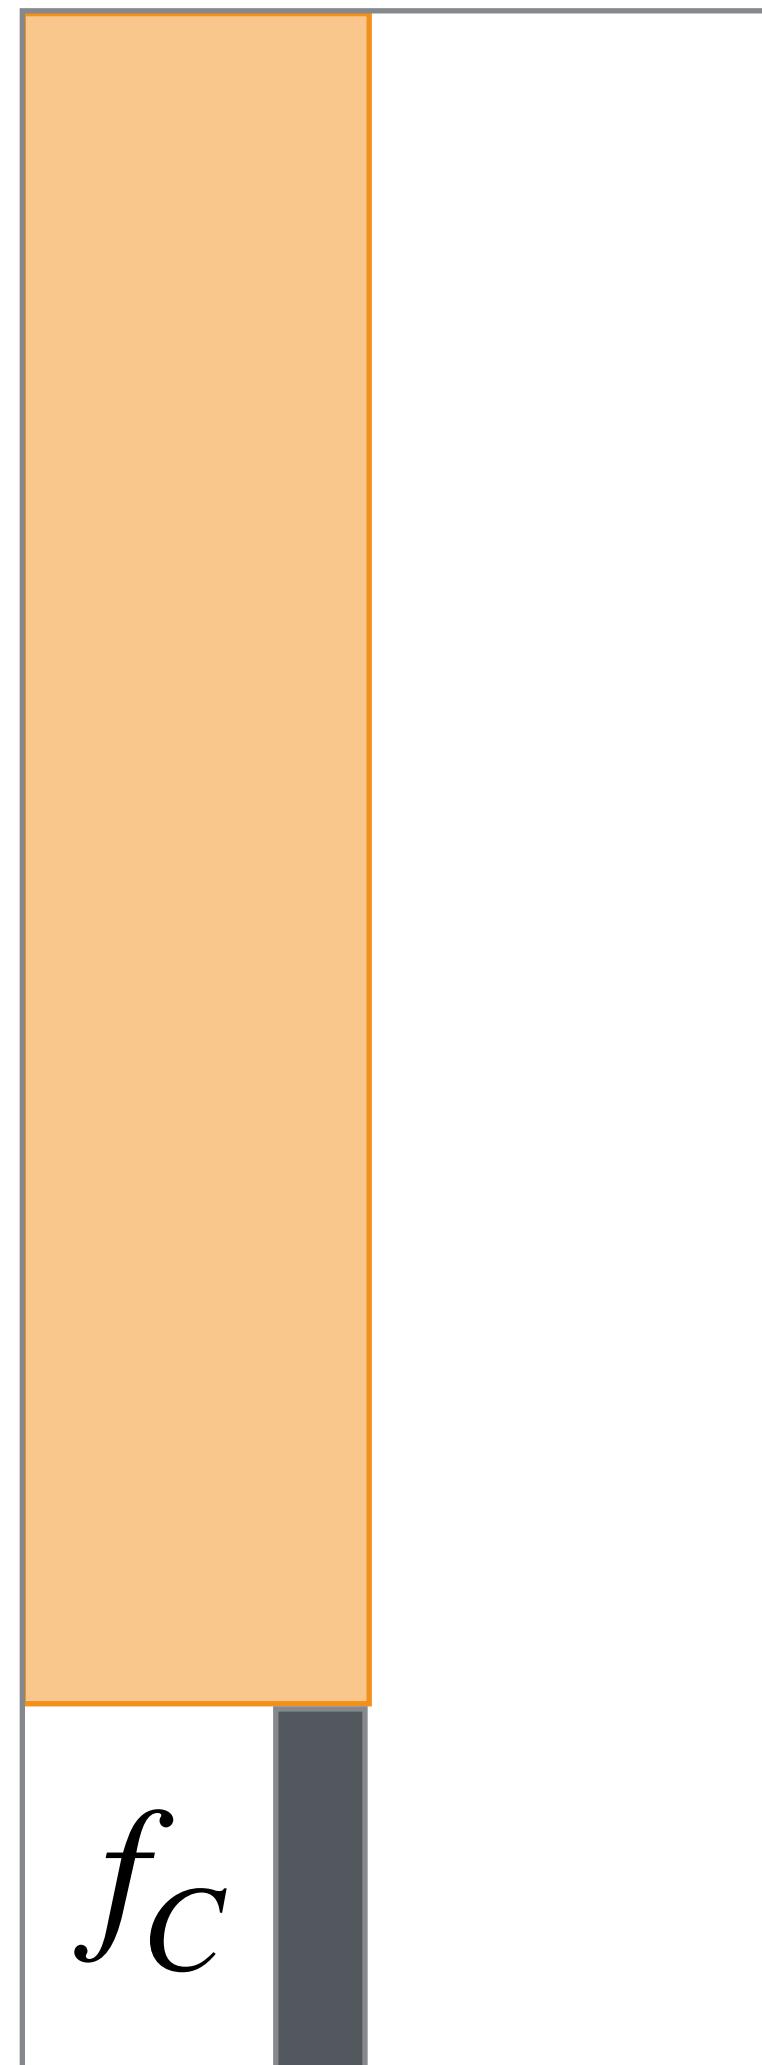
▶ Terminology

- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

▶ Non-blocking

Byzantine Quorum

set A



$|system| = n$ (servers)

▶ Terminology

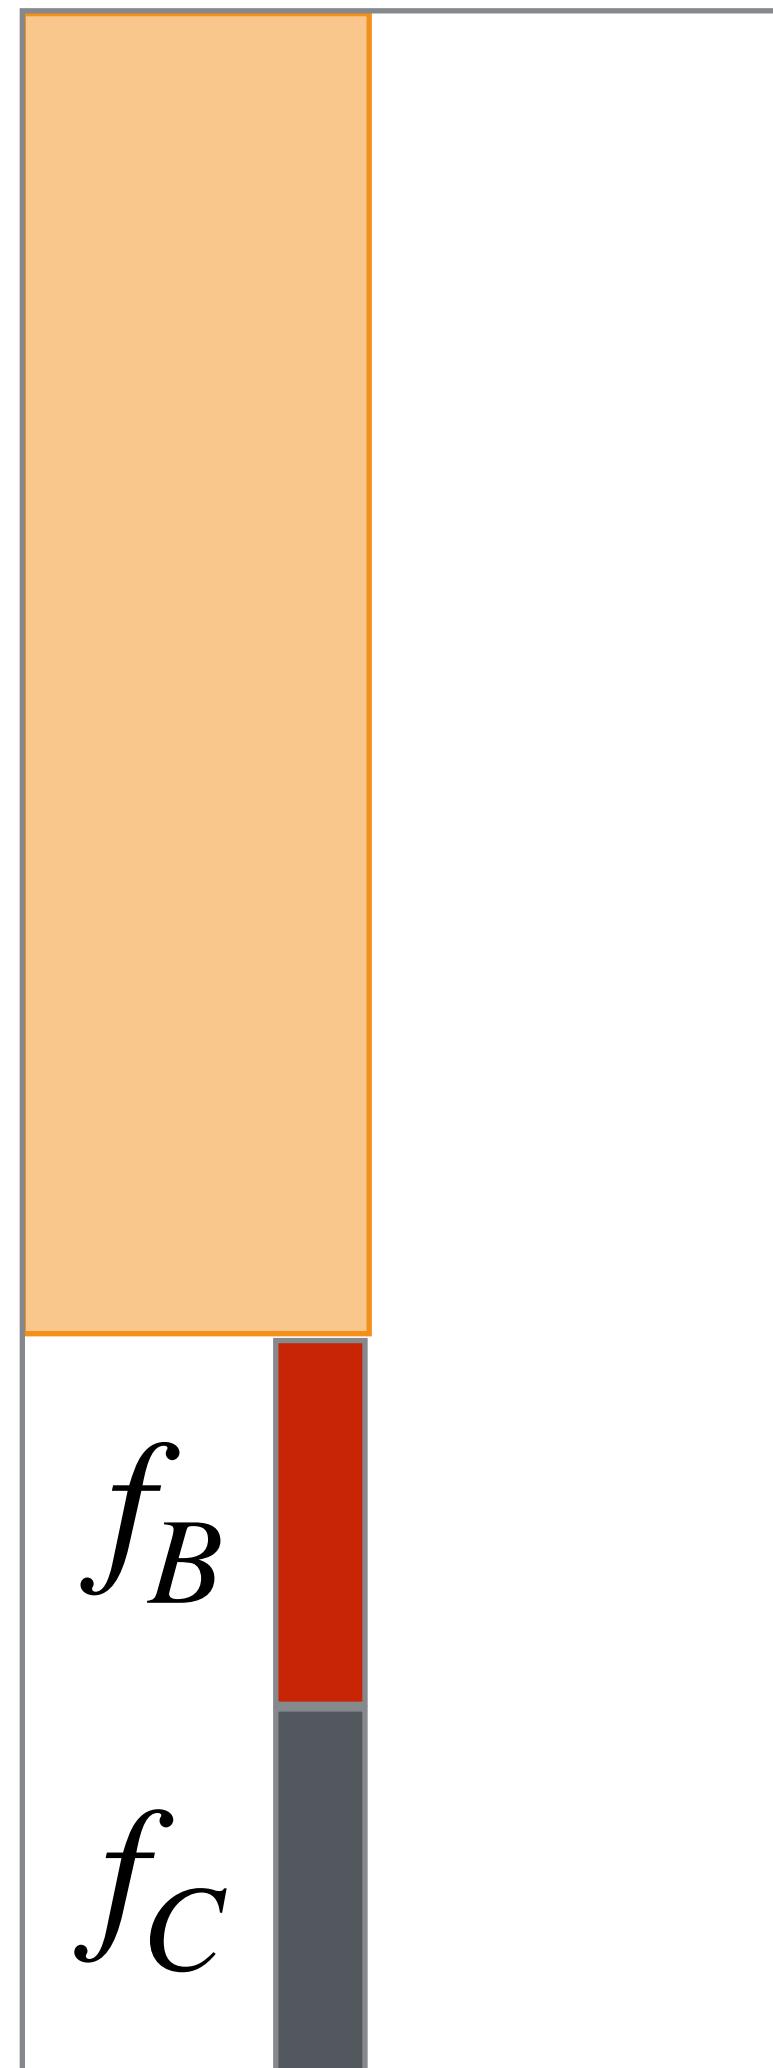
- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

▶ Non-blocking

- ▶ f_C have crashed

Byzantine Quorum

set A



$|\text{system}| = n$ (servers)

▶ Terminology

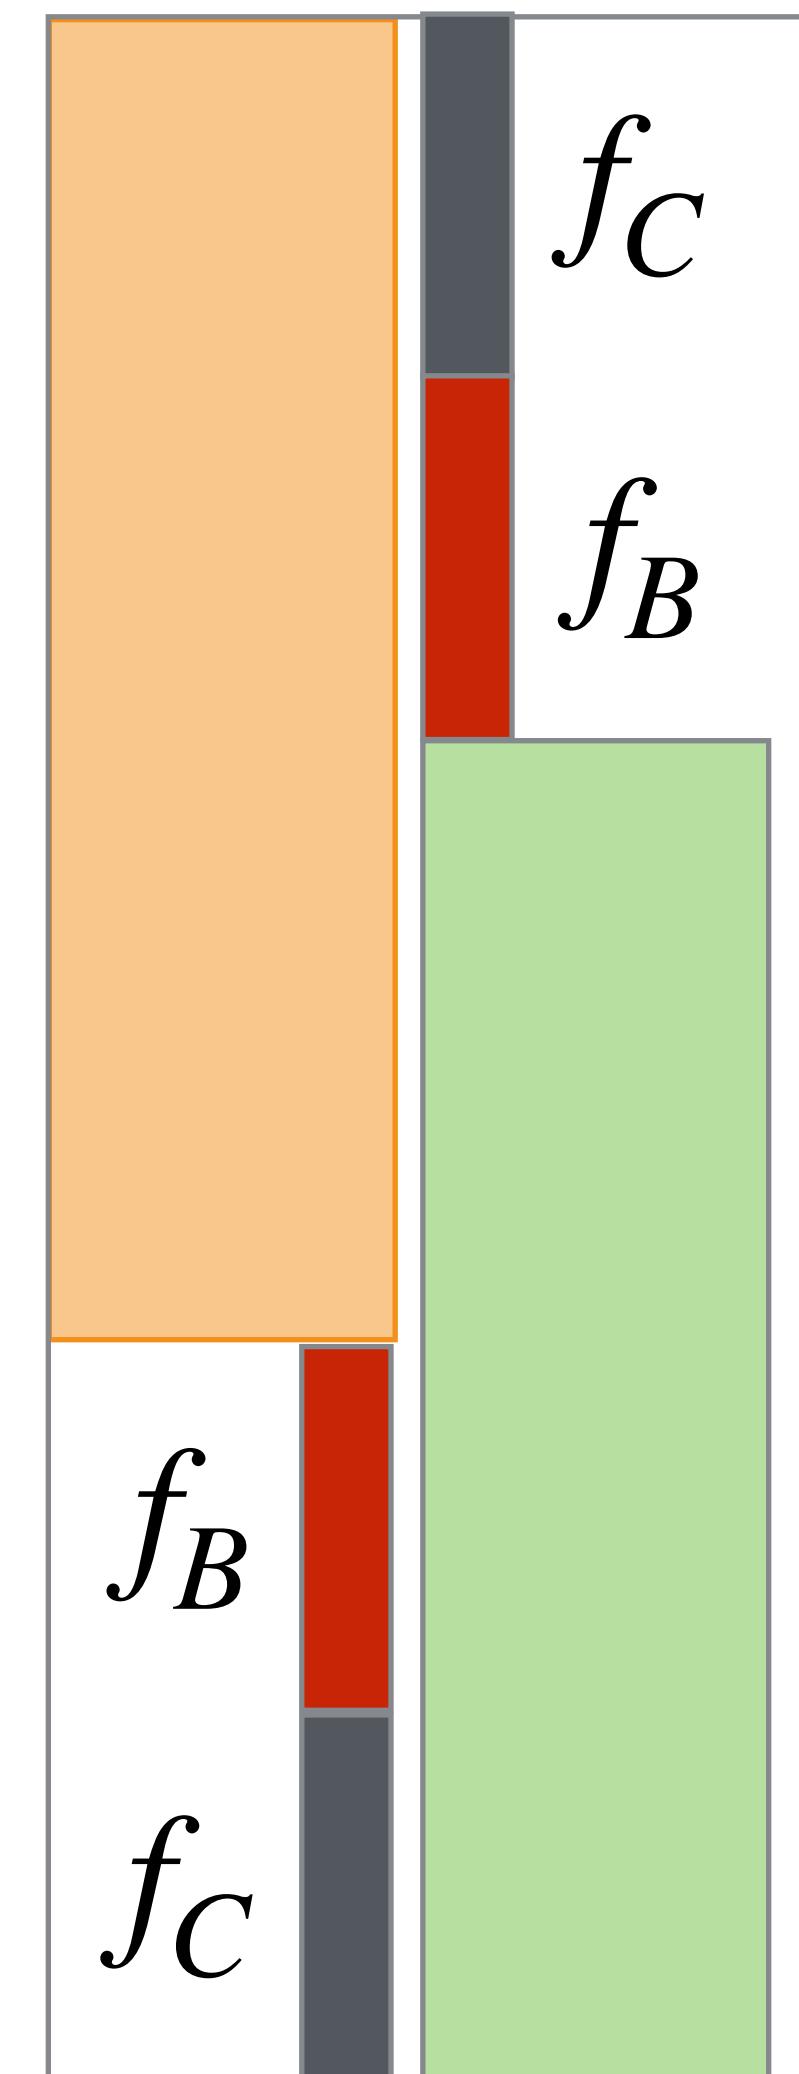
- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

▶ Non-blocking

- ▶ f_C have crashed
- ▶ f_B never reply

Byzantine Quorum

set A



$|\text{system}| = n$ (servers)

set B

$n-f$

f_B

f_C

f_B

f_C

f_B

Terminology

- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

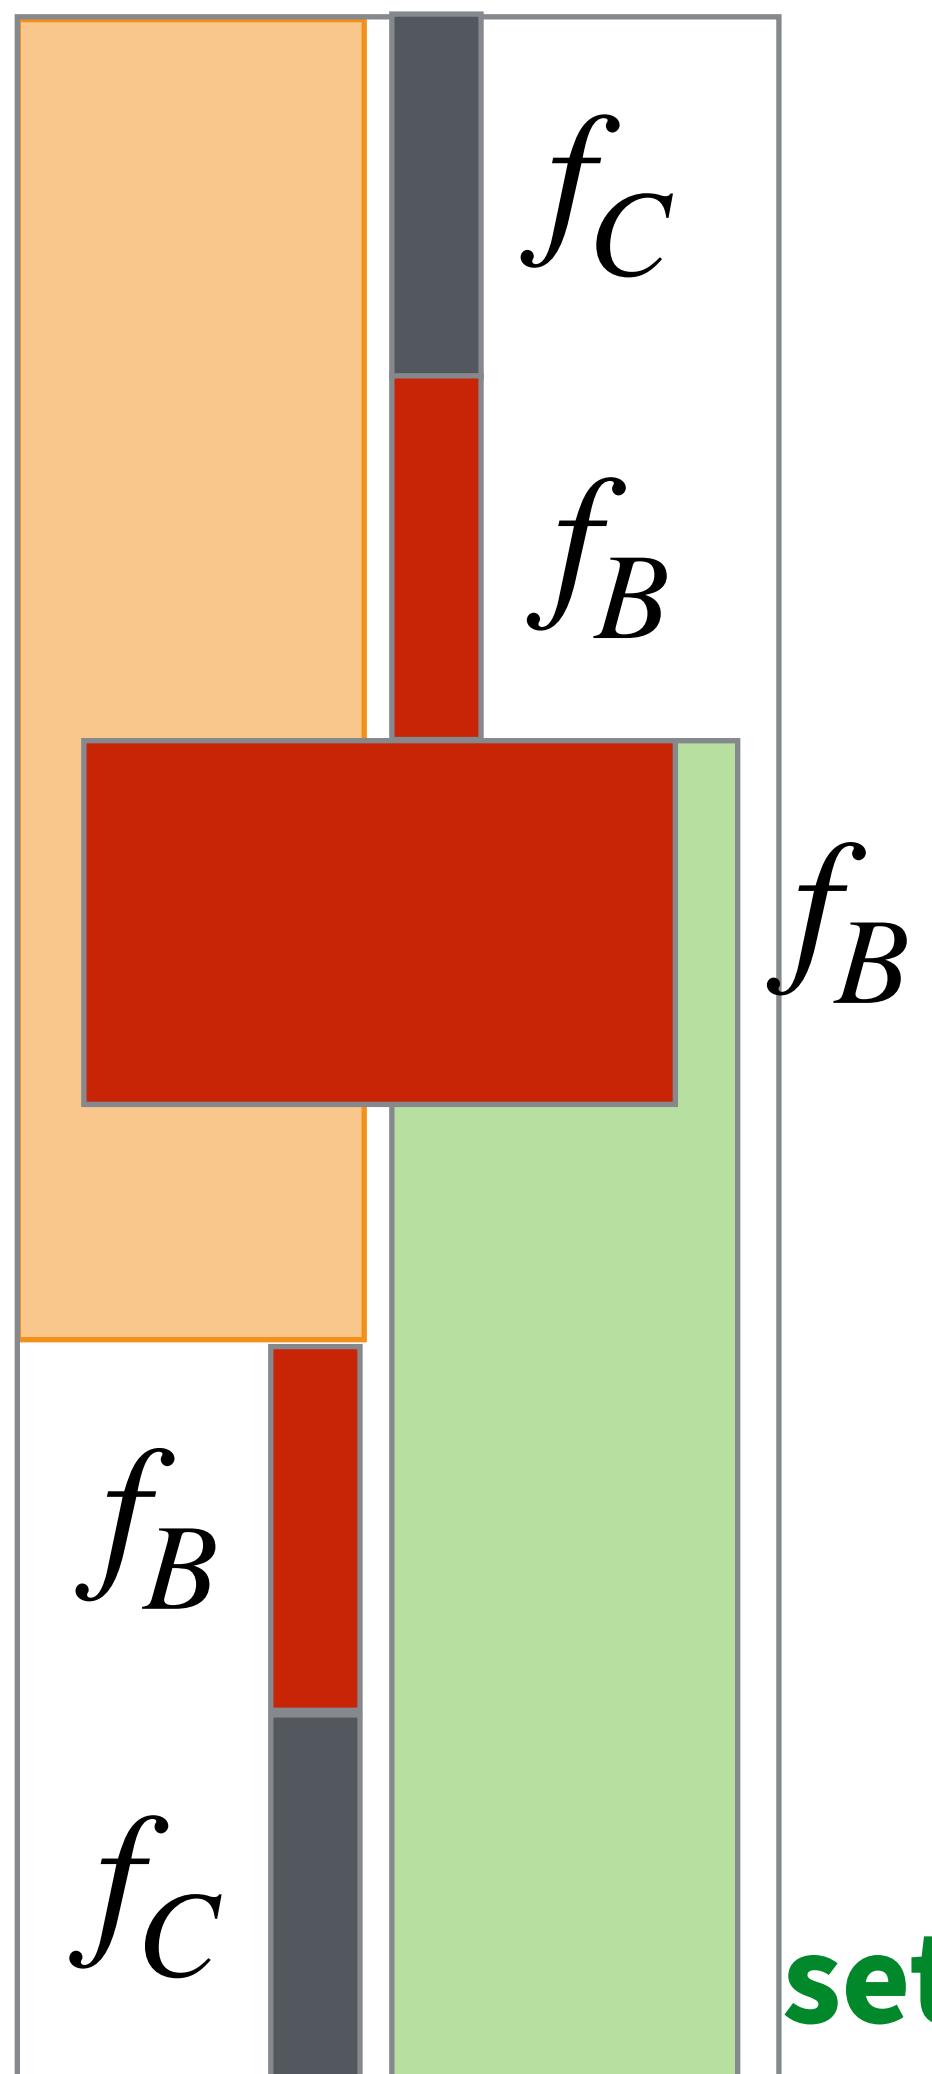
Non-blocking

- ▶ f_C have crashed
- ▶ f_B never reply

Intersection

Byzantine Quorum

set A



$|\text{system}| = n$ (servers)

▶ Terminology

- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

▶ Non-blocking

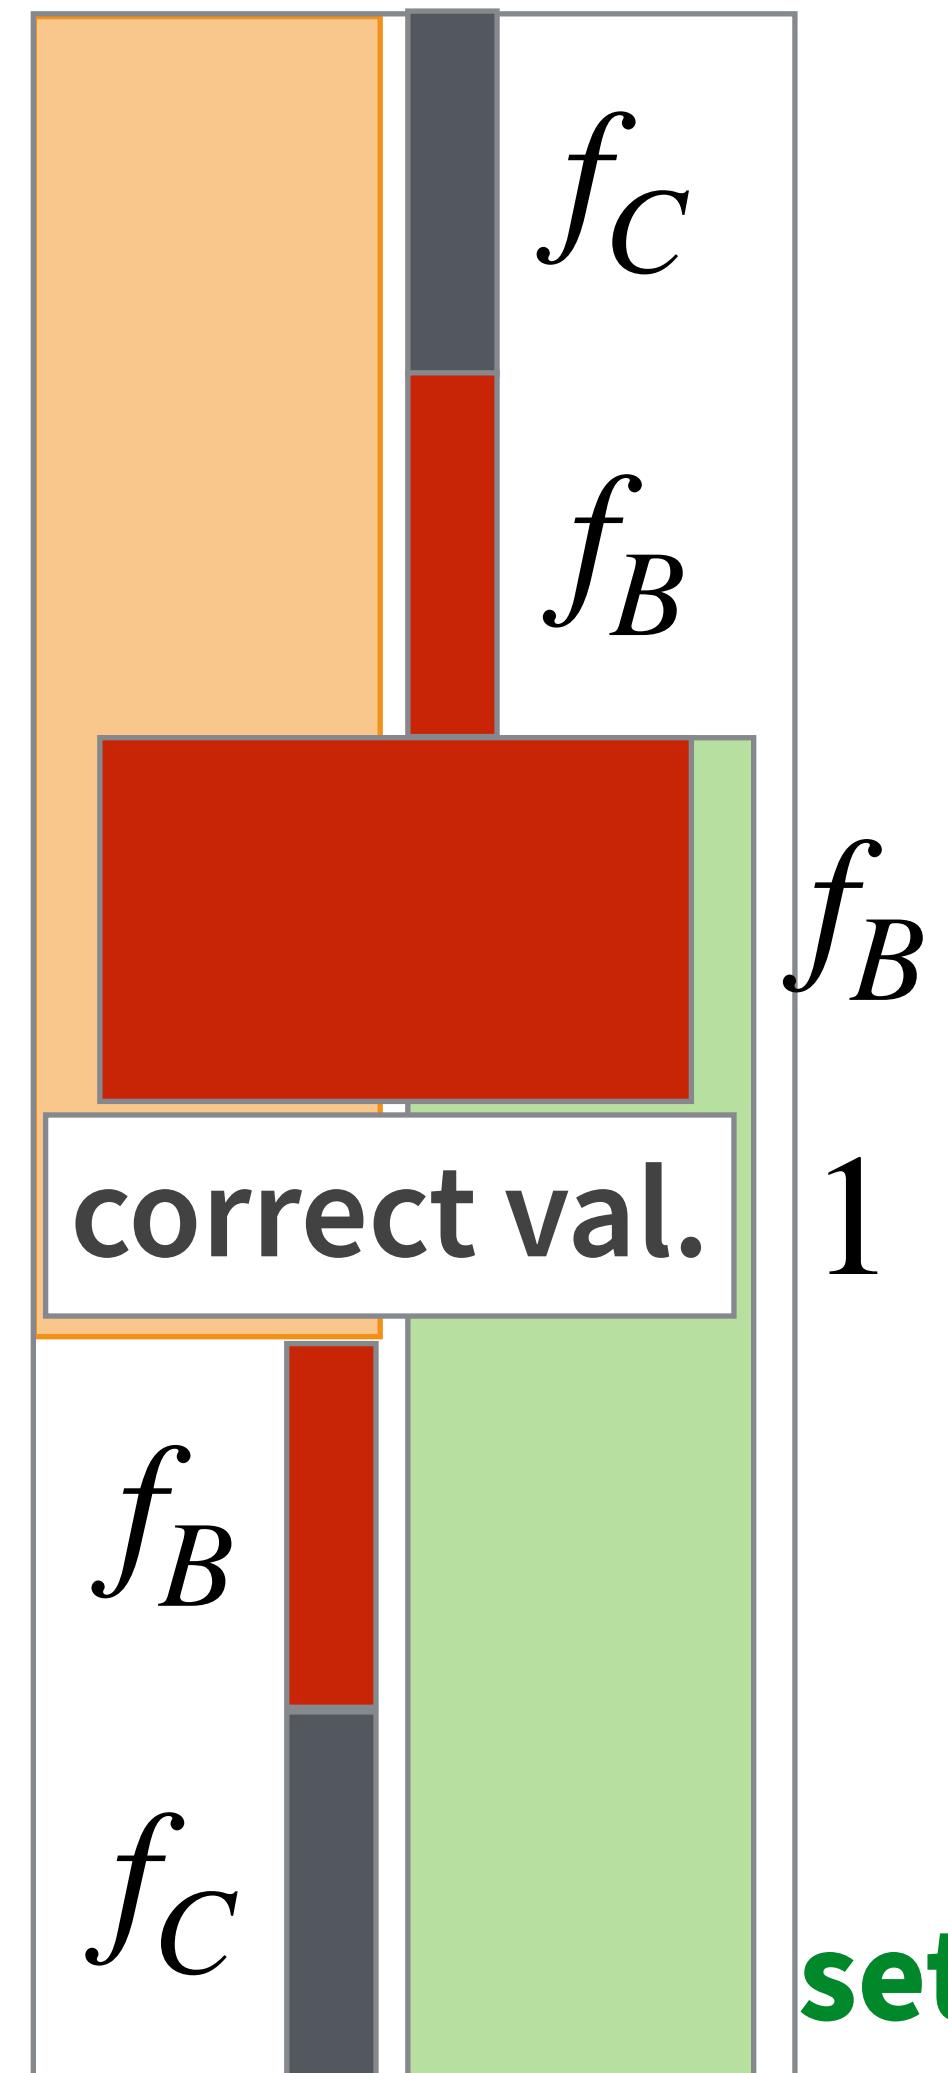
- ▶ f_C have crashed
- ▶ f_B never reply

▶ Intersection

- ▶ includes lies

Byzantine Quorum

set A



$|\text{system}| = n$ (servers)

▶ Terminology

- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

▶ Non-blocking

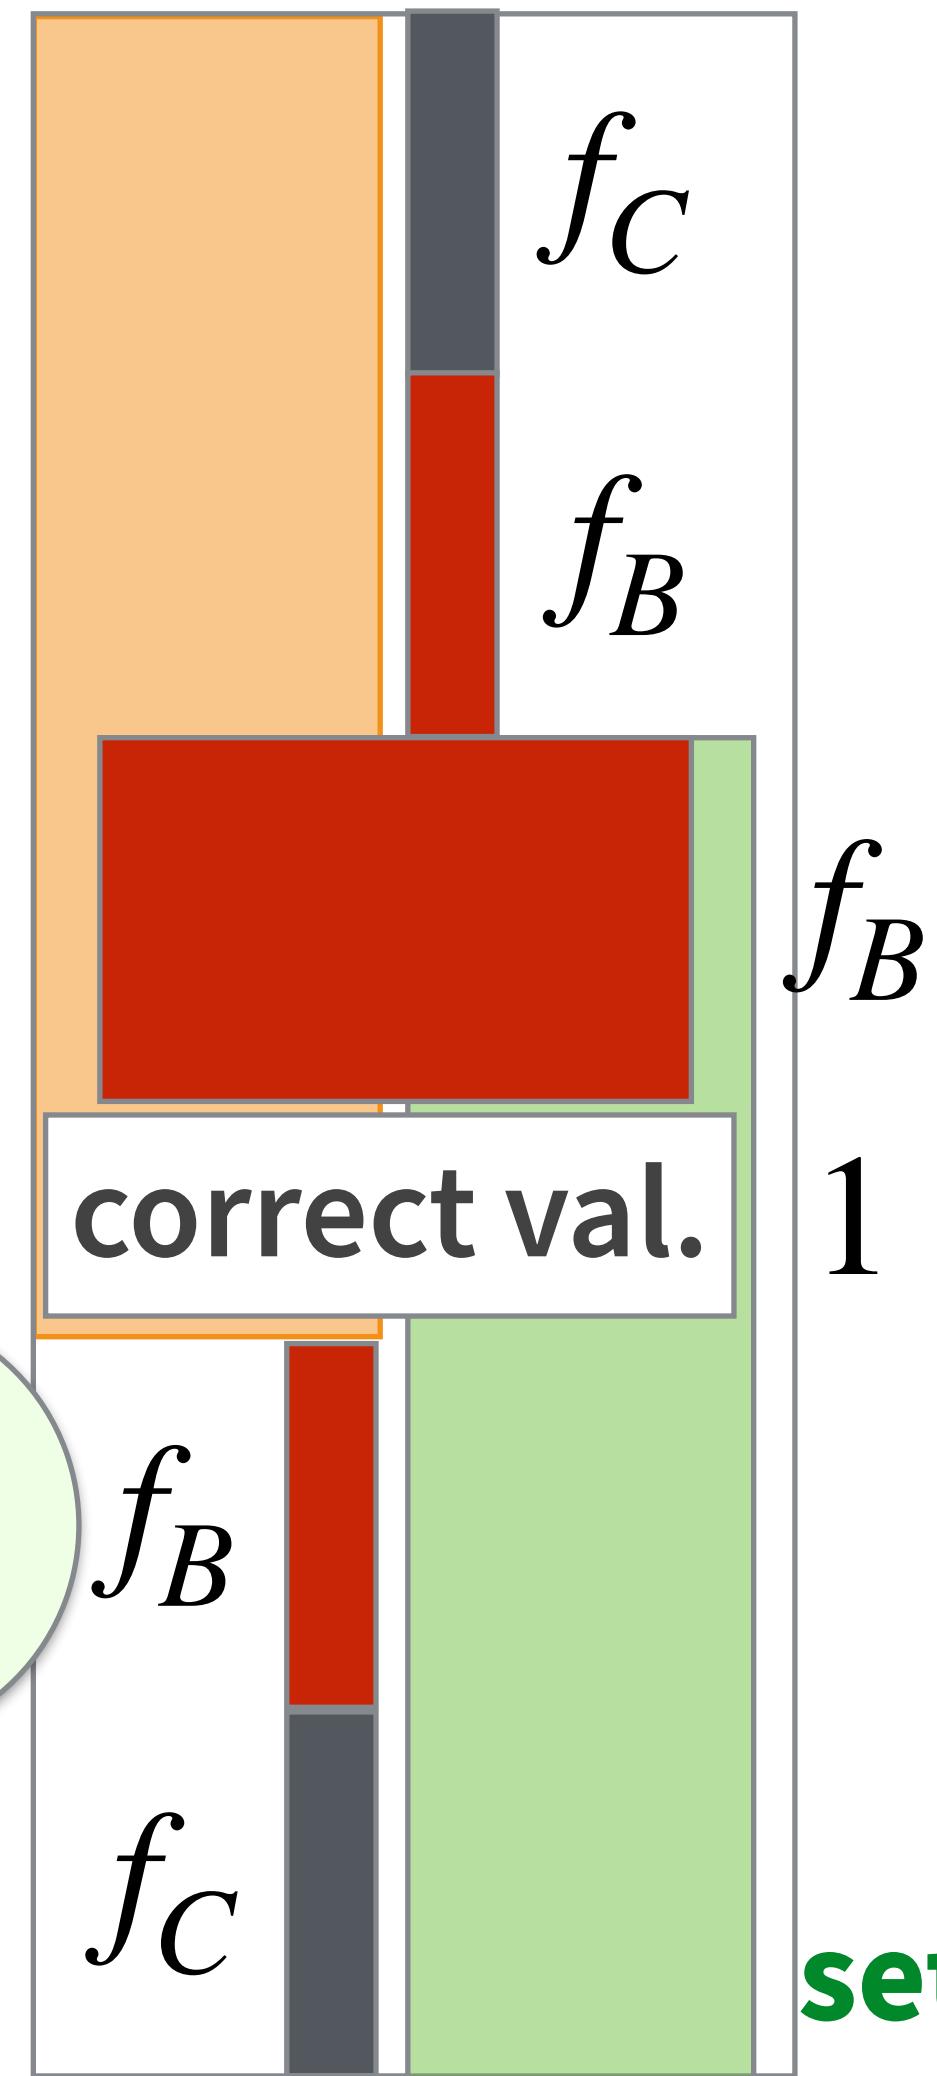
- ▶ f_C have crashed
- ▶ f_B never reply

▶ Intersection

- ▶ includes lies; at least one correct value

Byzantine Quorum

set A



$|\text{system}| = n$ (servers)

Terminology

- ▶ n servers; $f = f_C + f_B$
- ▶ f_C number of **crash** faults
- ▶ f_B number of **Byzantine** faults

Non-blocking

- ▶ f_C have crashed
- ▶ f_B never reply

$$n \geq 2f_C + 3f_B + 1$$

Intersection

- ▶ includes lies; at least one correct value

Practical Byzantine Fault Tolerance

▶ Context

- ▶ described in client/server model
- ▶ state-machine replication
- ▶ similarities with Paxos
- ▶ Byzantine-tolerant
- ▶ uses crypto hash function

[CL01] Castro, Liskov: **Practical Byzantine Fault-Tolerance and Proactive Recovery.** *ACM Trans. Comput. Syst.* 20(4): 398-461

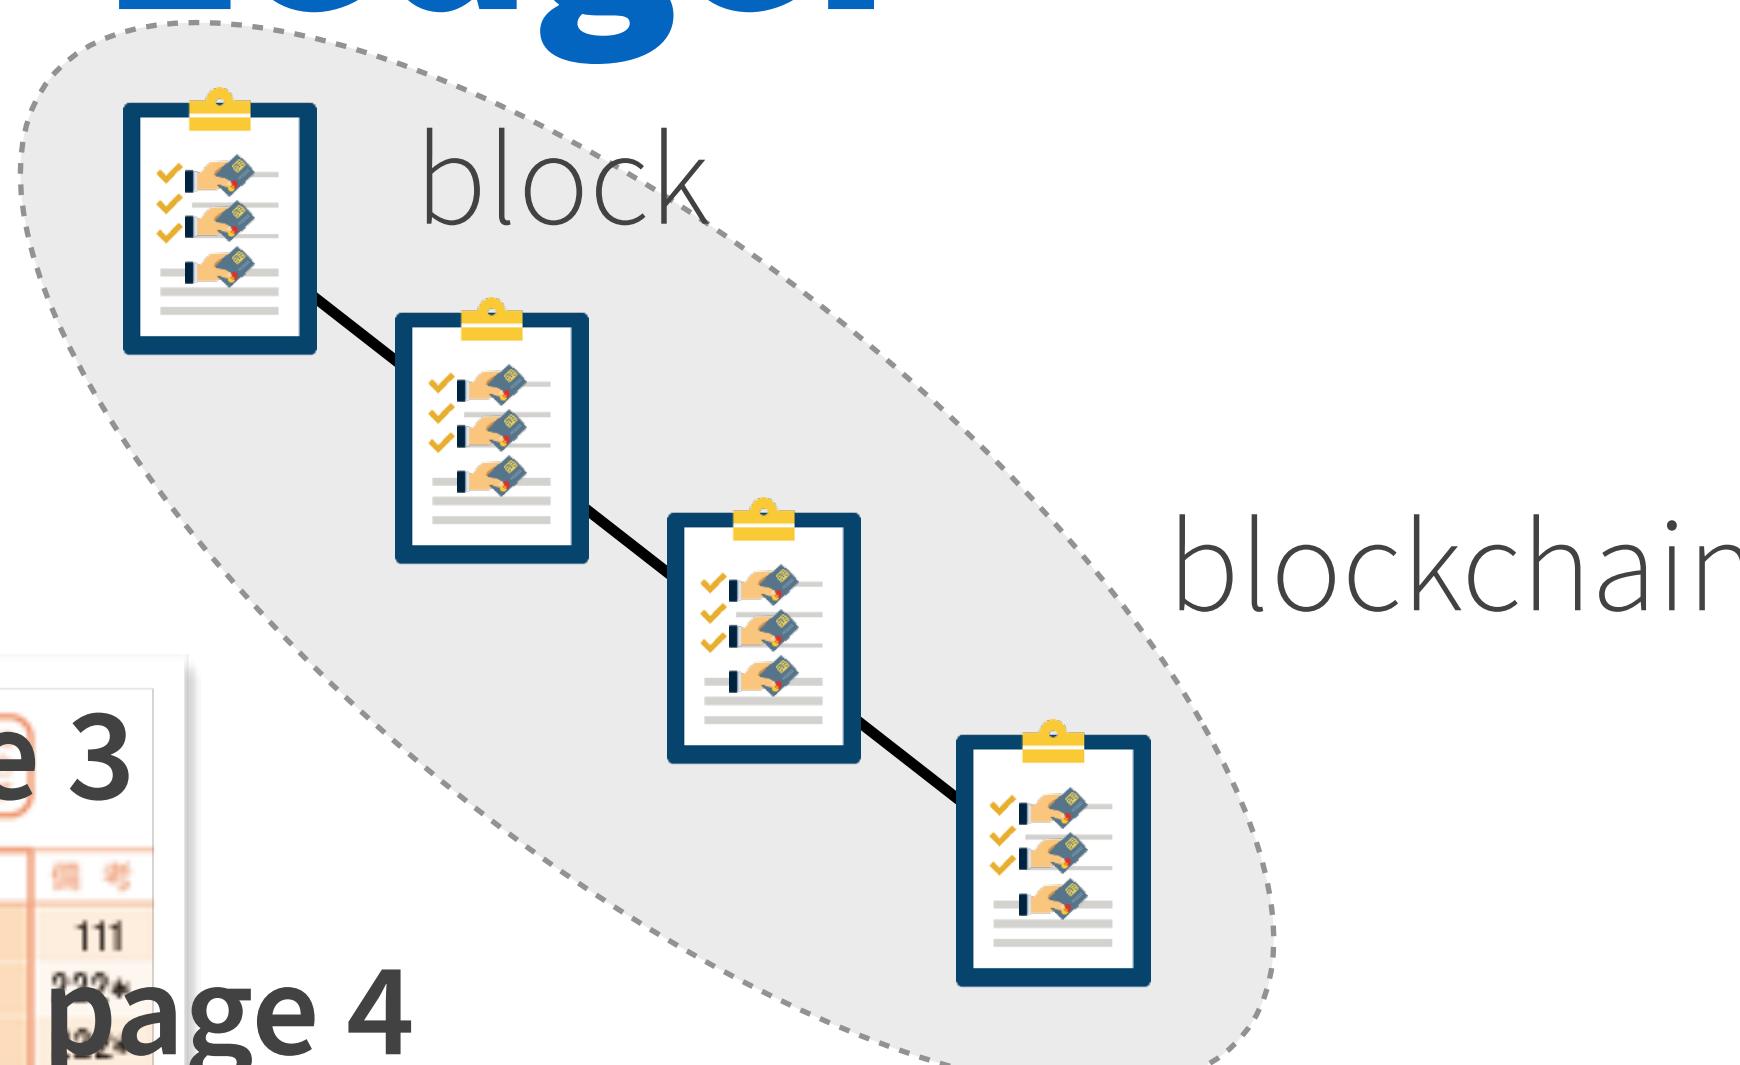
(2001)

Blockchain

Ledger

The image shows four pages of a ledger, each titled "普通預金-1 (兼お借入明細)". Each page has a header with a barcode and a table with columns: 年月日 (Date), 記号 (Number), お支払い金額(円) (Amount Paid Out), お預り金額(円) (Amount Received), 差引残高(円) (Balance), and 備考 (Remarks). The pages show a sequence of transactions.

	年月日	記号	お支払い金額(円)	お預り金額(円)	差引残高(円)	備考
1	24-12-15	100	-10,000	CD	*232,650	111
2	24-12-25	100	-10,000	CD	*222,650	222*
3	25-01-07	100	-20,000	CD	*202,650	222*
4	25-01-10	100	-100,000	CD	*102,650	333*
5	25-02-14	900 フク	+35,000	CD	*137,650	333*
6	25-02-15	900 フク	20,000	CD	*157,650	333*
7	25-02-25	100	+5,000	CD	*152,650	222*
8	25-02-28	100	+23,000	CD	*129,650	222*
9	25-03-03	100	+5,000	CD	*124,650	333*
10	25-03-10	900 フク	+135,000	CD	*259,650	333*
11	25-03-14	100	+8,000	CD	*251,650	333*
12	25-03-15	100	+12,000	CD	*239,650	333*



append page

This image shows a single ledger page containing all the transactions from the four pages above. The header includes a note: "差引残高の金額頭部に「マイナス(-)」がある場合はお借入残高を表わします". The table structure is identical to the ones above.

	年月日	記号	お支払い金額(円)	お預り金額(円)	差引残高(円)	備考
1	24-12-15	100	-10,000	CD	*232,650	111
2	24-12-25	100	-10,000	CD	*222,650	222*
3	25-01-07	100	-20,000	CD	*202,650	222*
4	25-01-10	100	-100,000	CD	*102,650	333*
5	25-02-14	900 フク	+35,000	CD	*137,650	333*
6	25-02-15	900 フク	20,000	CD	*157,650	333*
7	25-02-25	100	+5,000	CD	*152,650	222*
8	25-02-28	100	+23,000	CD	*129,650	222*
9	25-03-03	100	+5,000	CD	*124,650	333*
10	25-03-10	900 フク	+135,000	CD	*259,650	333*
11	25-03-14	100	+8,000	CD	*251,650	333*
12	25-03-15	100	+12,000	CD	*239,650	333*



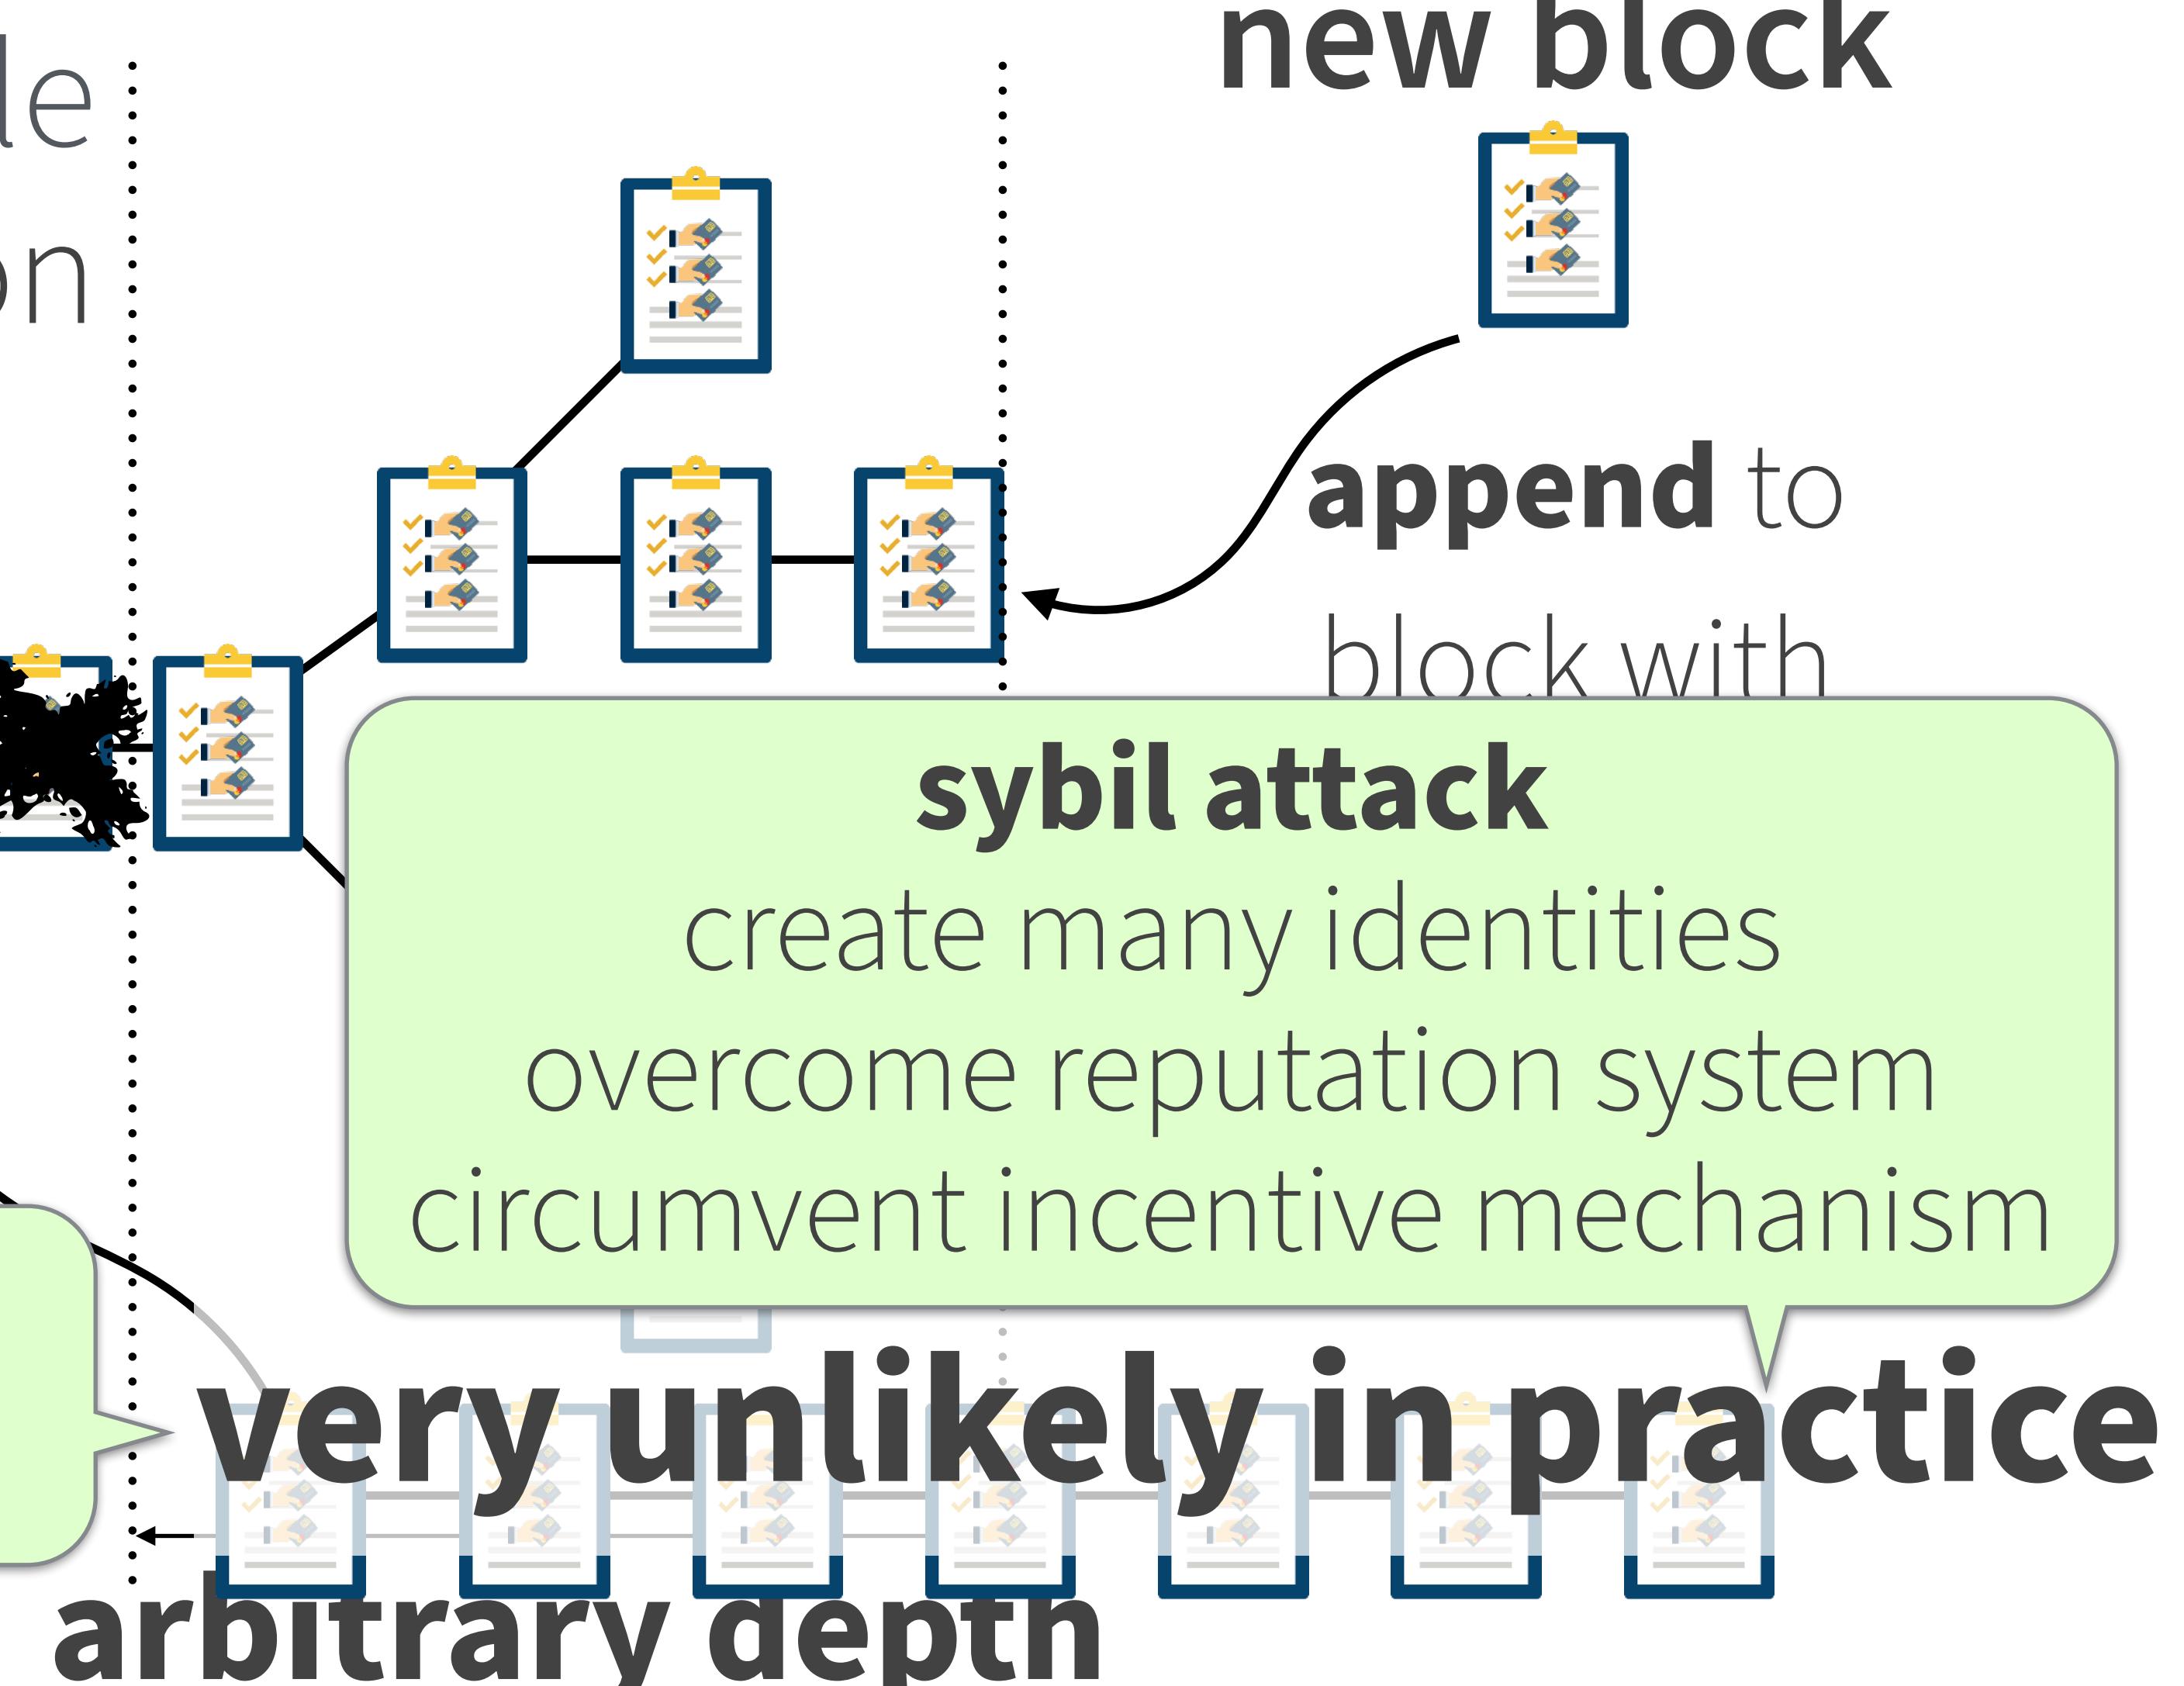
Distributed Ledger

assumed final/stable
& agreed upon

Root



finality
not guaranteed

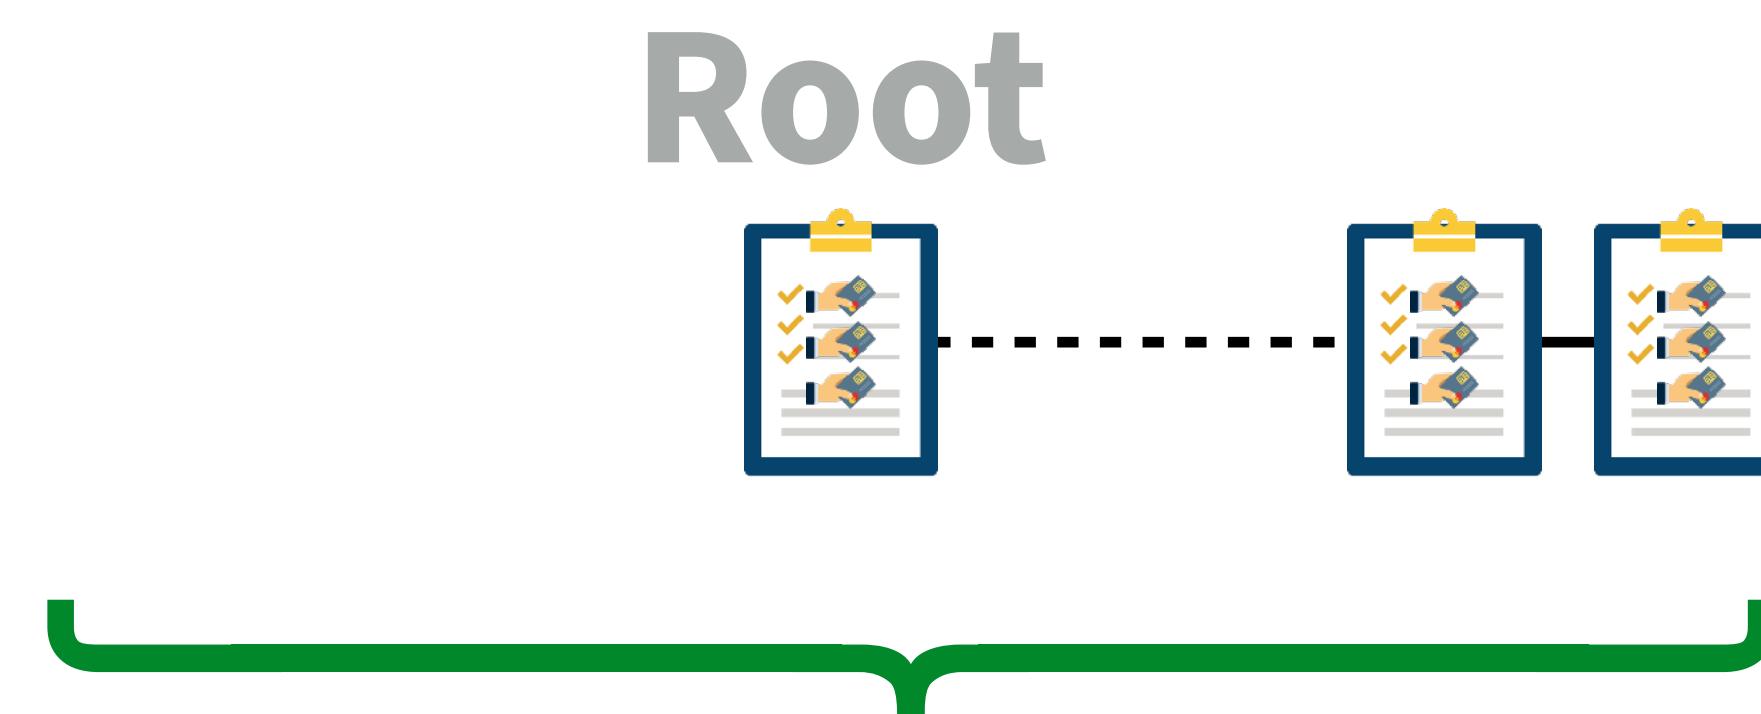


(typ. 5-6 blocks for Bitcoin)

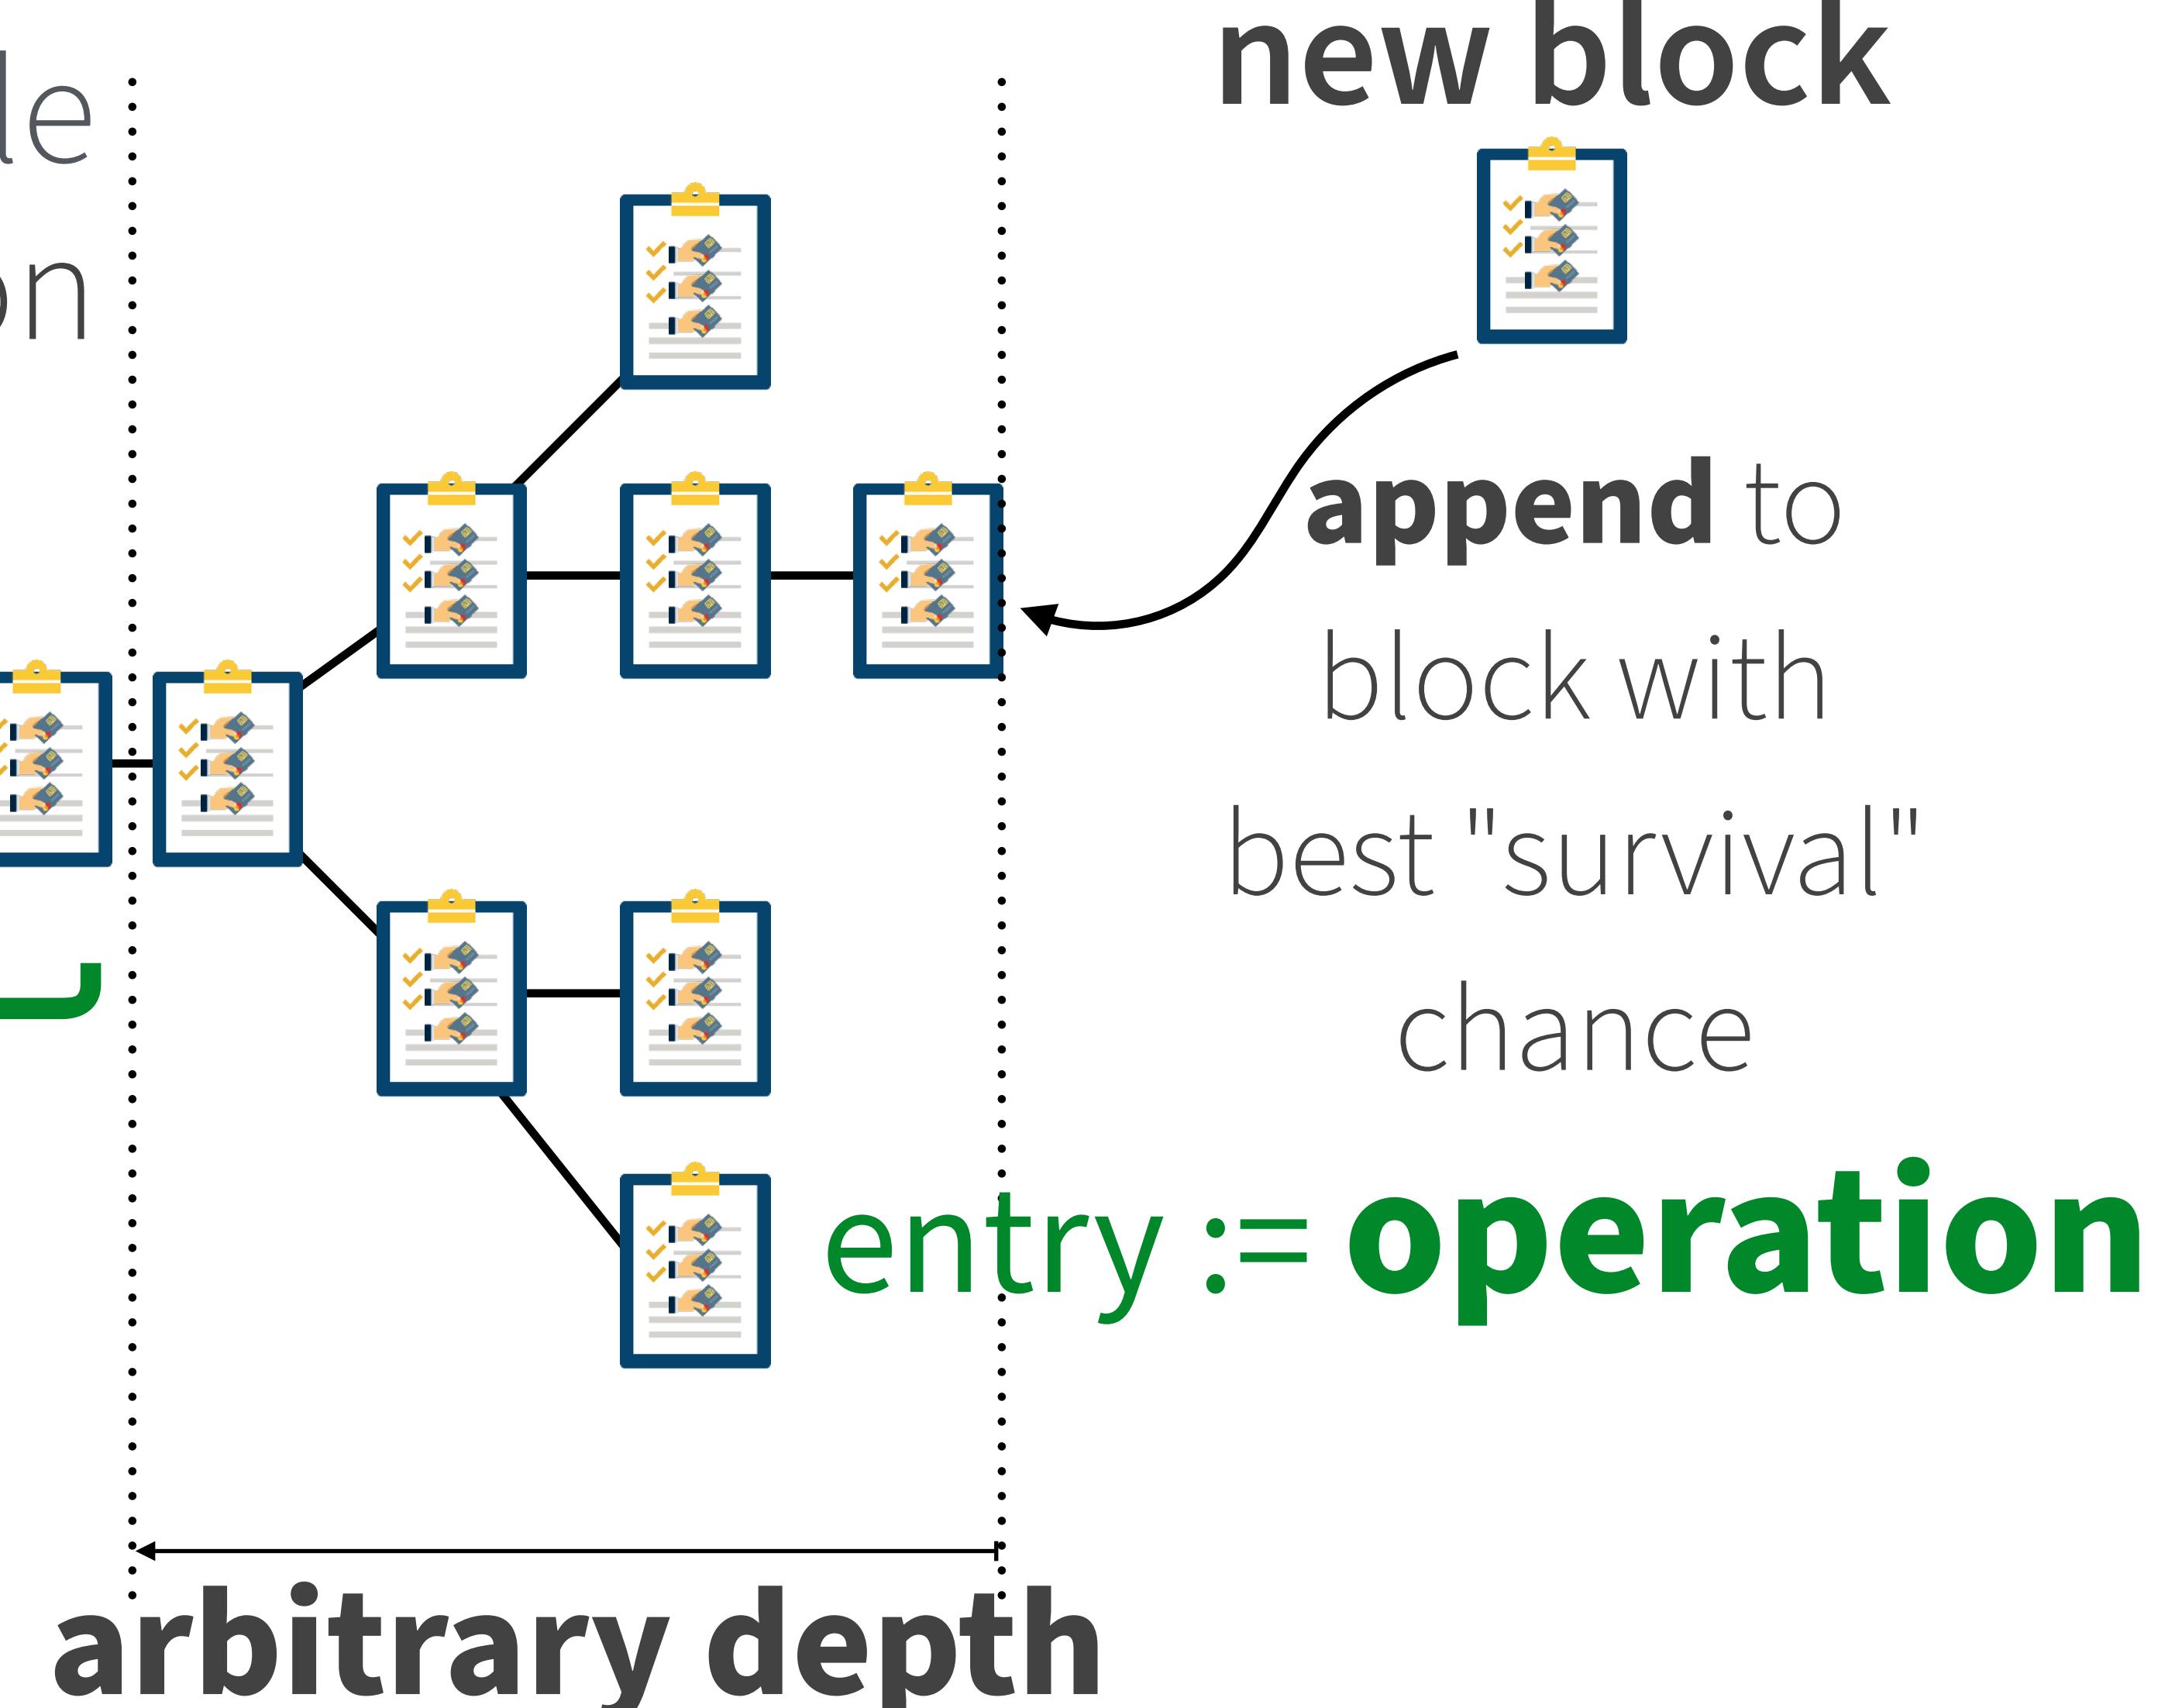


Distributed Ledger

assumed final/stable
& agreed upon



Total Order
on operations



(typ. 5-6 blocks for Bitcoin)

Comparison

▶ Byzantine fault-tolerance

- ▶ strong consistency
- ▶ **always safe**; mostly live
- ▶ **partition**: majority-only
- ▶ finality: commit => permanent

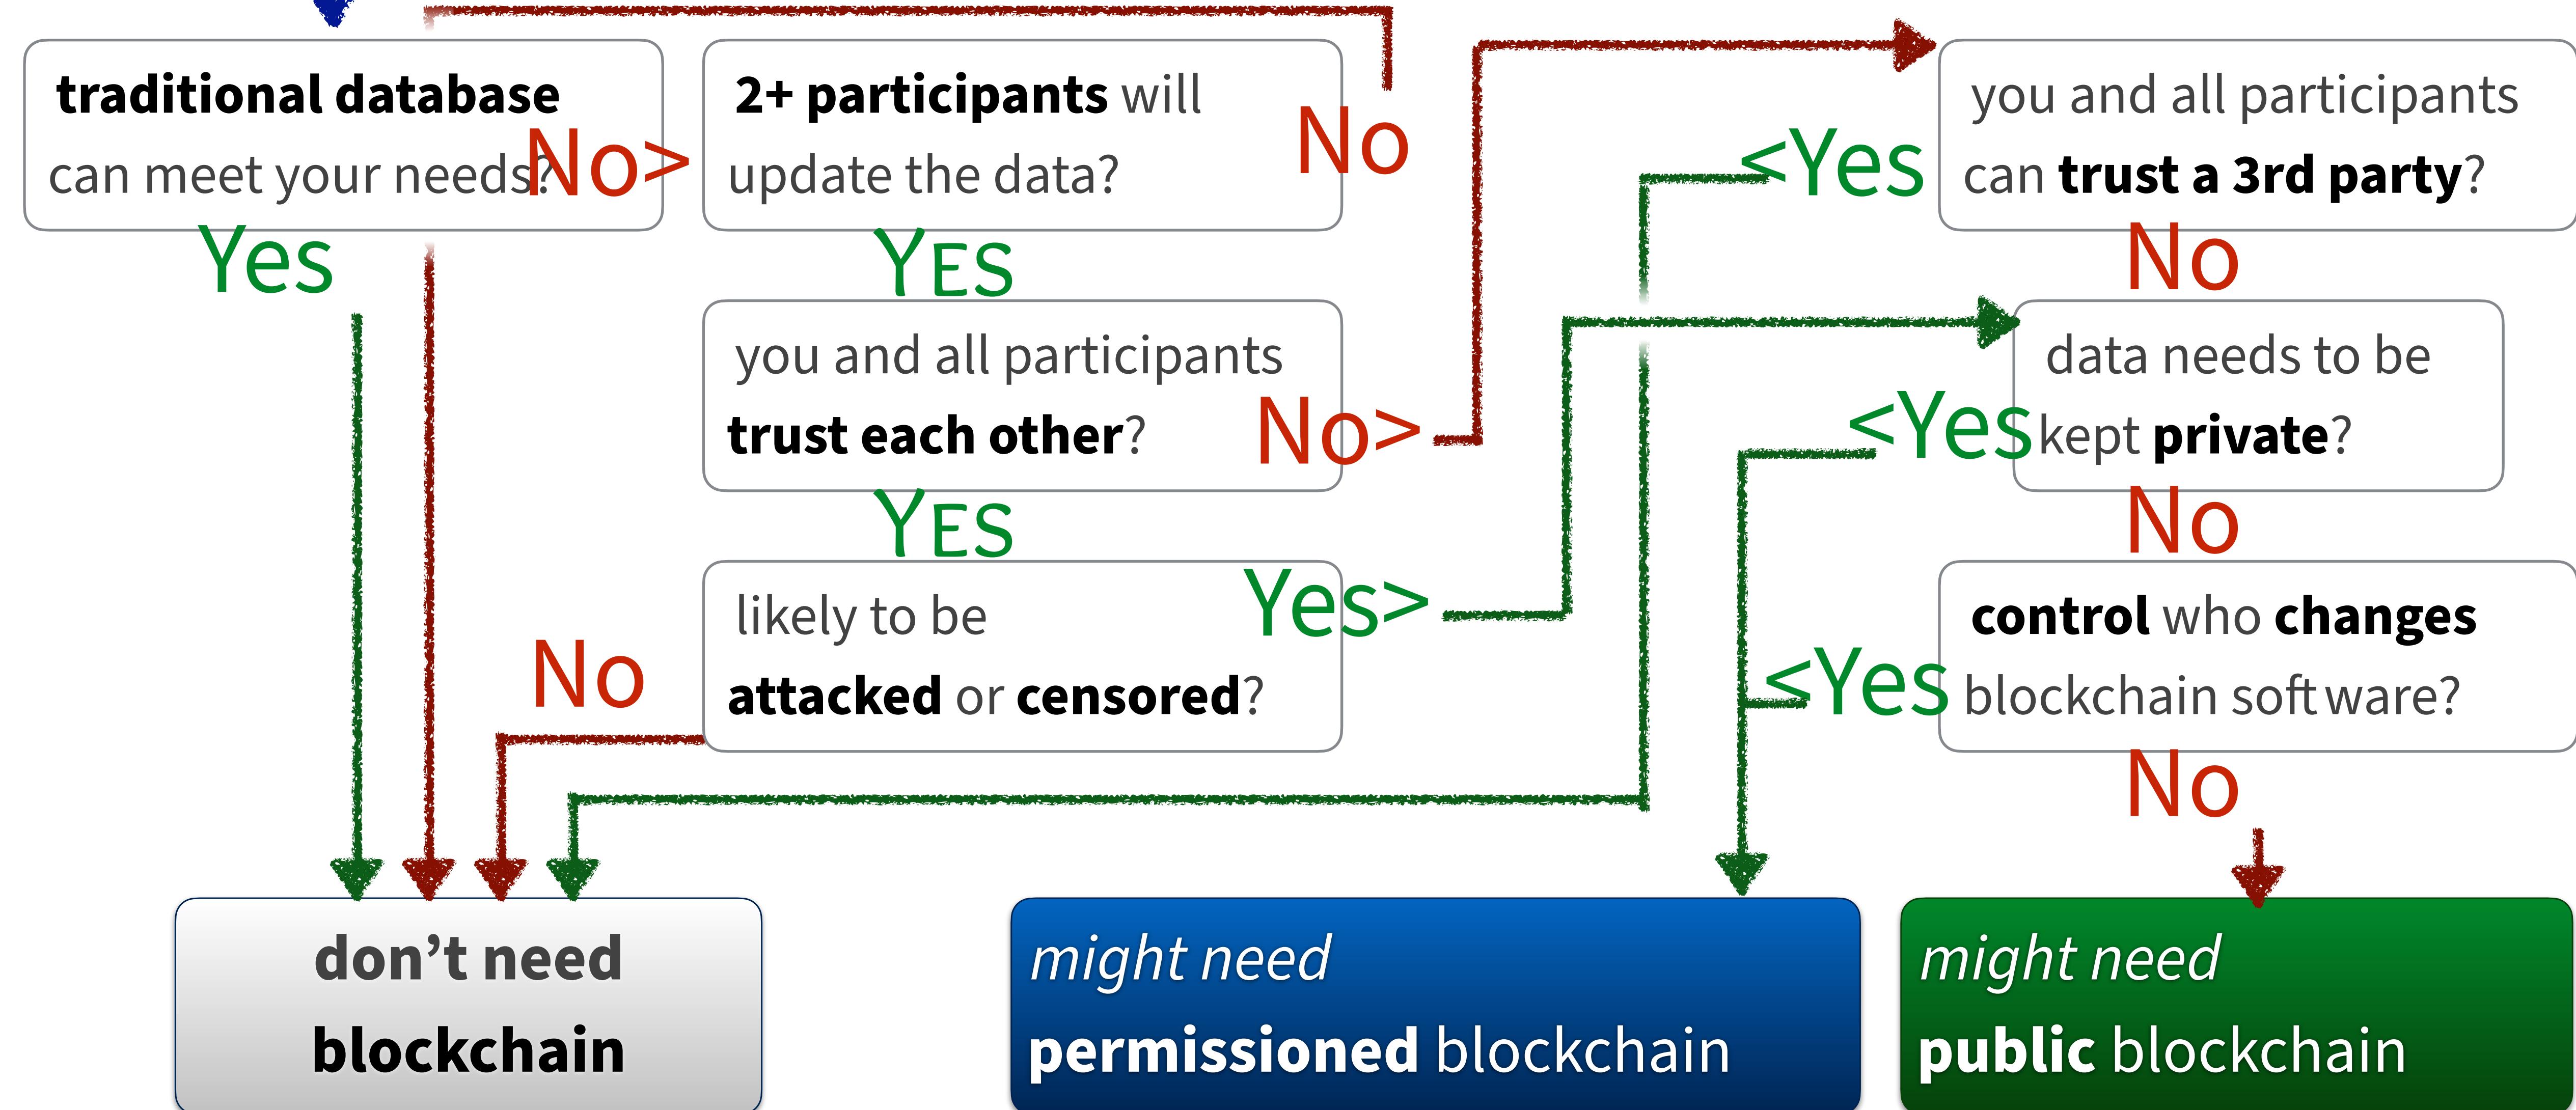
▶ Blockchain

- ▶ mostly safe; **always live**
- ▶ **partition**: loss at merge
- ▶ **finality**: high probability

▶ Overall

- ▶ each serves purpose
- ▶ address different issues
- ▶ not mutually exclusive
- ▶ BFT used in blockchain (PoS)
- ▶ tradeoff: CAP theorem
 - ▶ consistency
 - ▶ availability
 - ▶ partition-tolerance

I Want a Blockchain!



Source: M.E.Peck. Do you need a blockchain? *IEEE Spectrum*, 54(10):38-39, Oct 2017.

Summing Up

▶ **Distributed Ledger**

- ▶ state-machine replication
- ▶ quorums, agreement
- ▶ problem equivalence

▶ **Byzantine fault-tolerance**

- ▶ Byzantine faults
- ▶ strong consistency
- ▶ finality

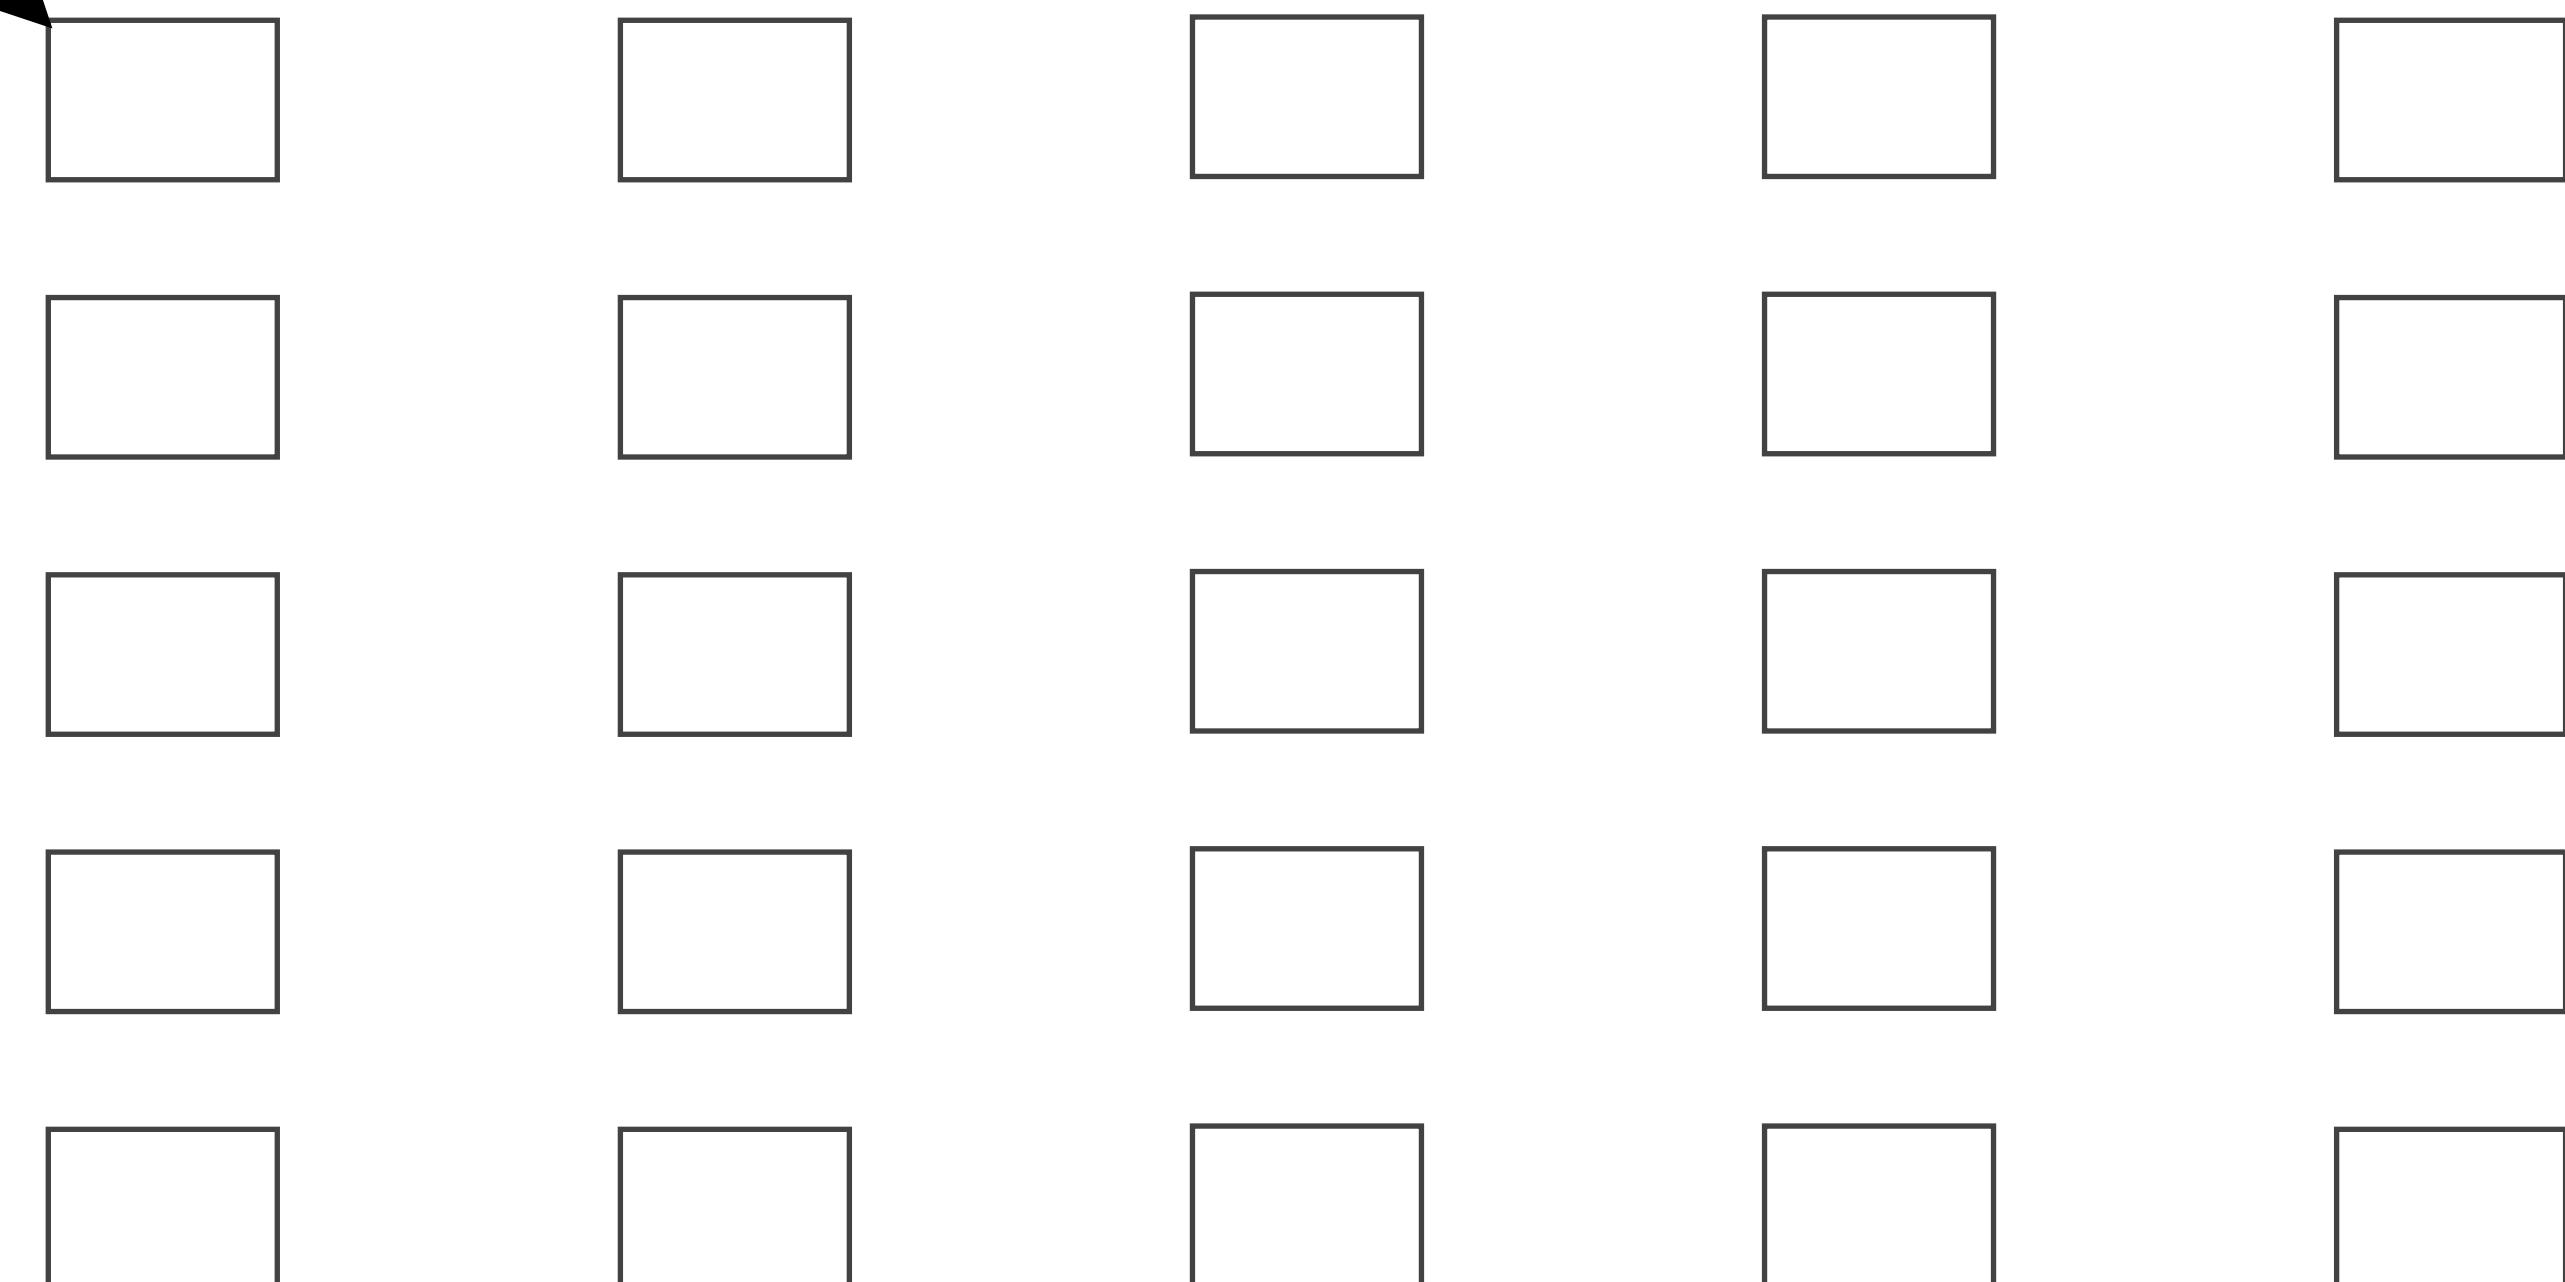
▶ **Blockchain**

- ▶ distributed ledger
- ▶ eventual consistency
- ▶ no finality

Self-Stabilization

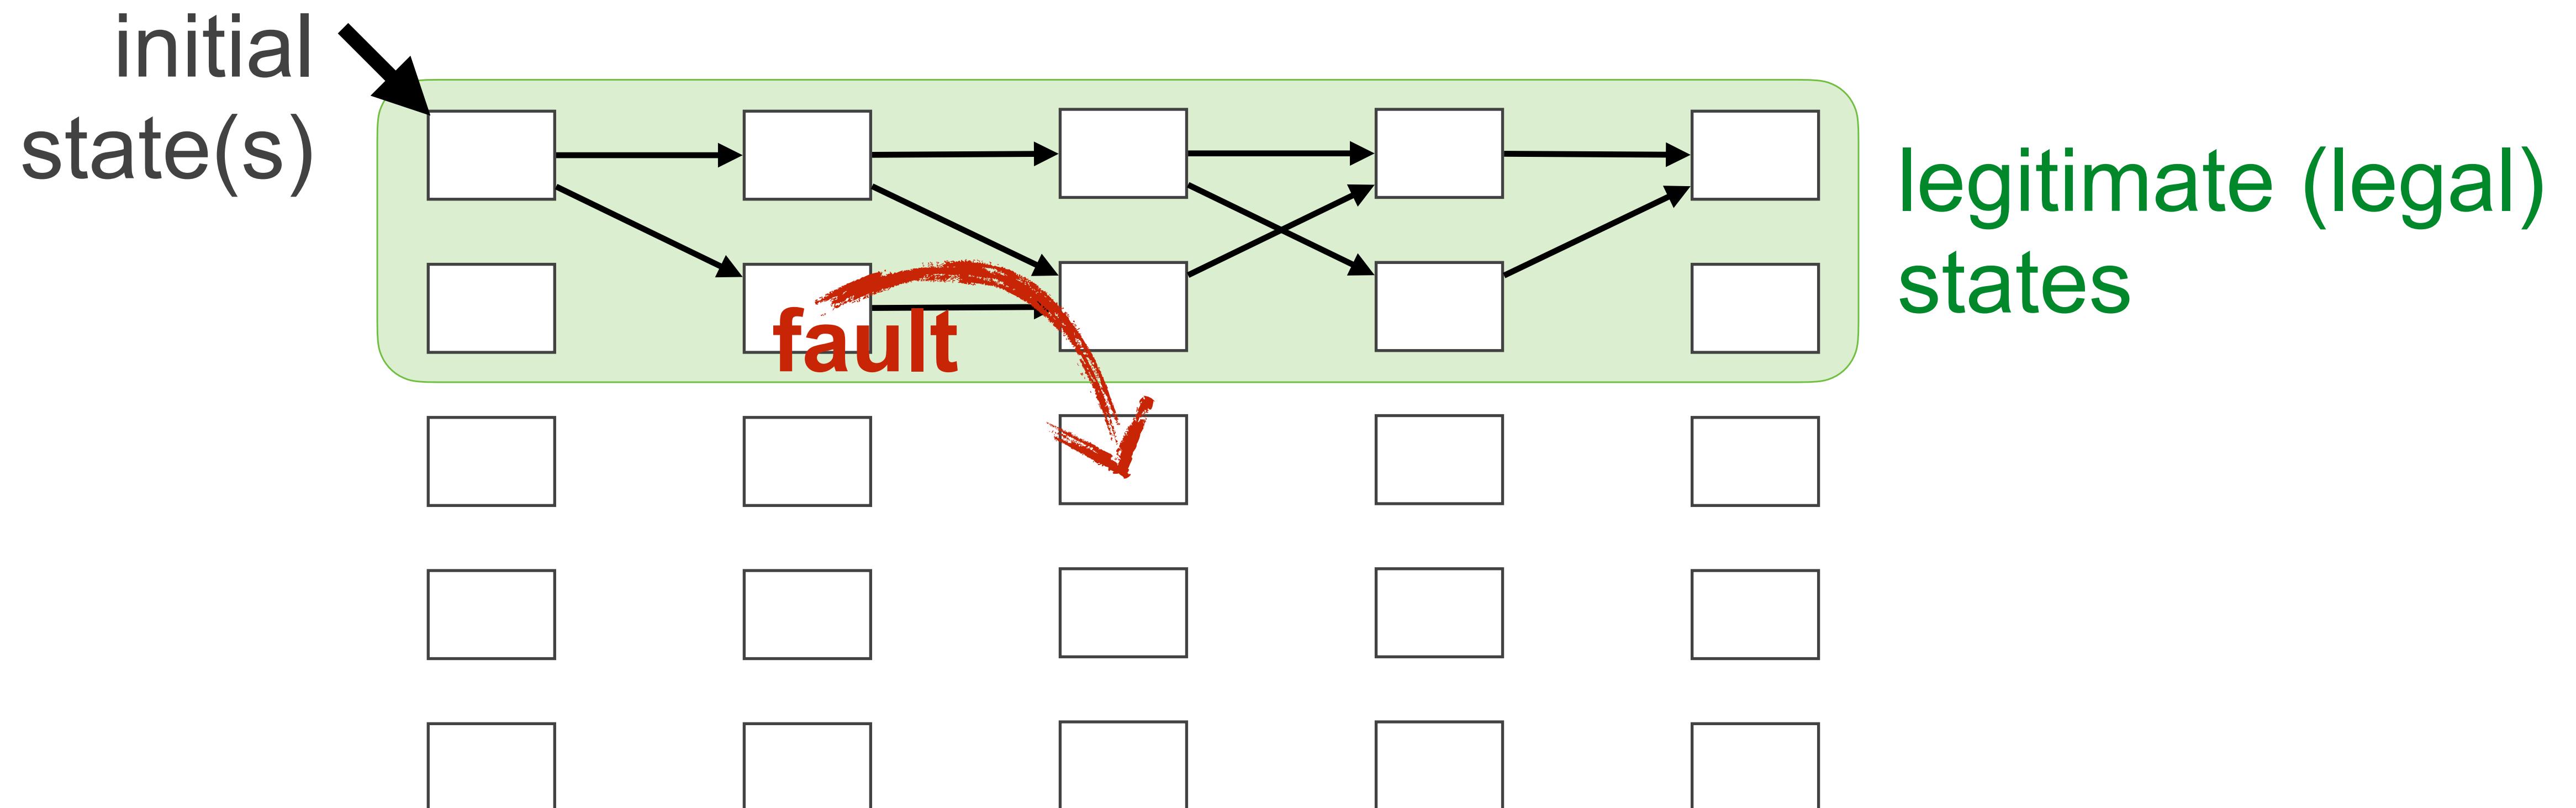
Systems

initial
state(s)



system state

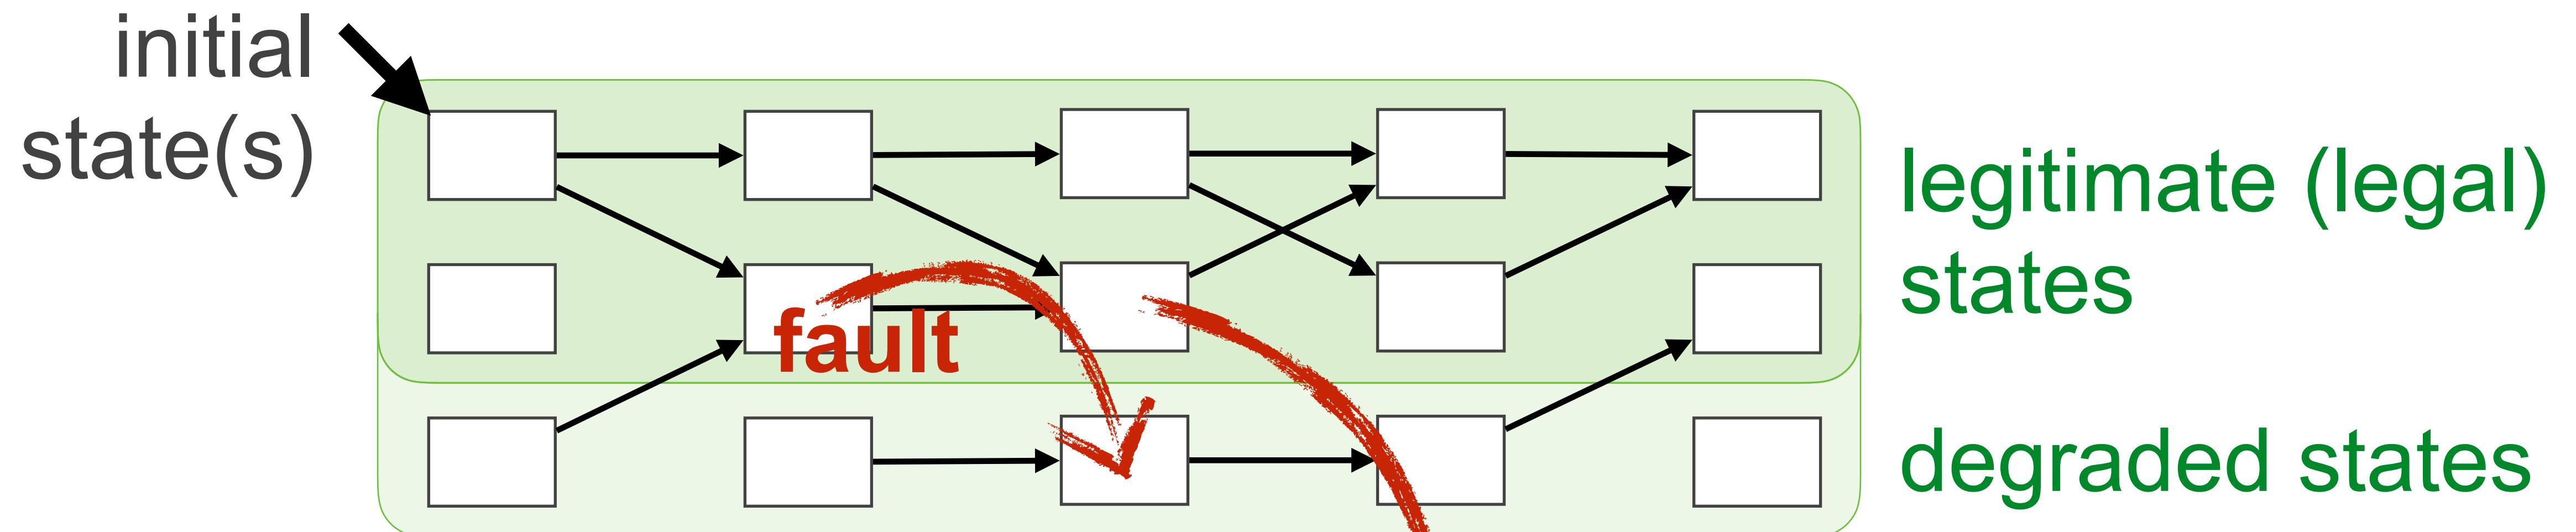
Systems



system state

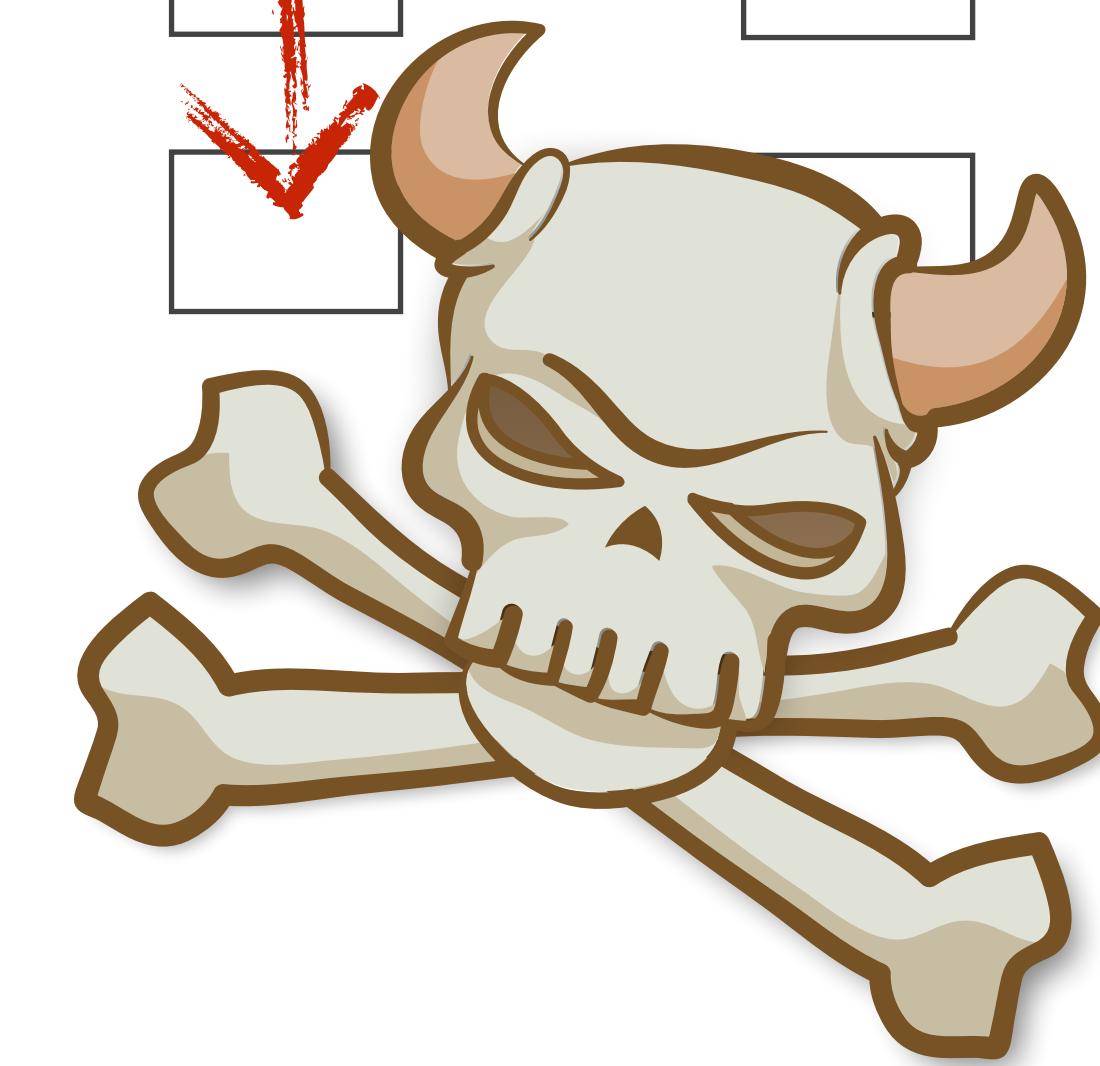
algorithm transitions

Classical Fault-Tolerance

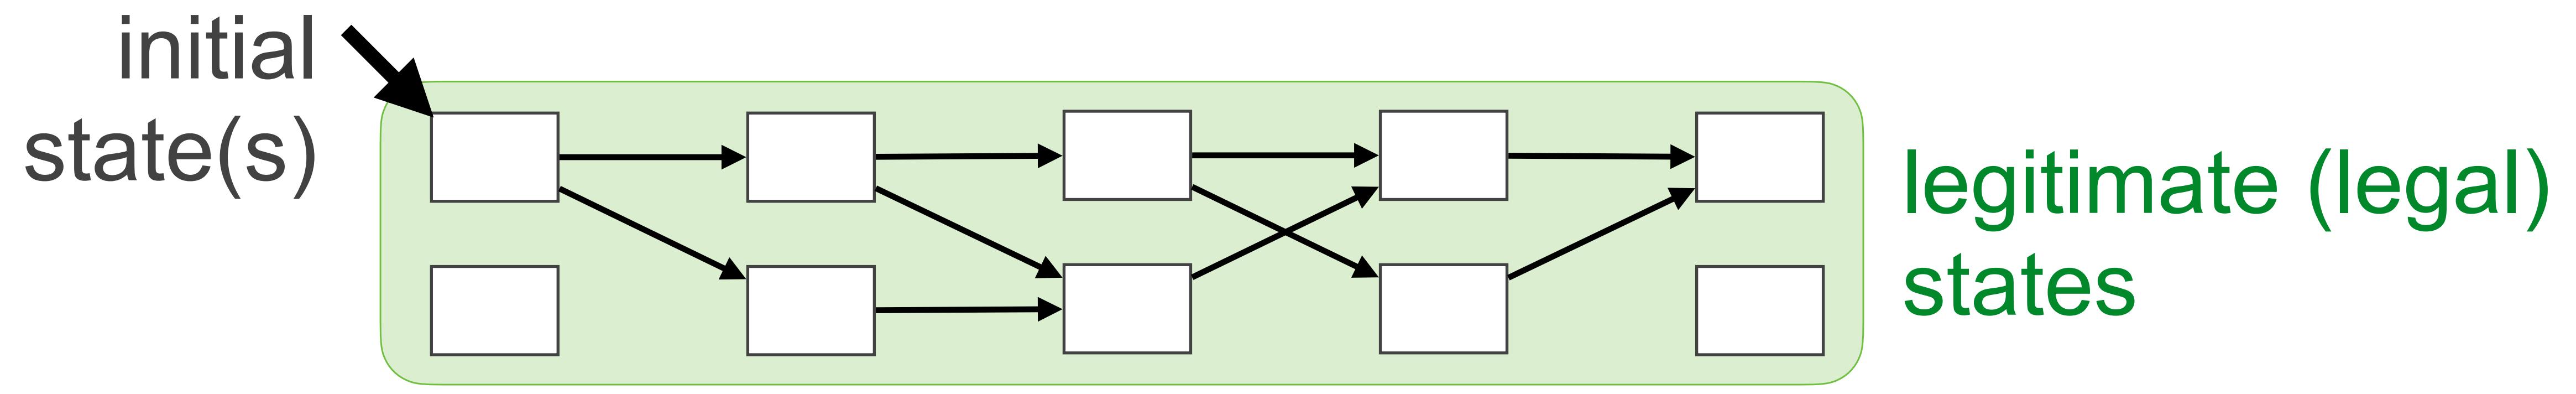


□ system state

→ algorithm transitions



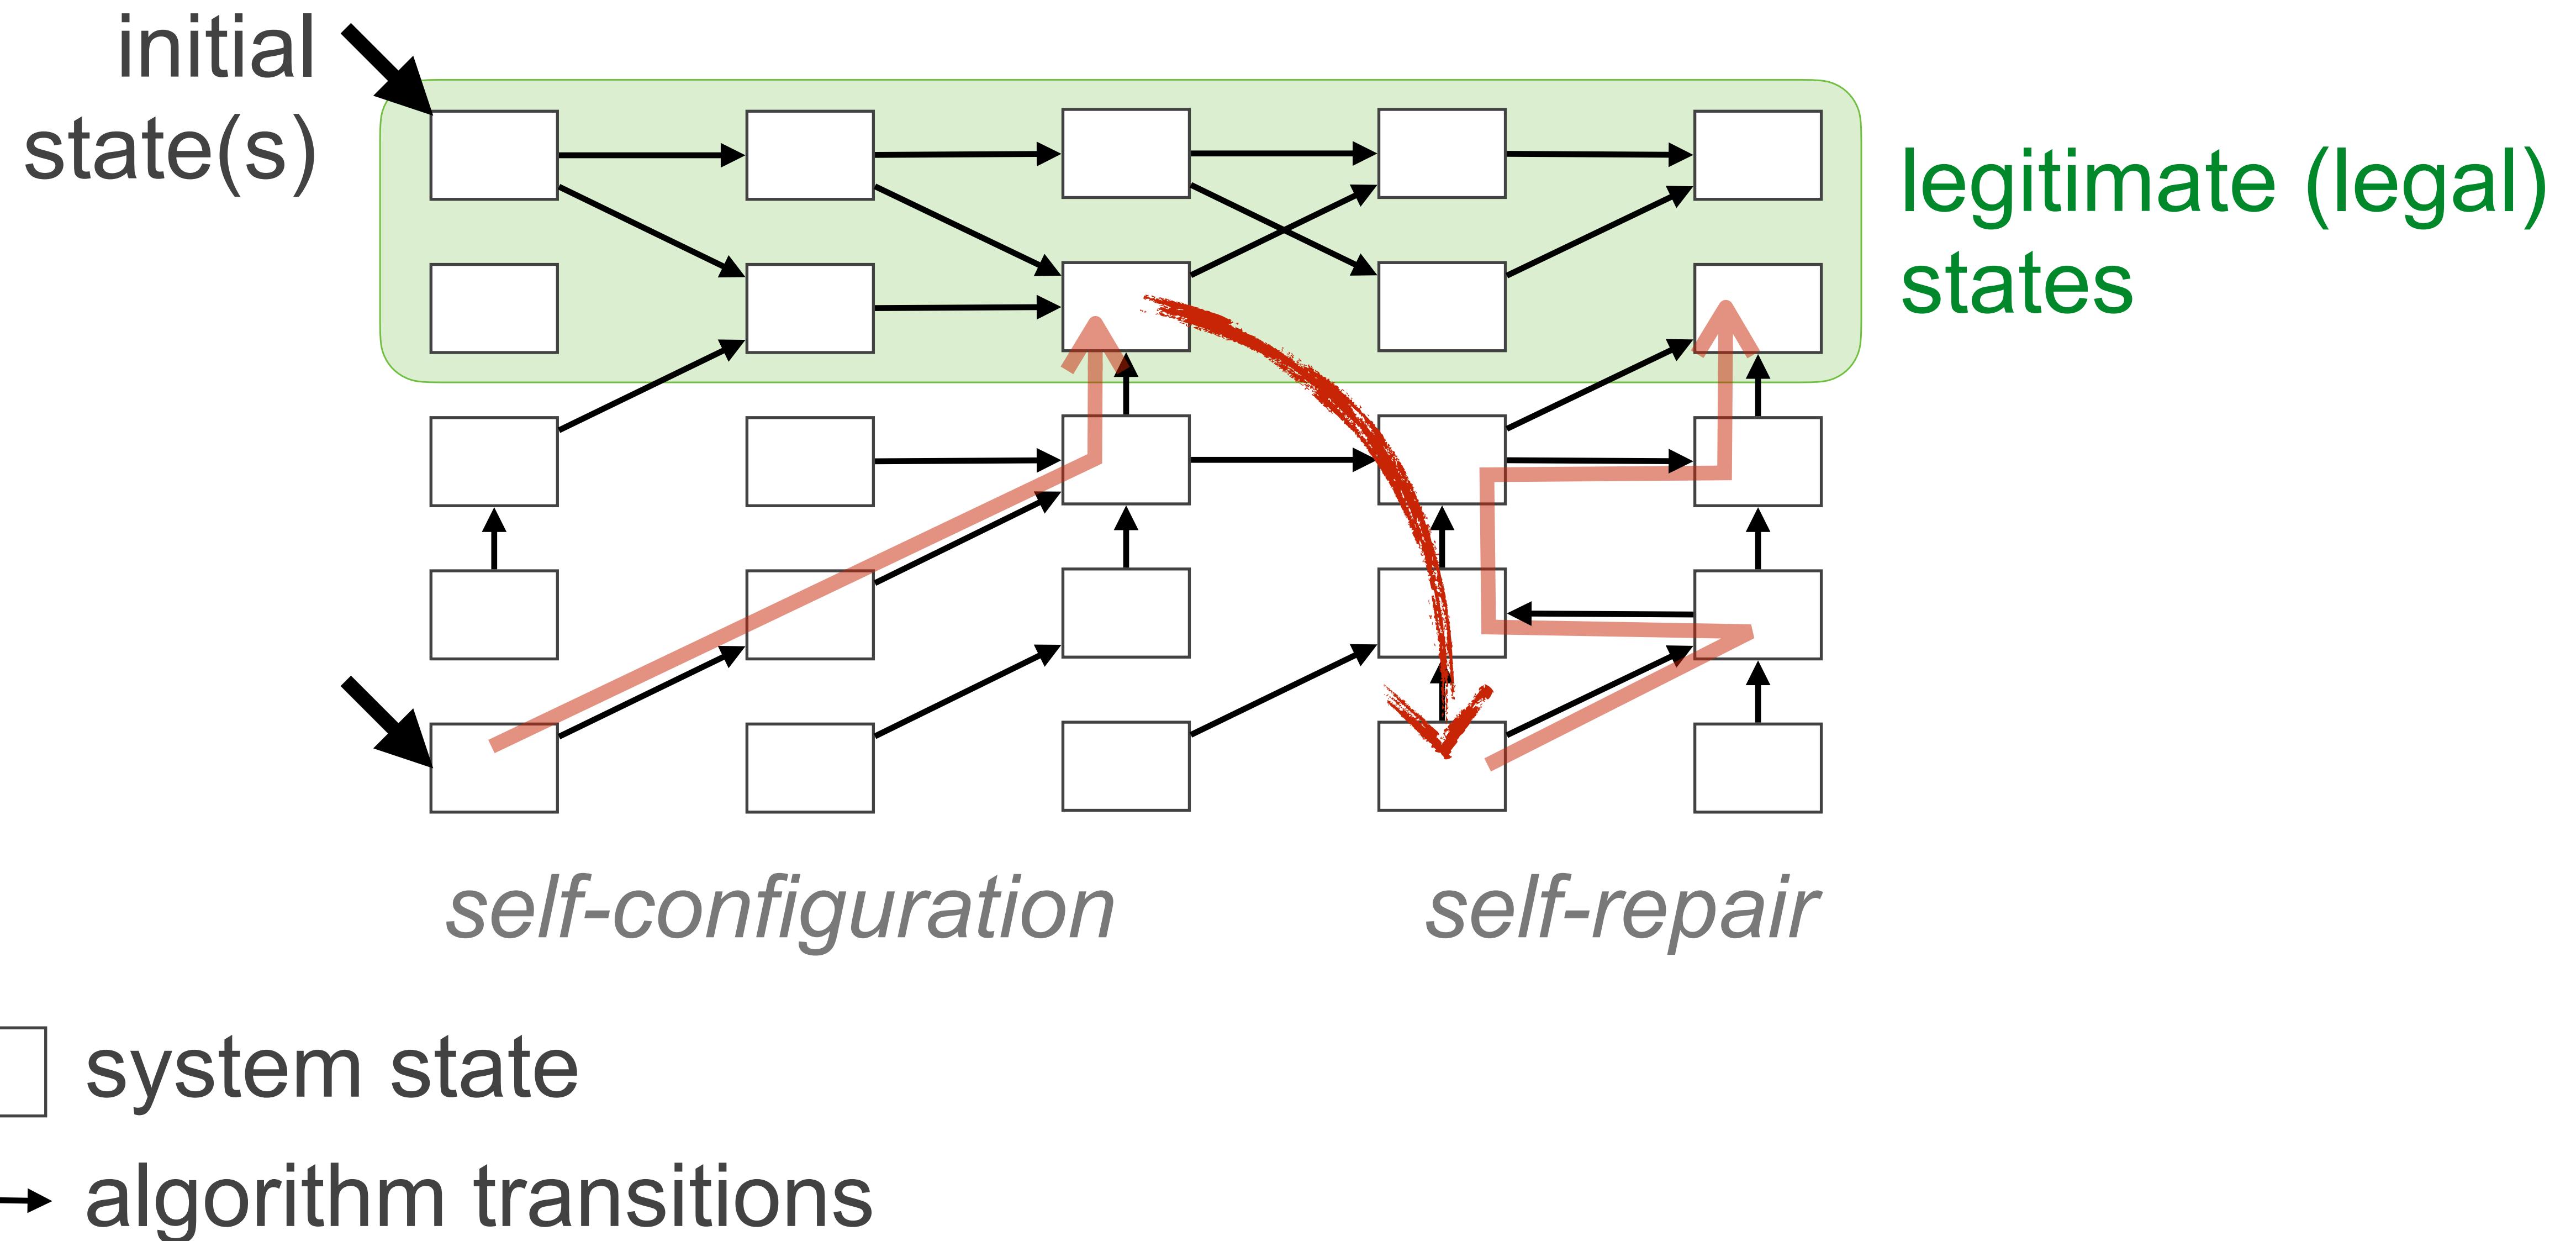
Self-Stabilization



system state

→ algorithm transitions

Self-Stabilization



Self-Stabilization

▶ Definition

- ▶ a system \mathbf{S} is **self-stabilizing** with respect to predicate \mathbf{P} if it satisfies:

▶ Closure

- ▶ \mathbf{P} is closed under the execution of \mathbf{S} .
- ▶ *Intuitive*: always legal state —> legal state

▶ Convergence

- ▶ starting from an arbitrary global state, \mathbf{S} is guaranteed to reach a global state satisfying \mathbf{P} within a finite number of state transitions.
- ▶ *Intuitive*: any state eventually —> legal state

Illustration: BFS tree

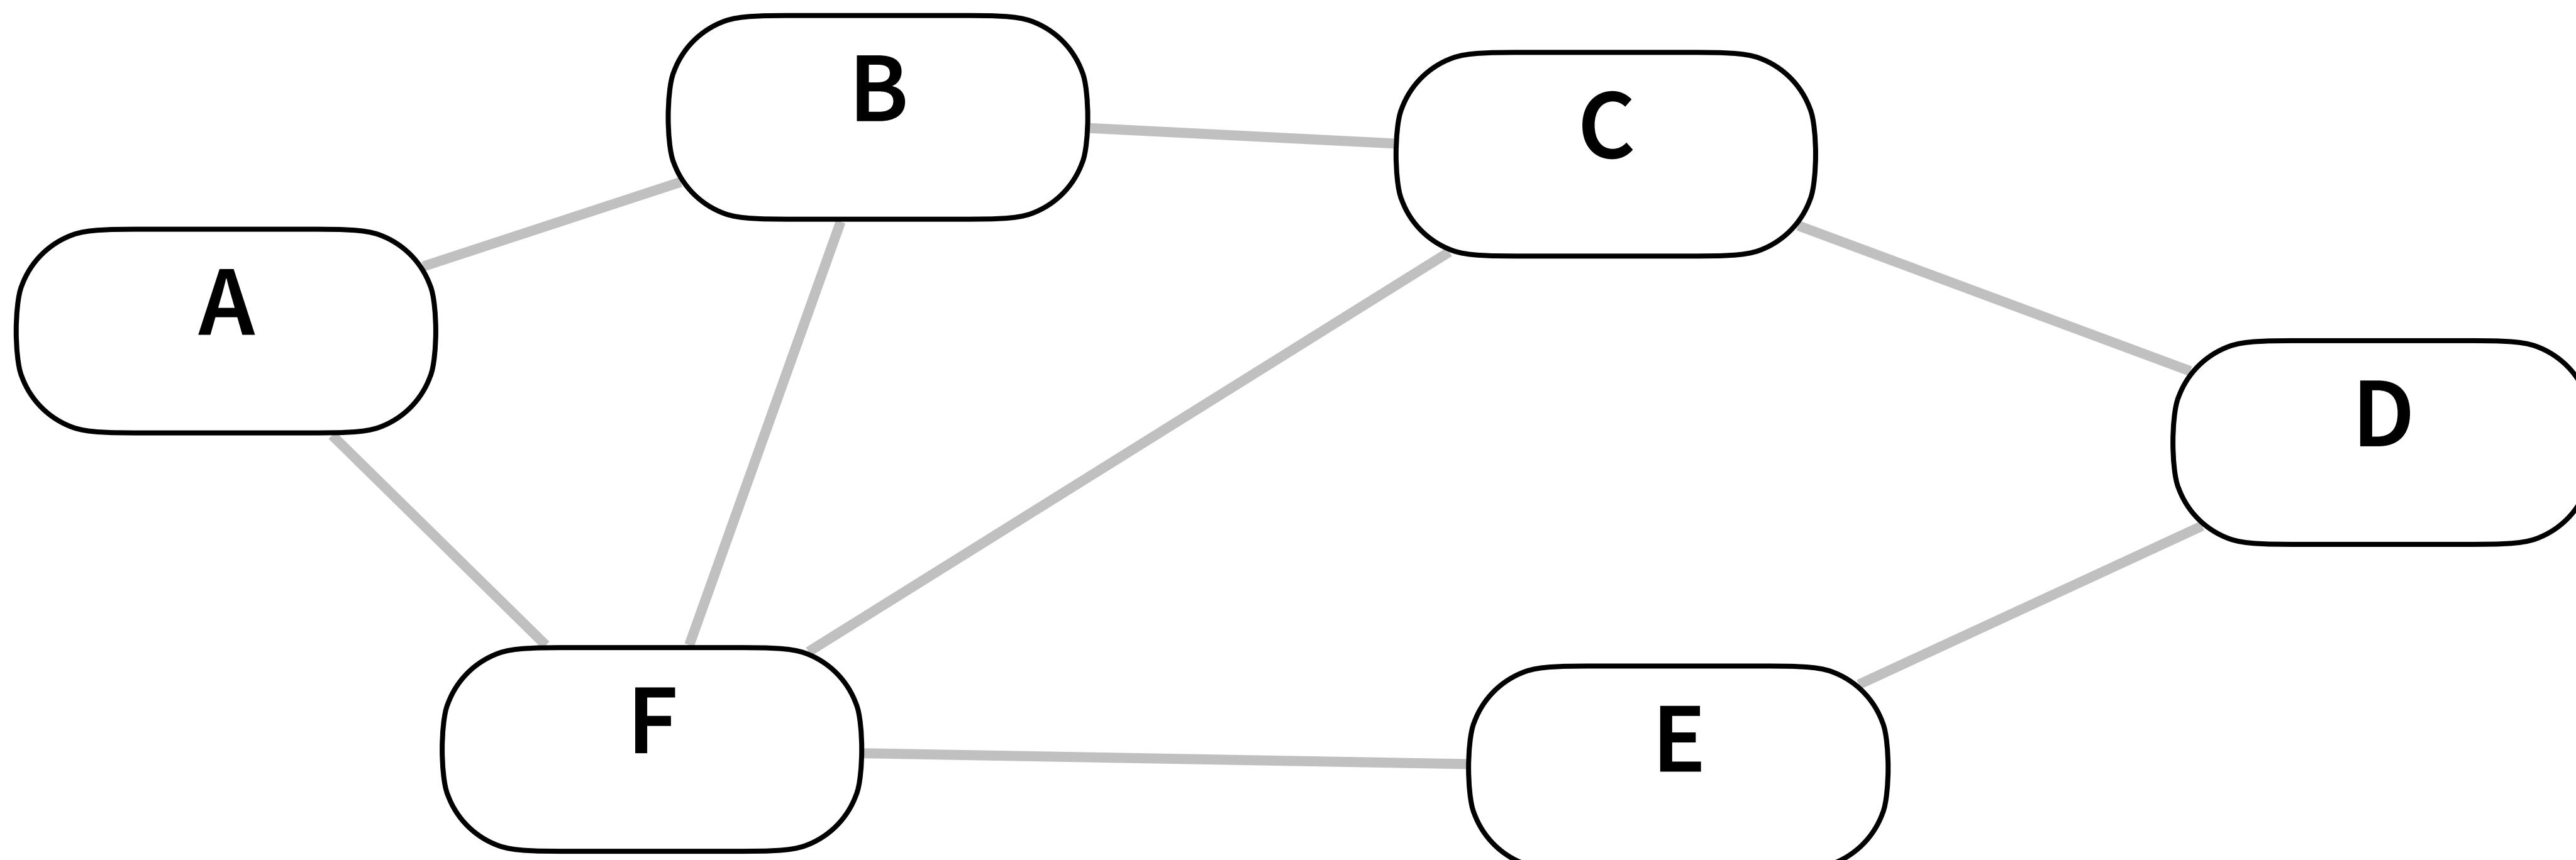


Illustration: BFS tree

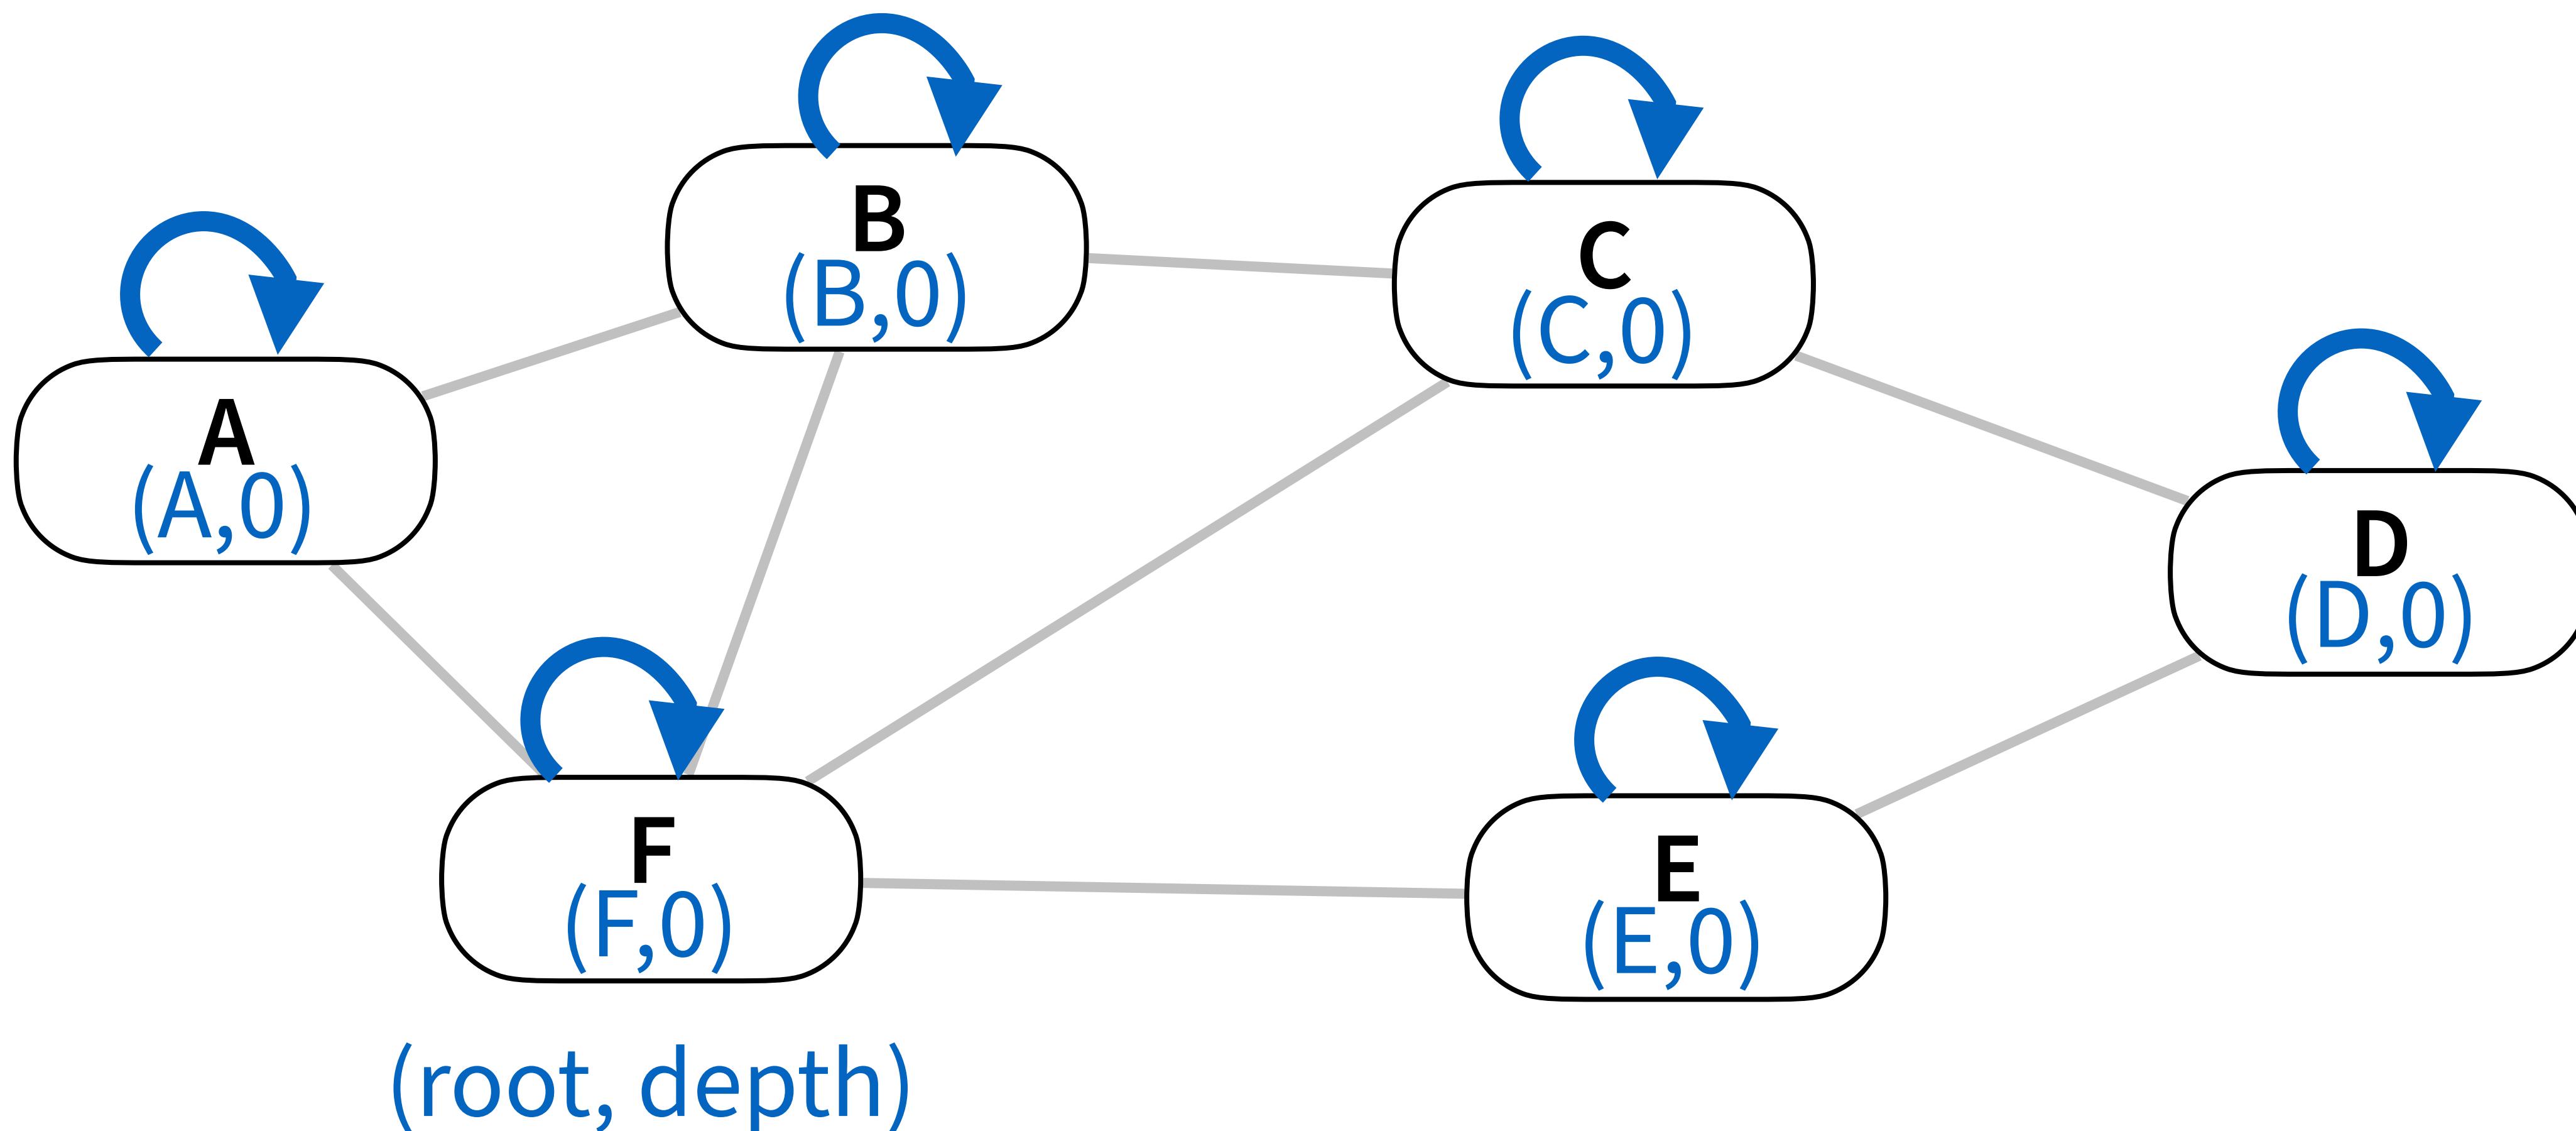


Illustration: BFS tree

$\min \{(\text{root}, \text{depth})\}$

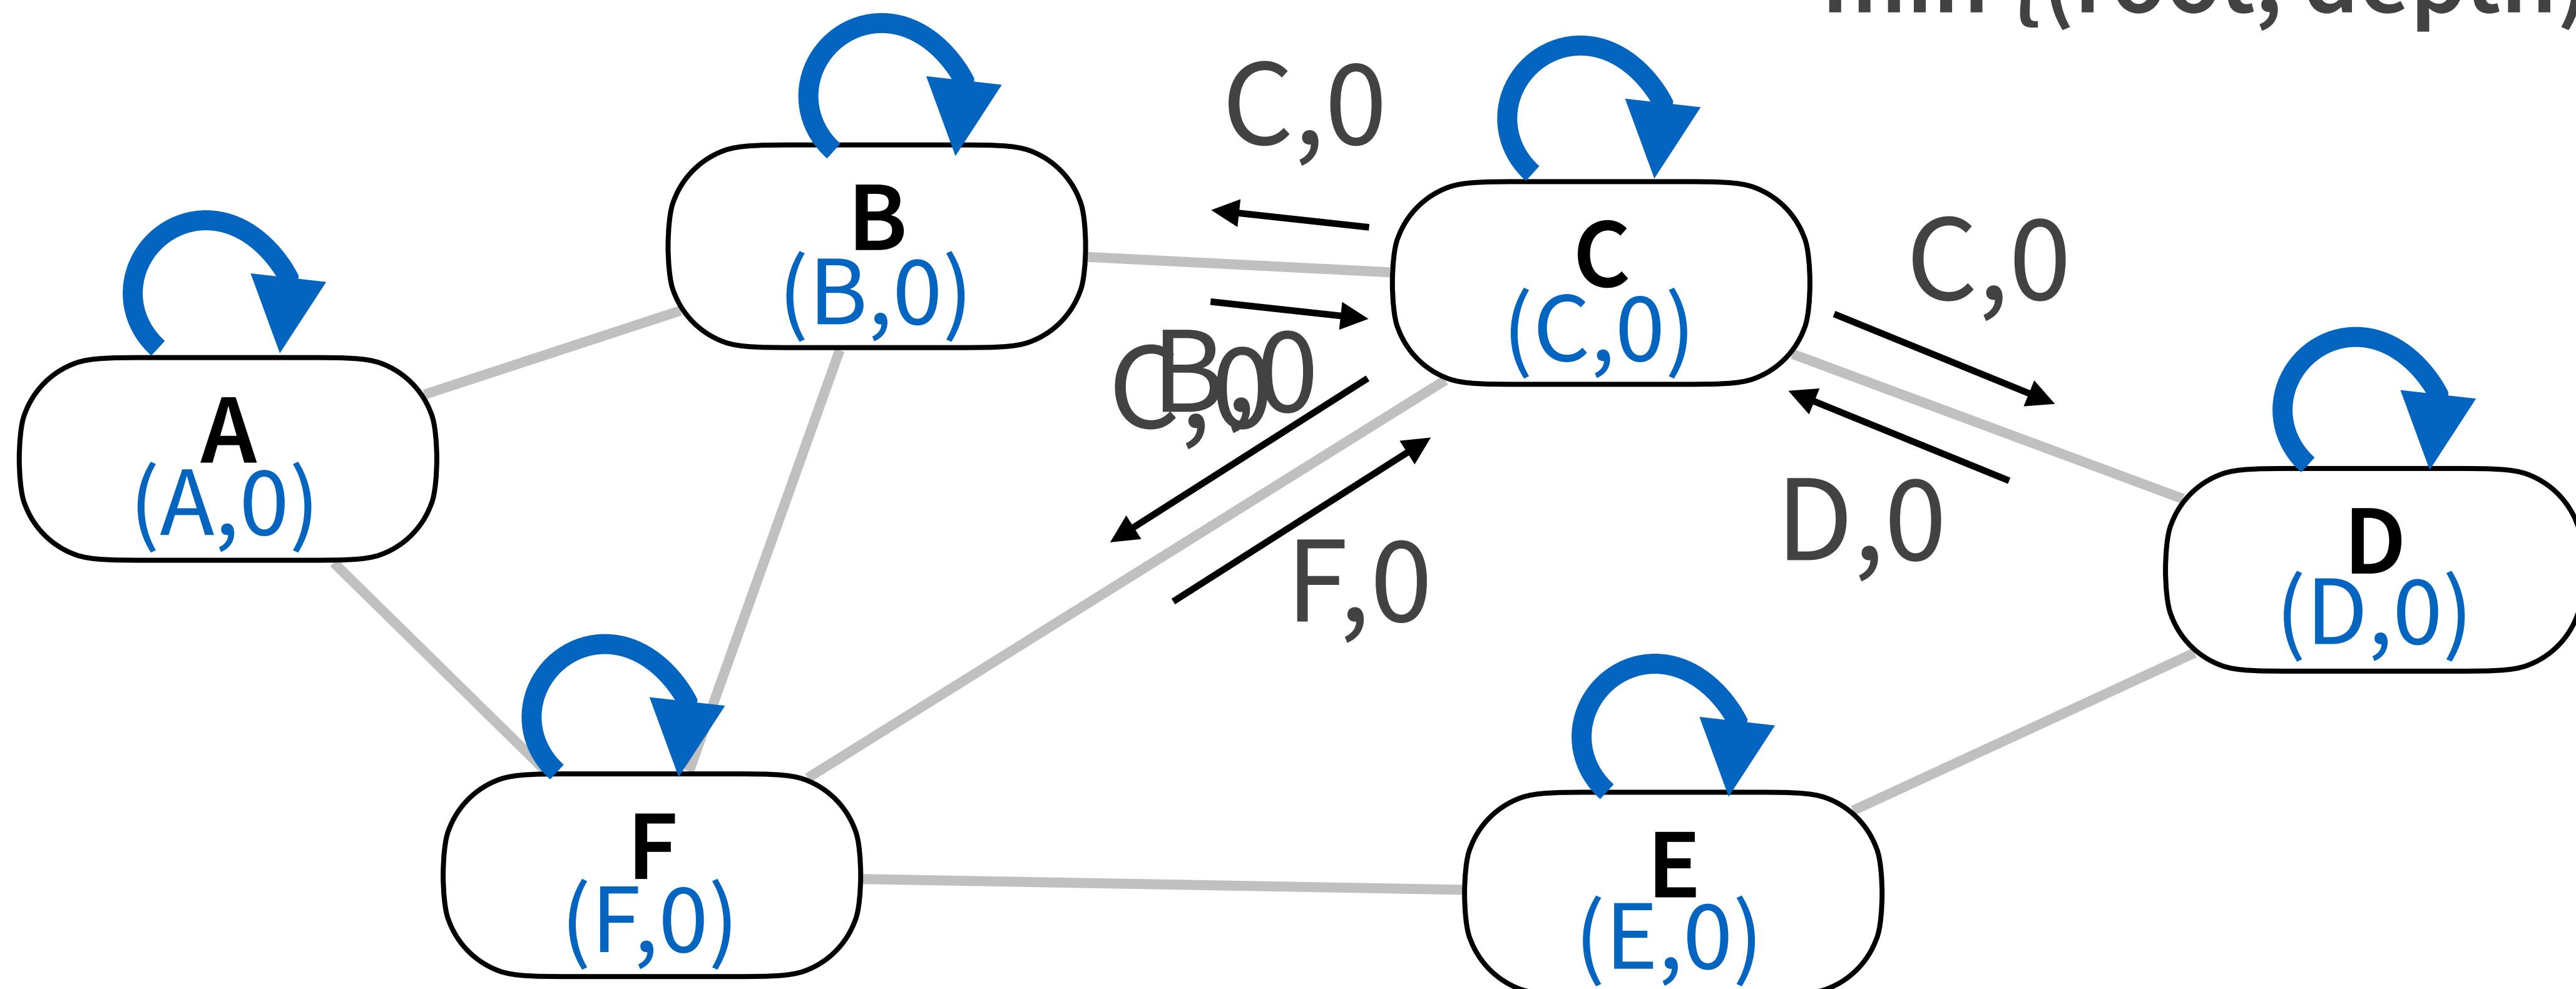


Illustration: BFS tree

$\min \{(\text{root}, \text{depth})\}$

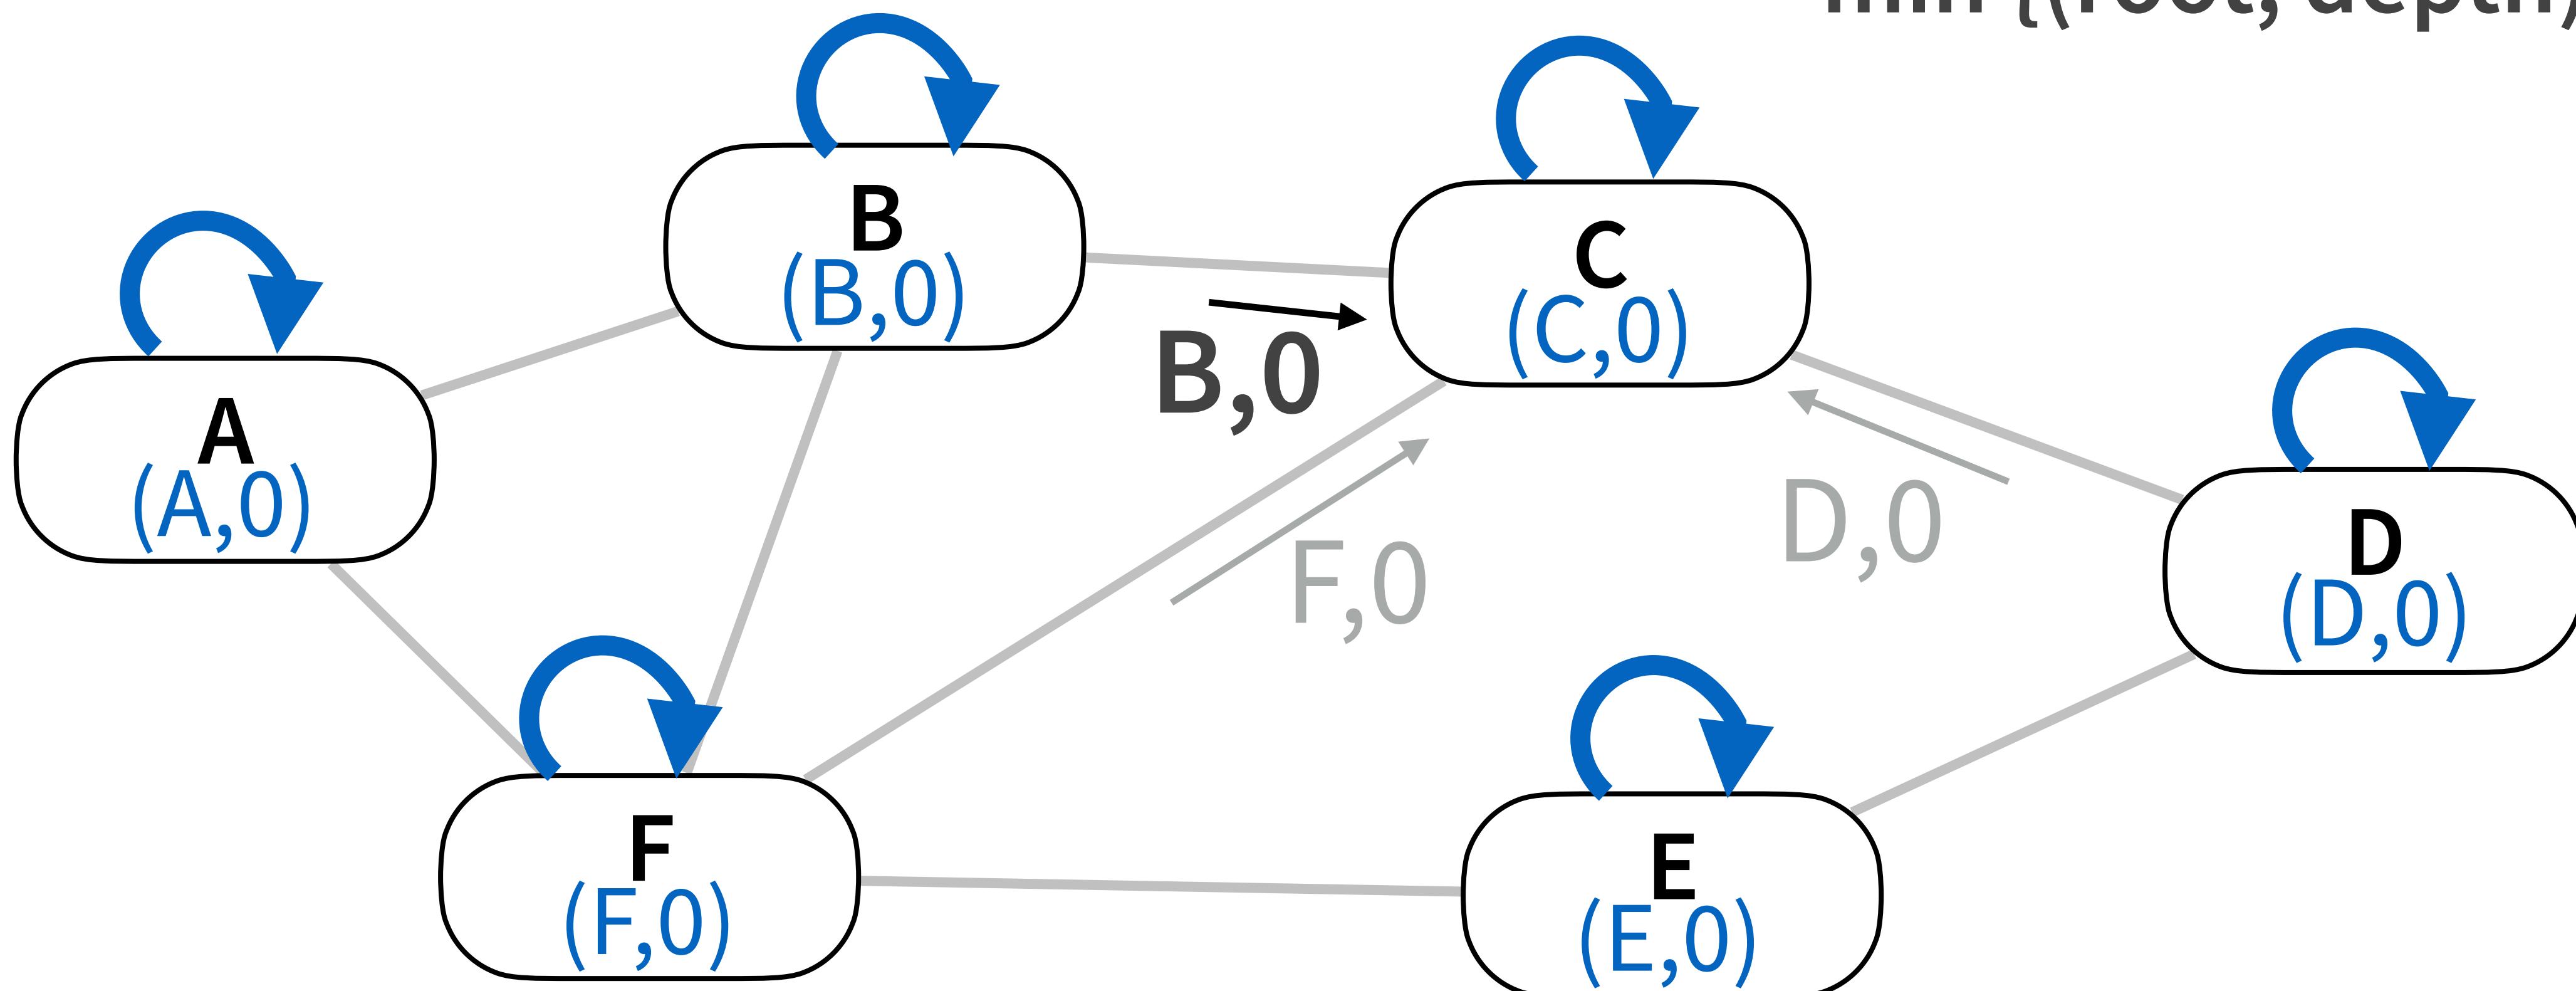


Illustration: BFS tree

$\min \{(\text{root}, \text{depth})\}$

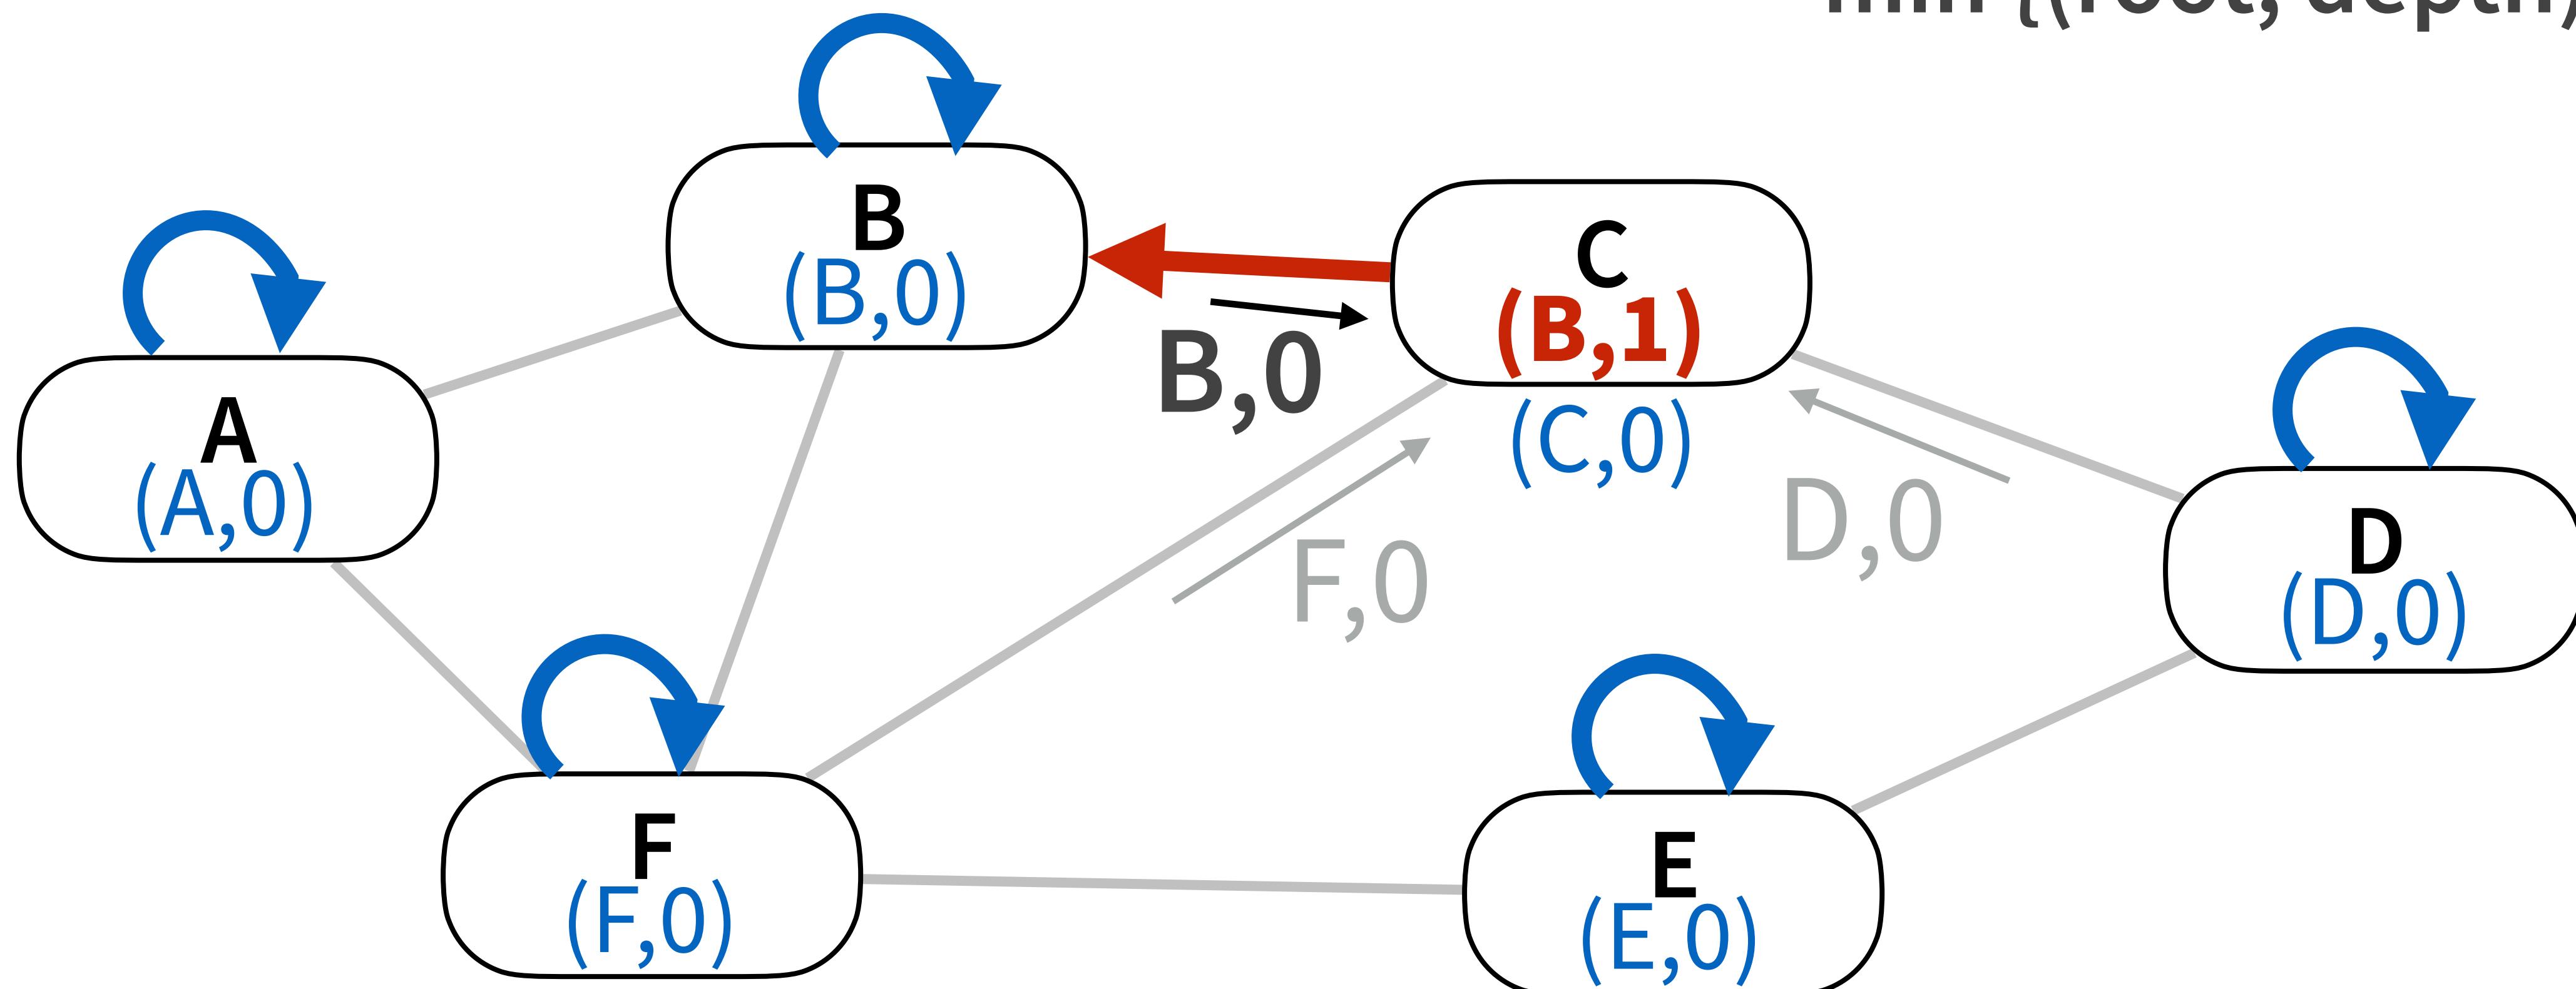


Illustration: BFS tree

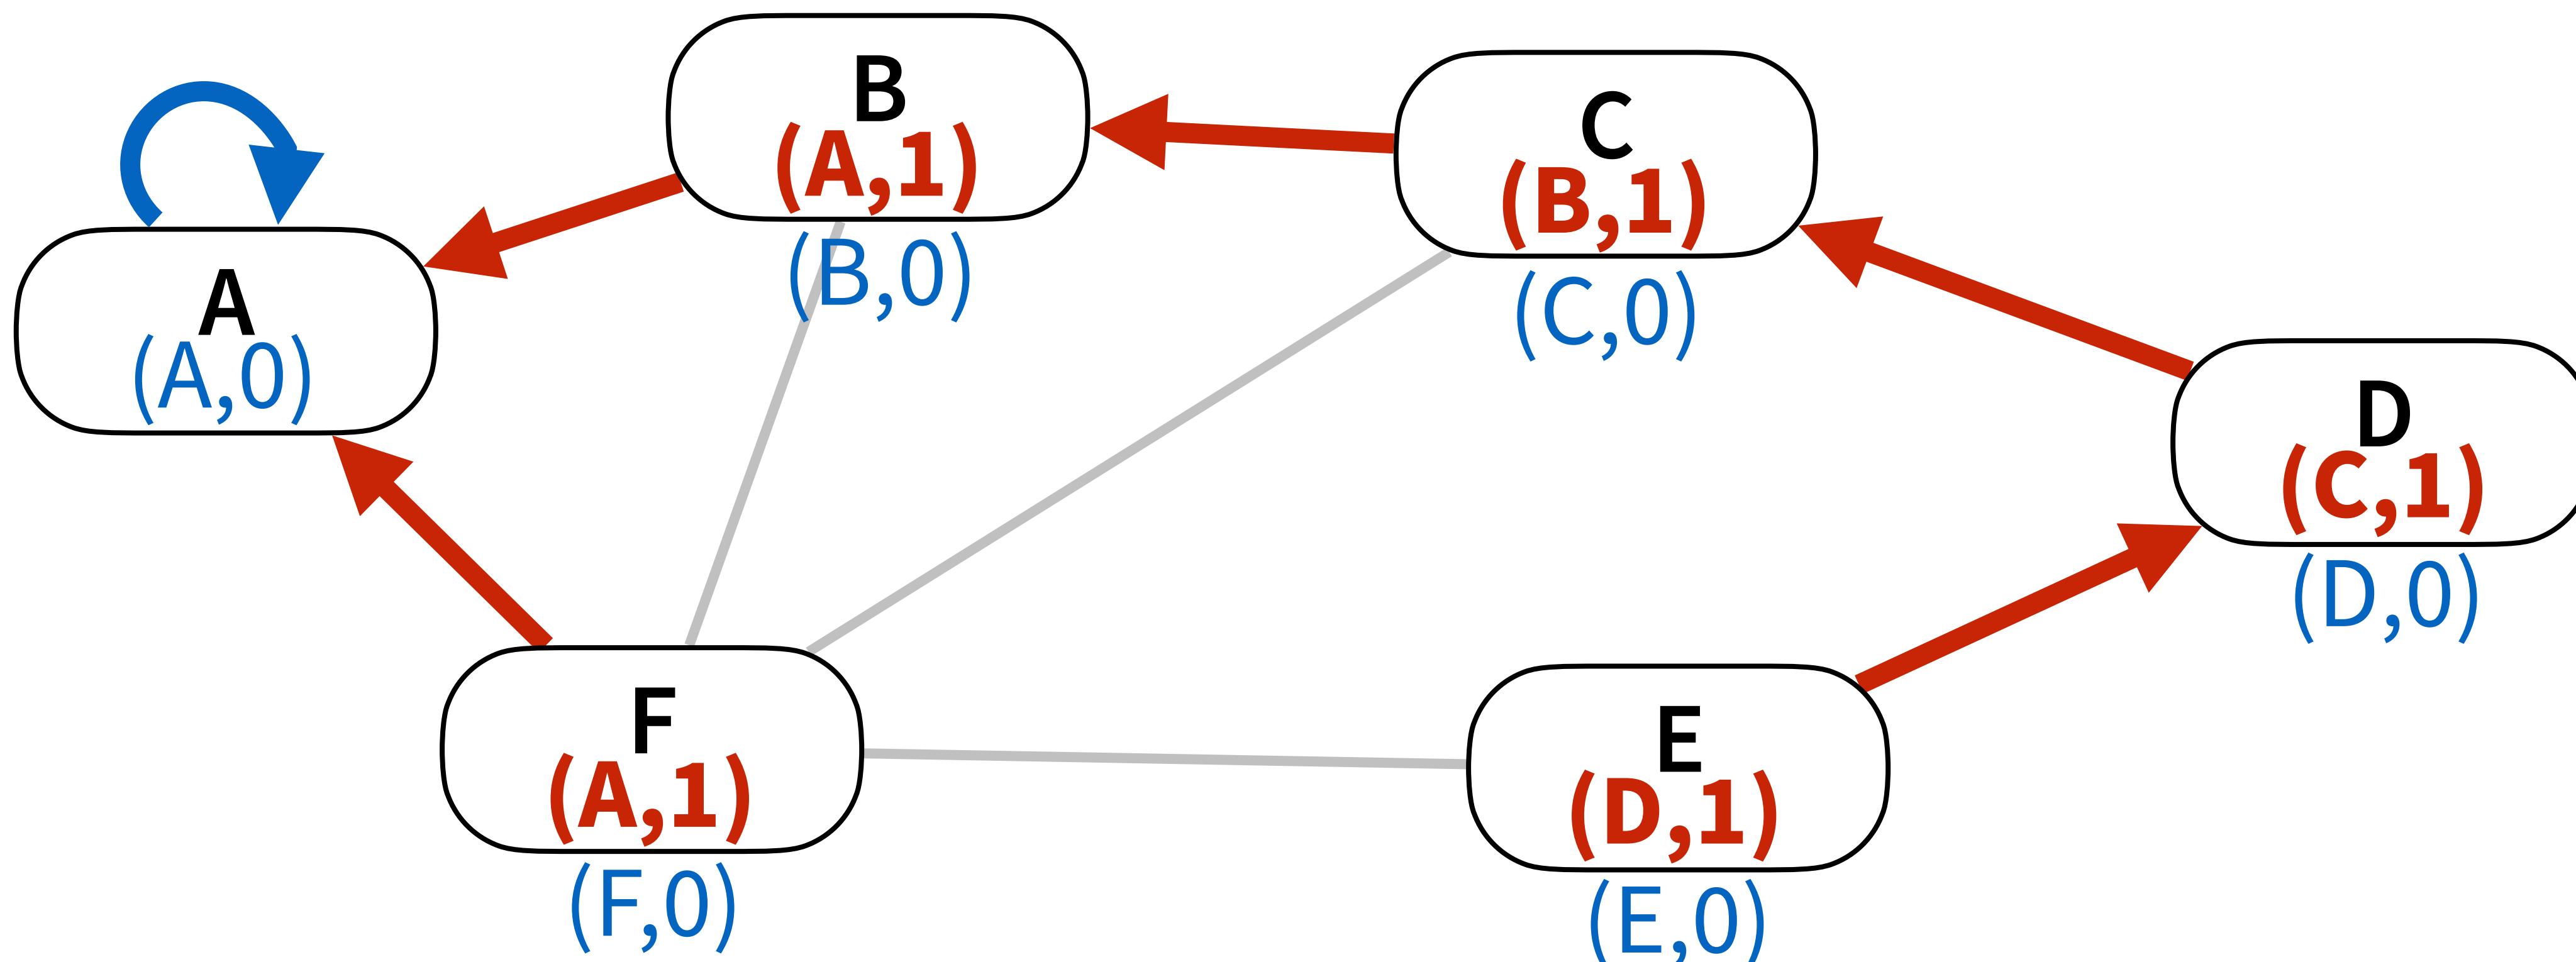


Illustration: BFS tree

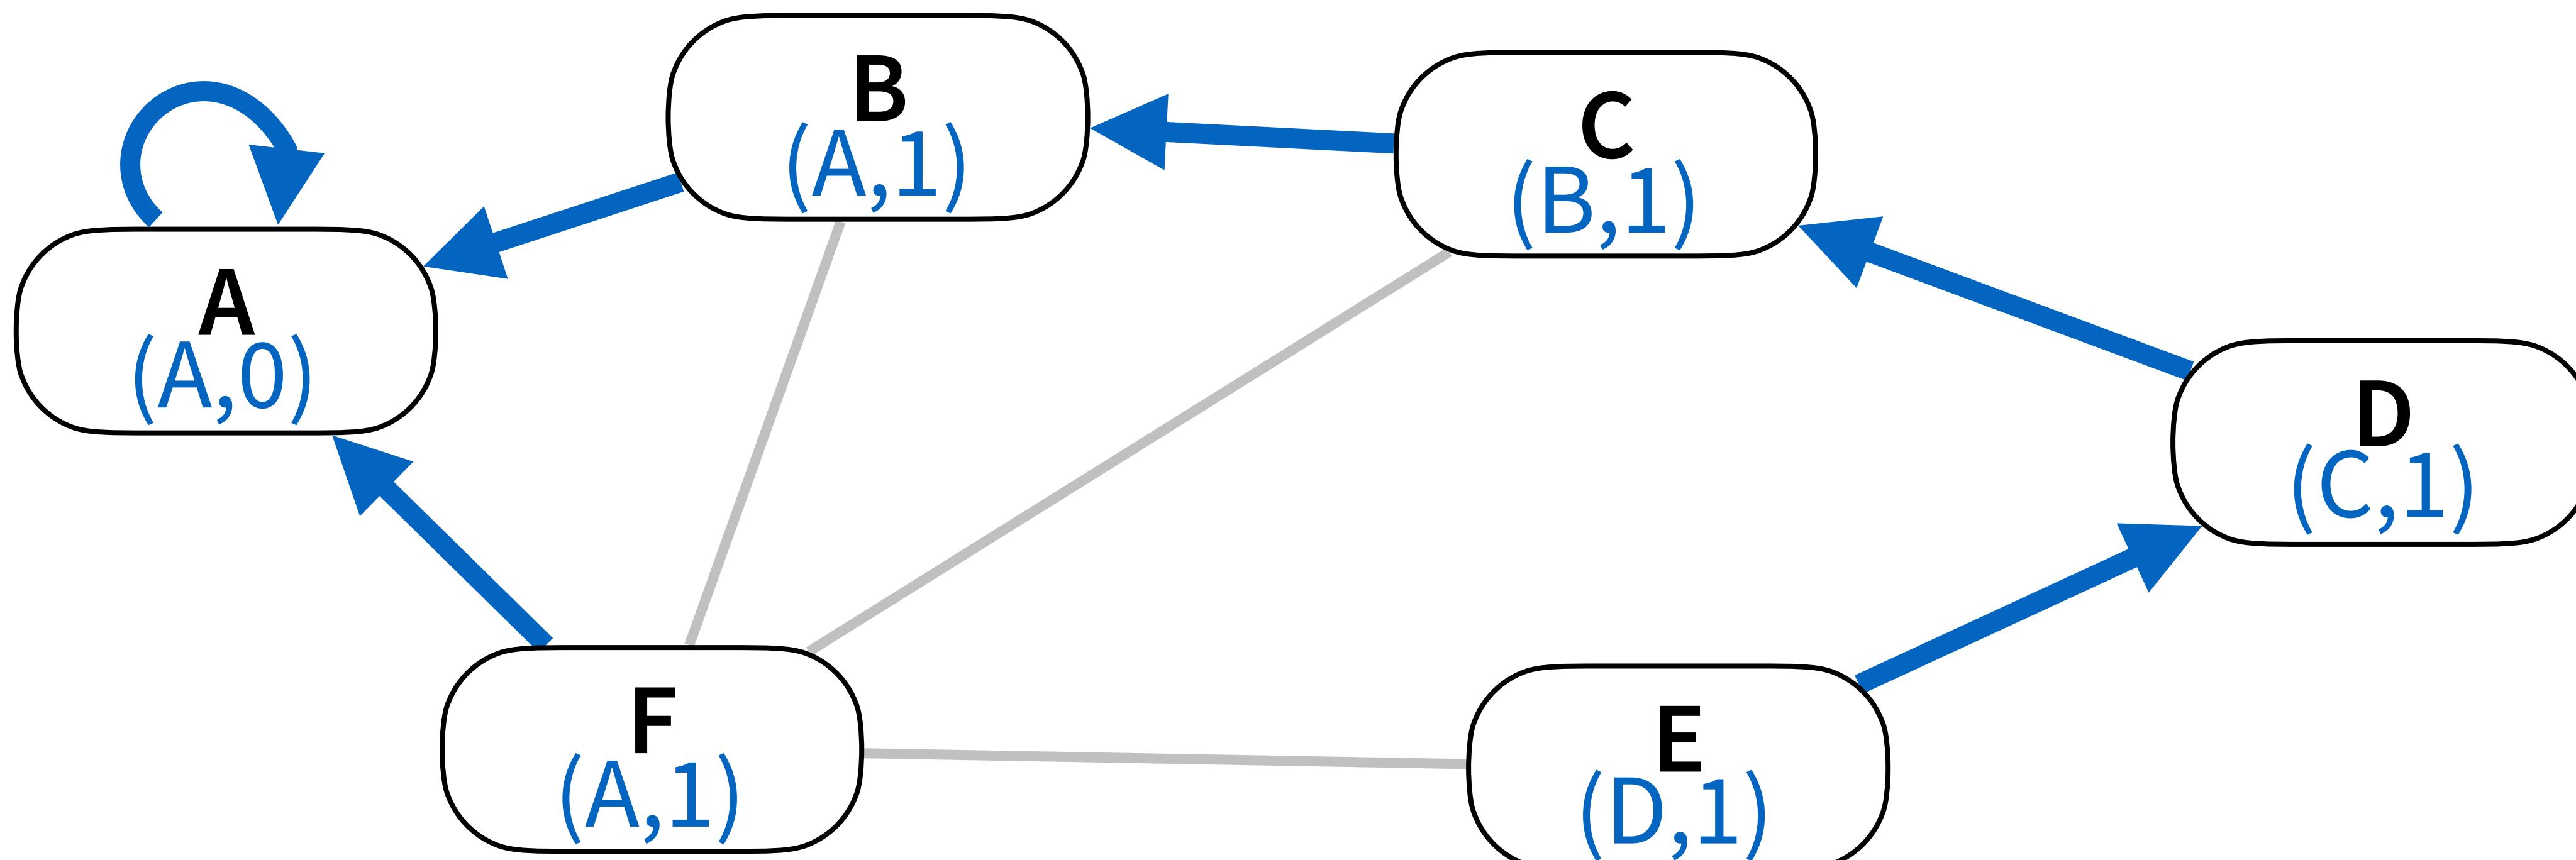


Illustration: BFS tree

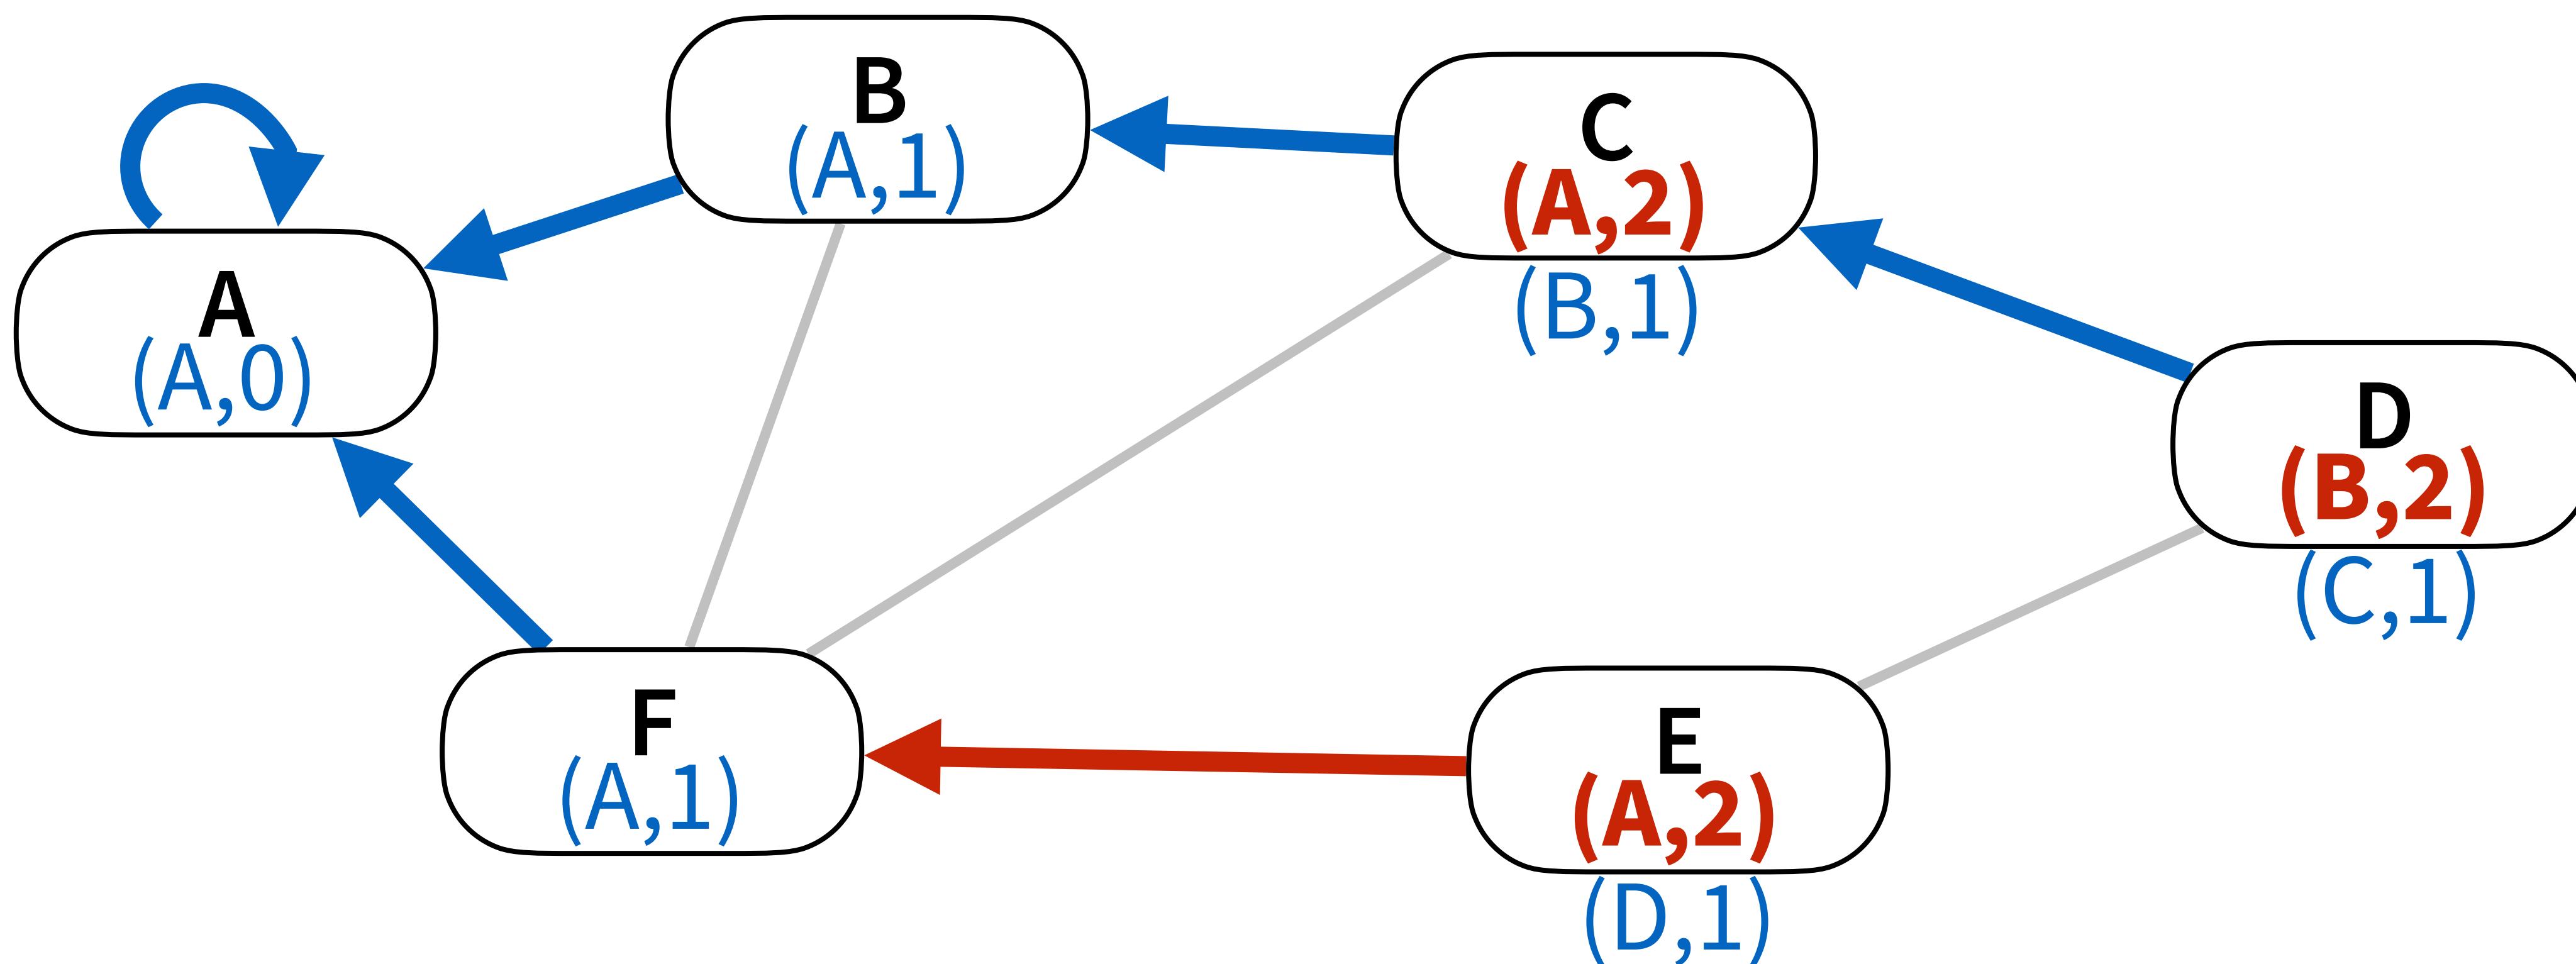


Illustration: BFS tree

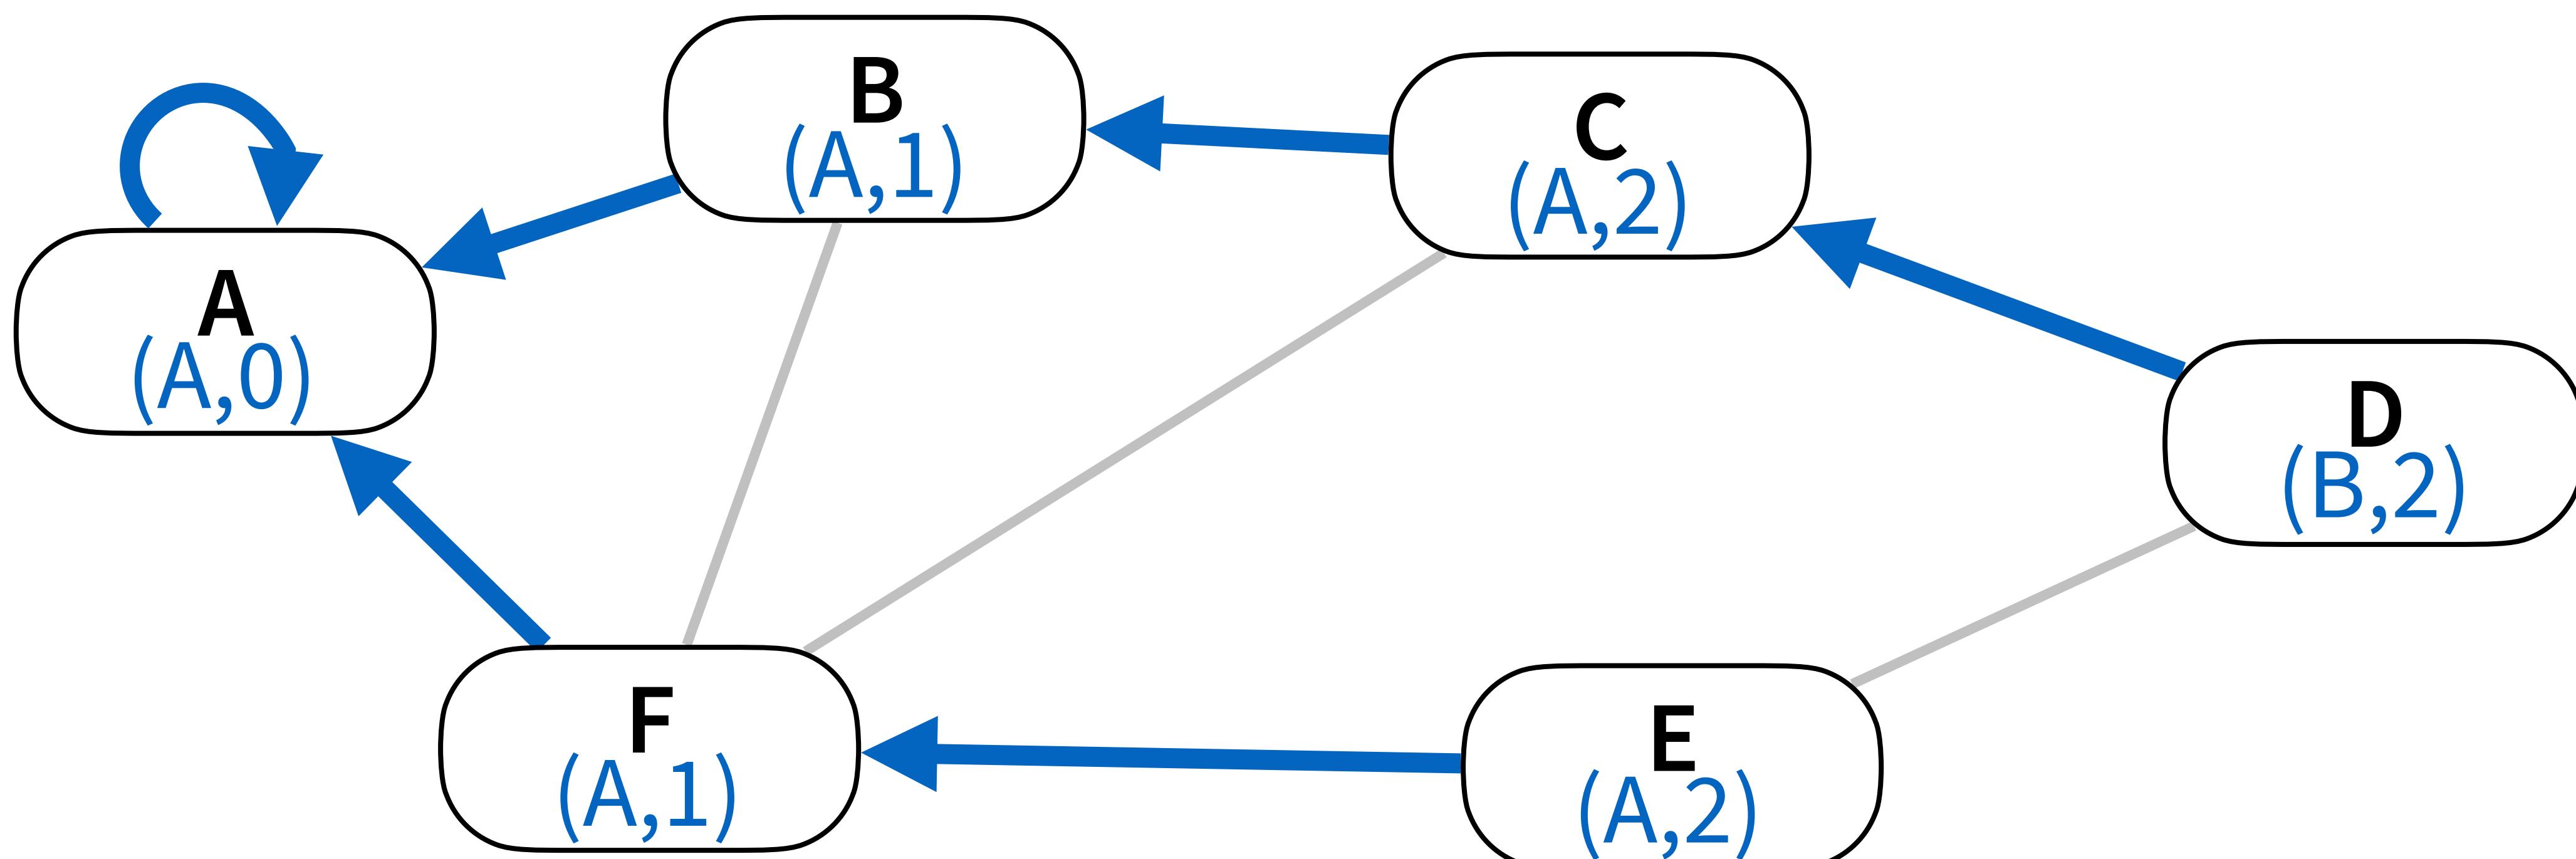


Illustration: BFS tree

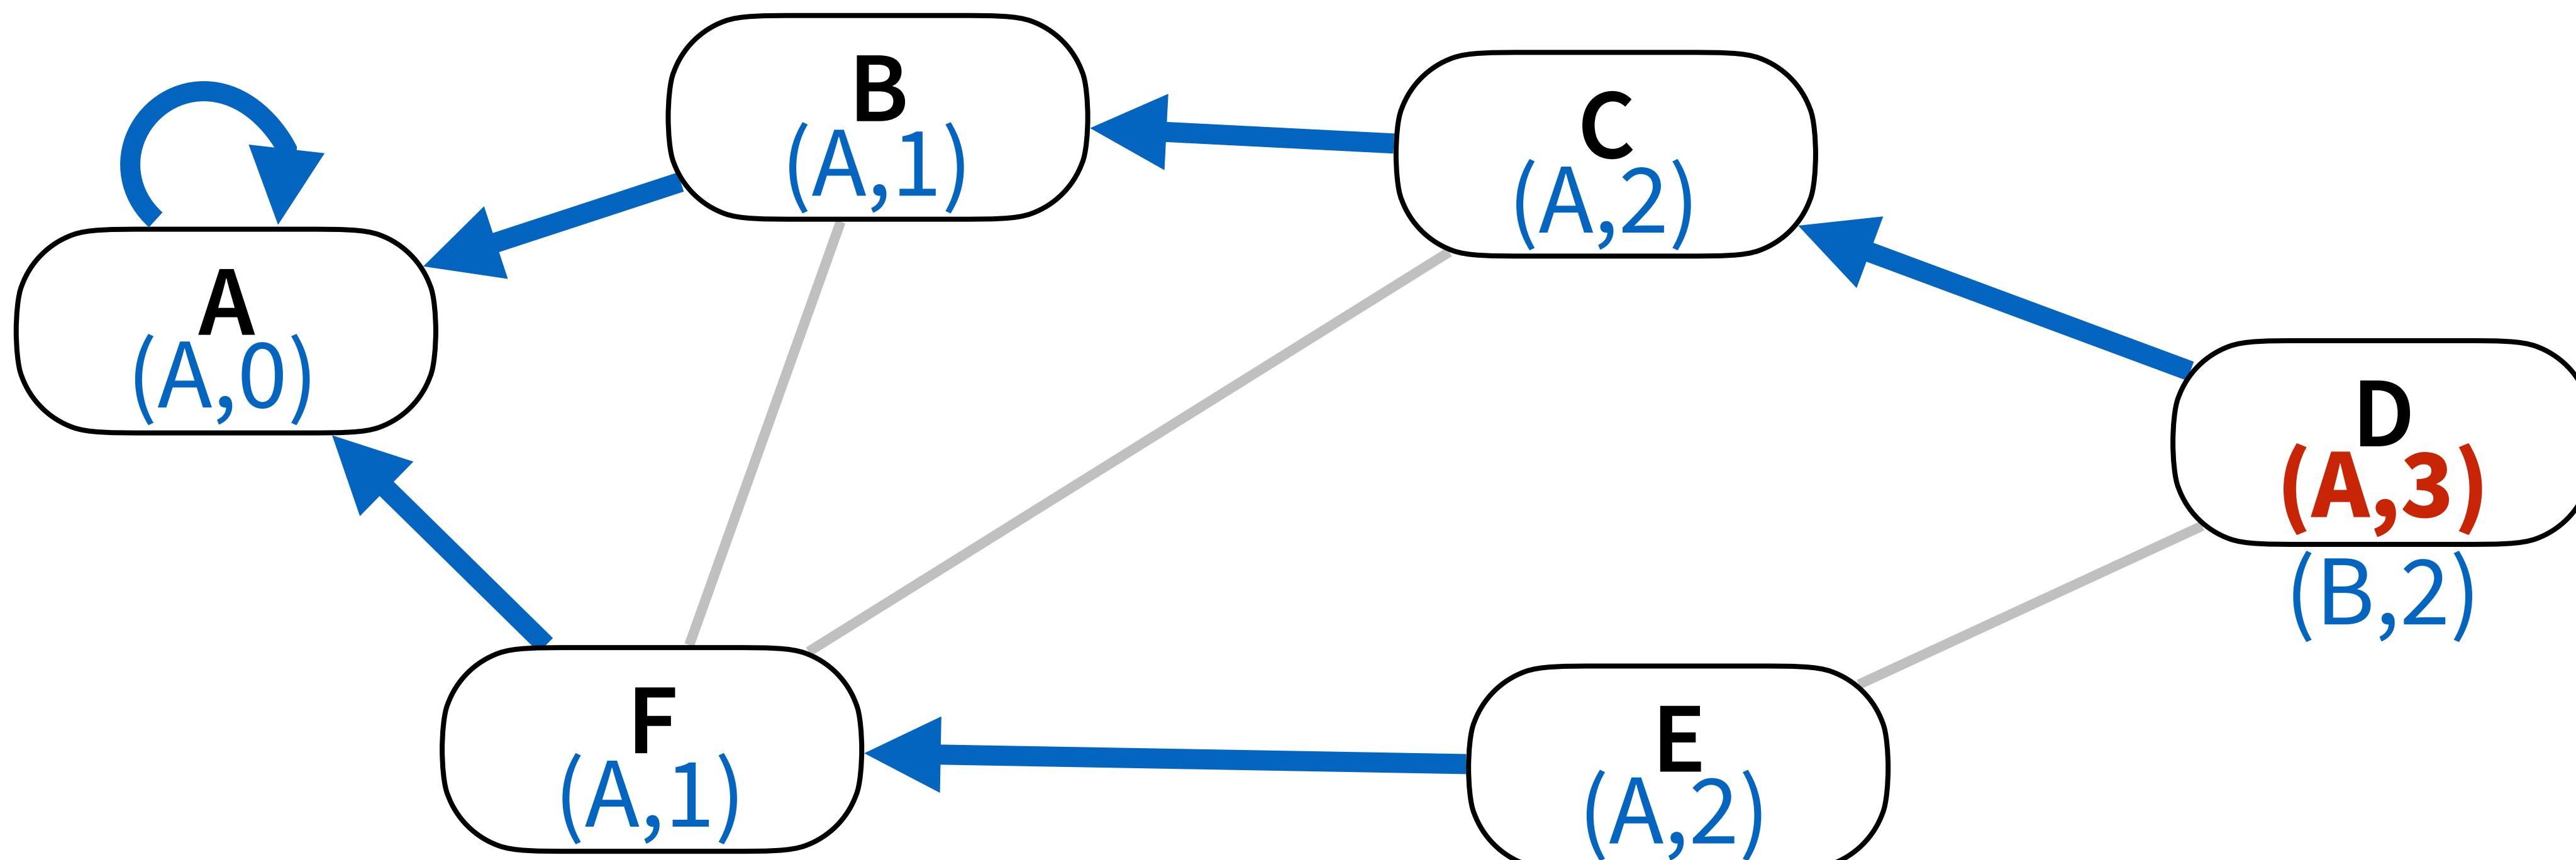


Illustration: BFS tree

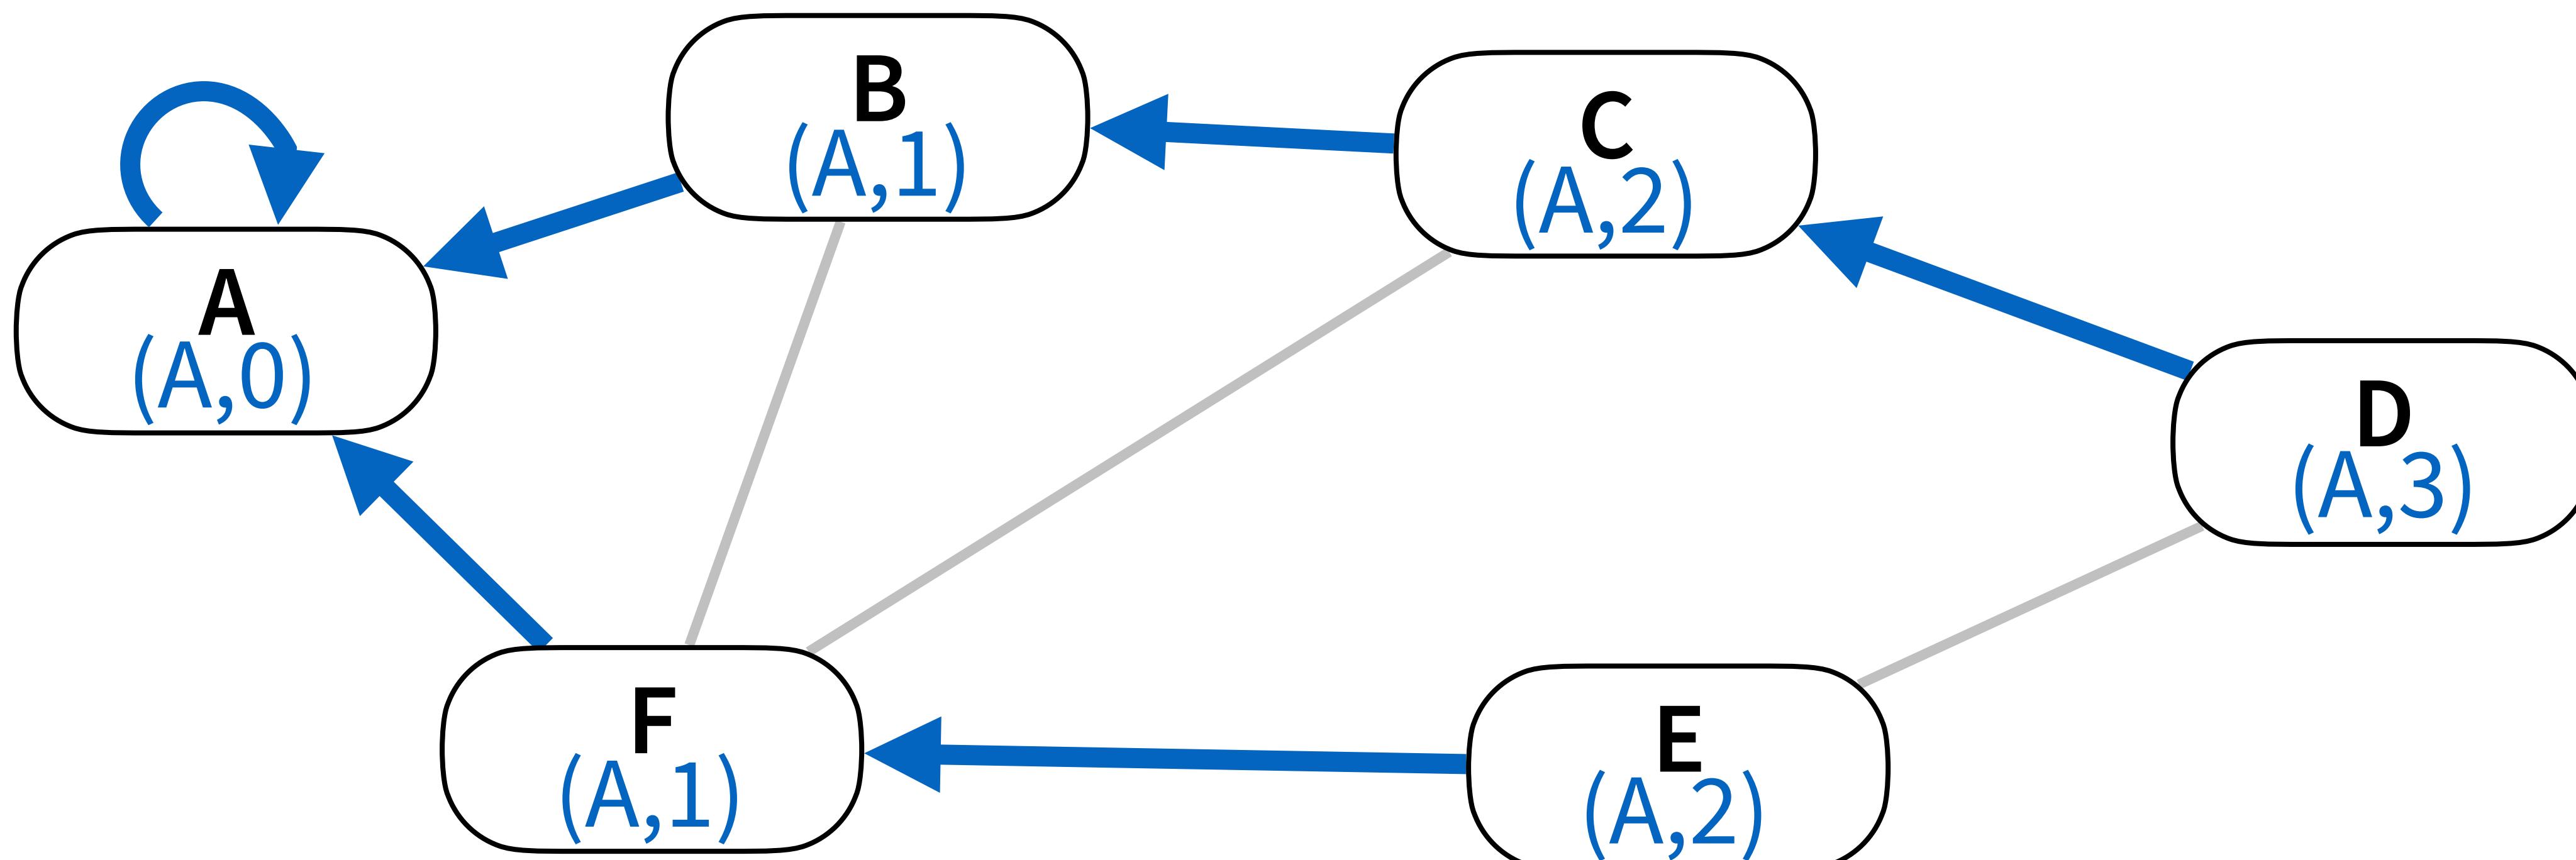


Illustration: BFS tree

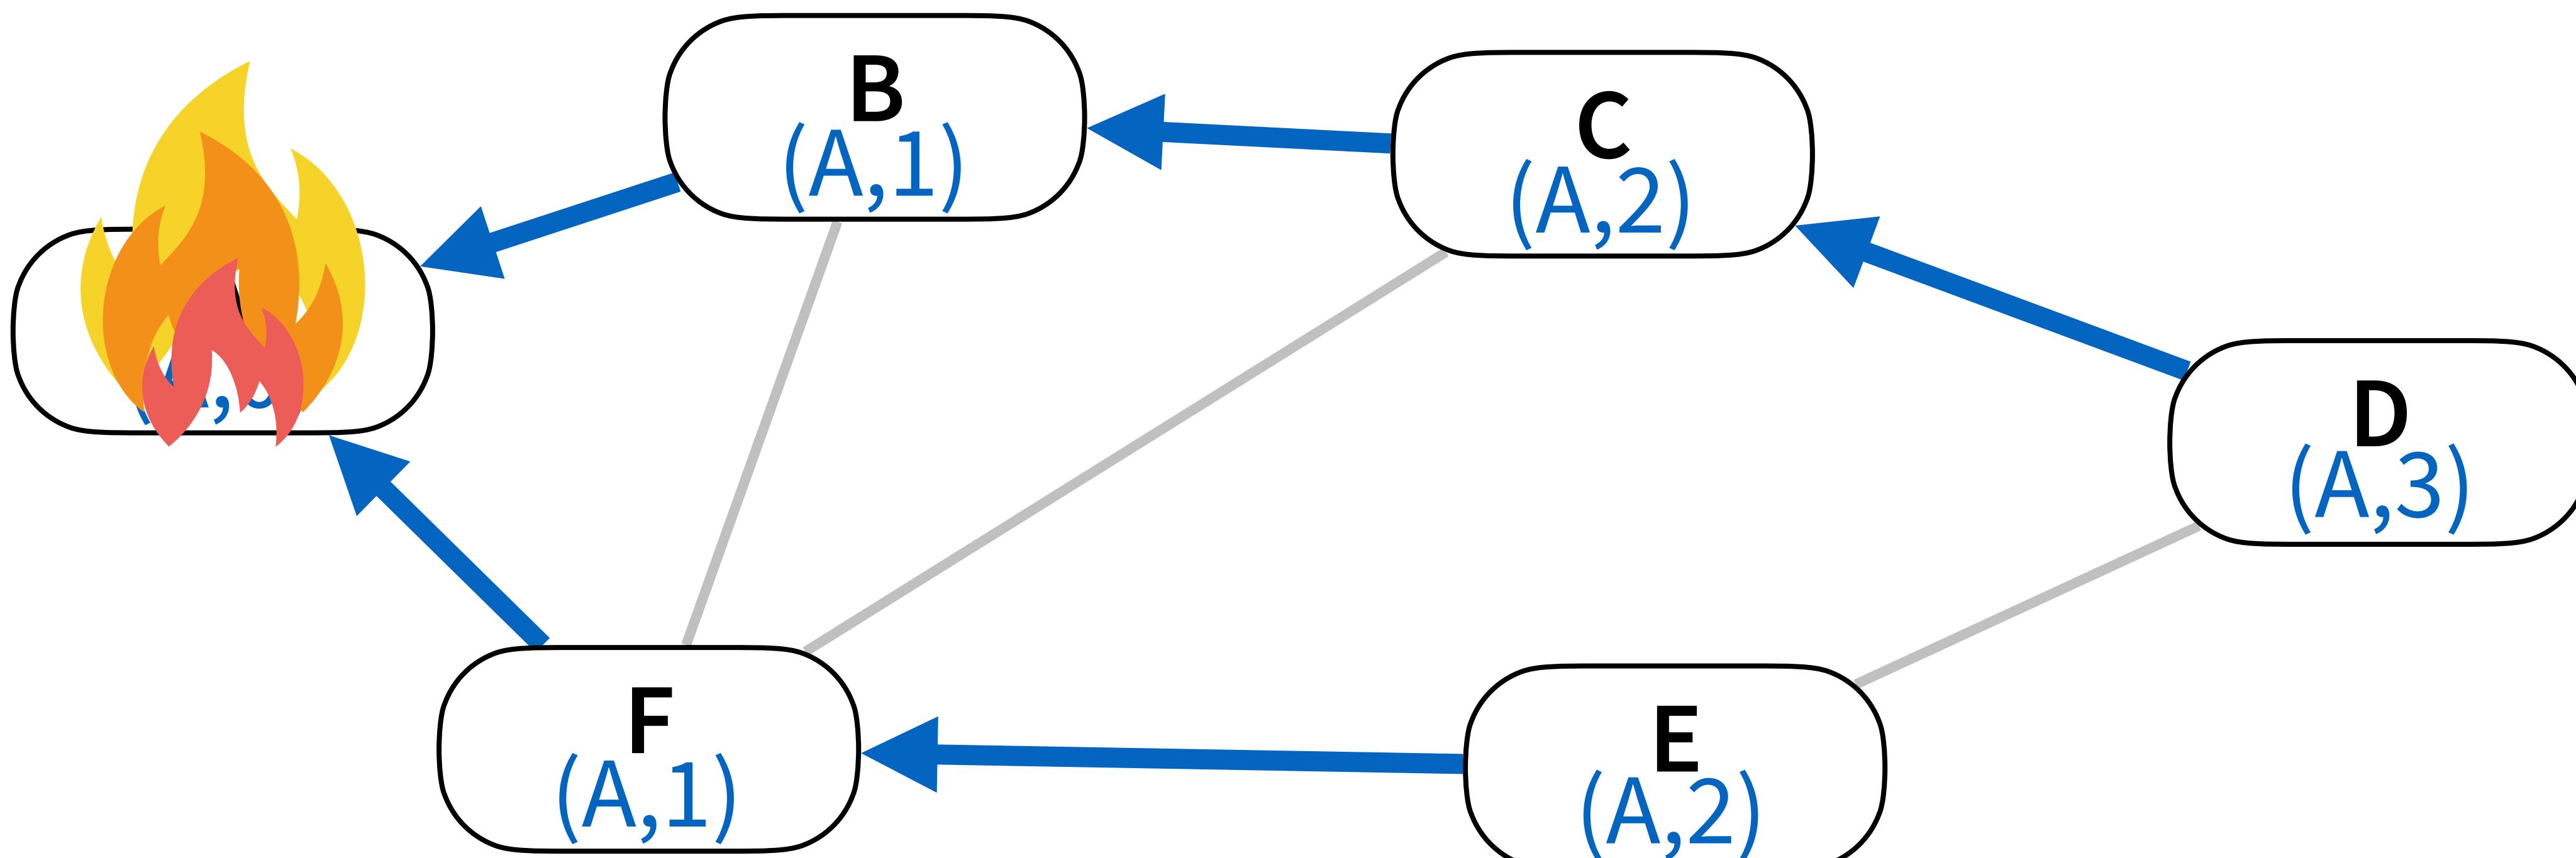


Illustration: BFS tree

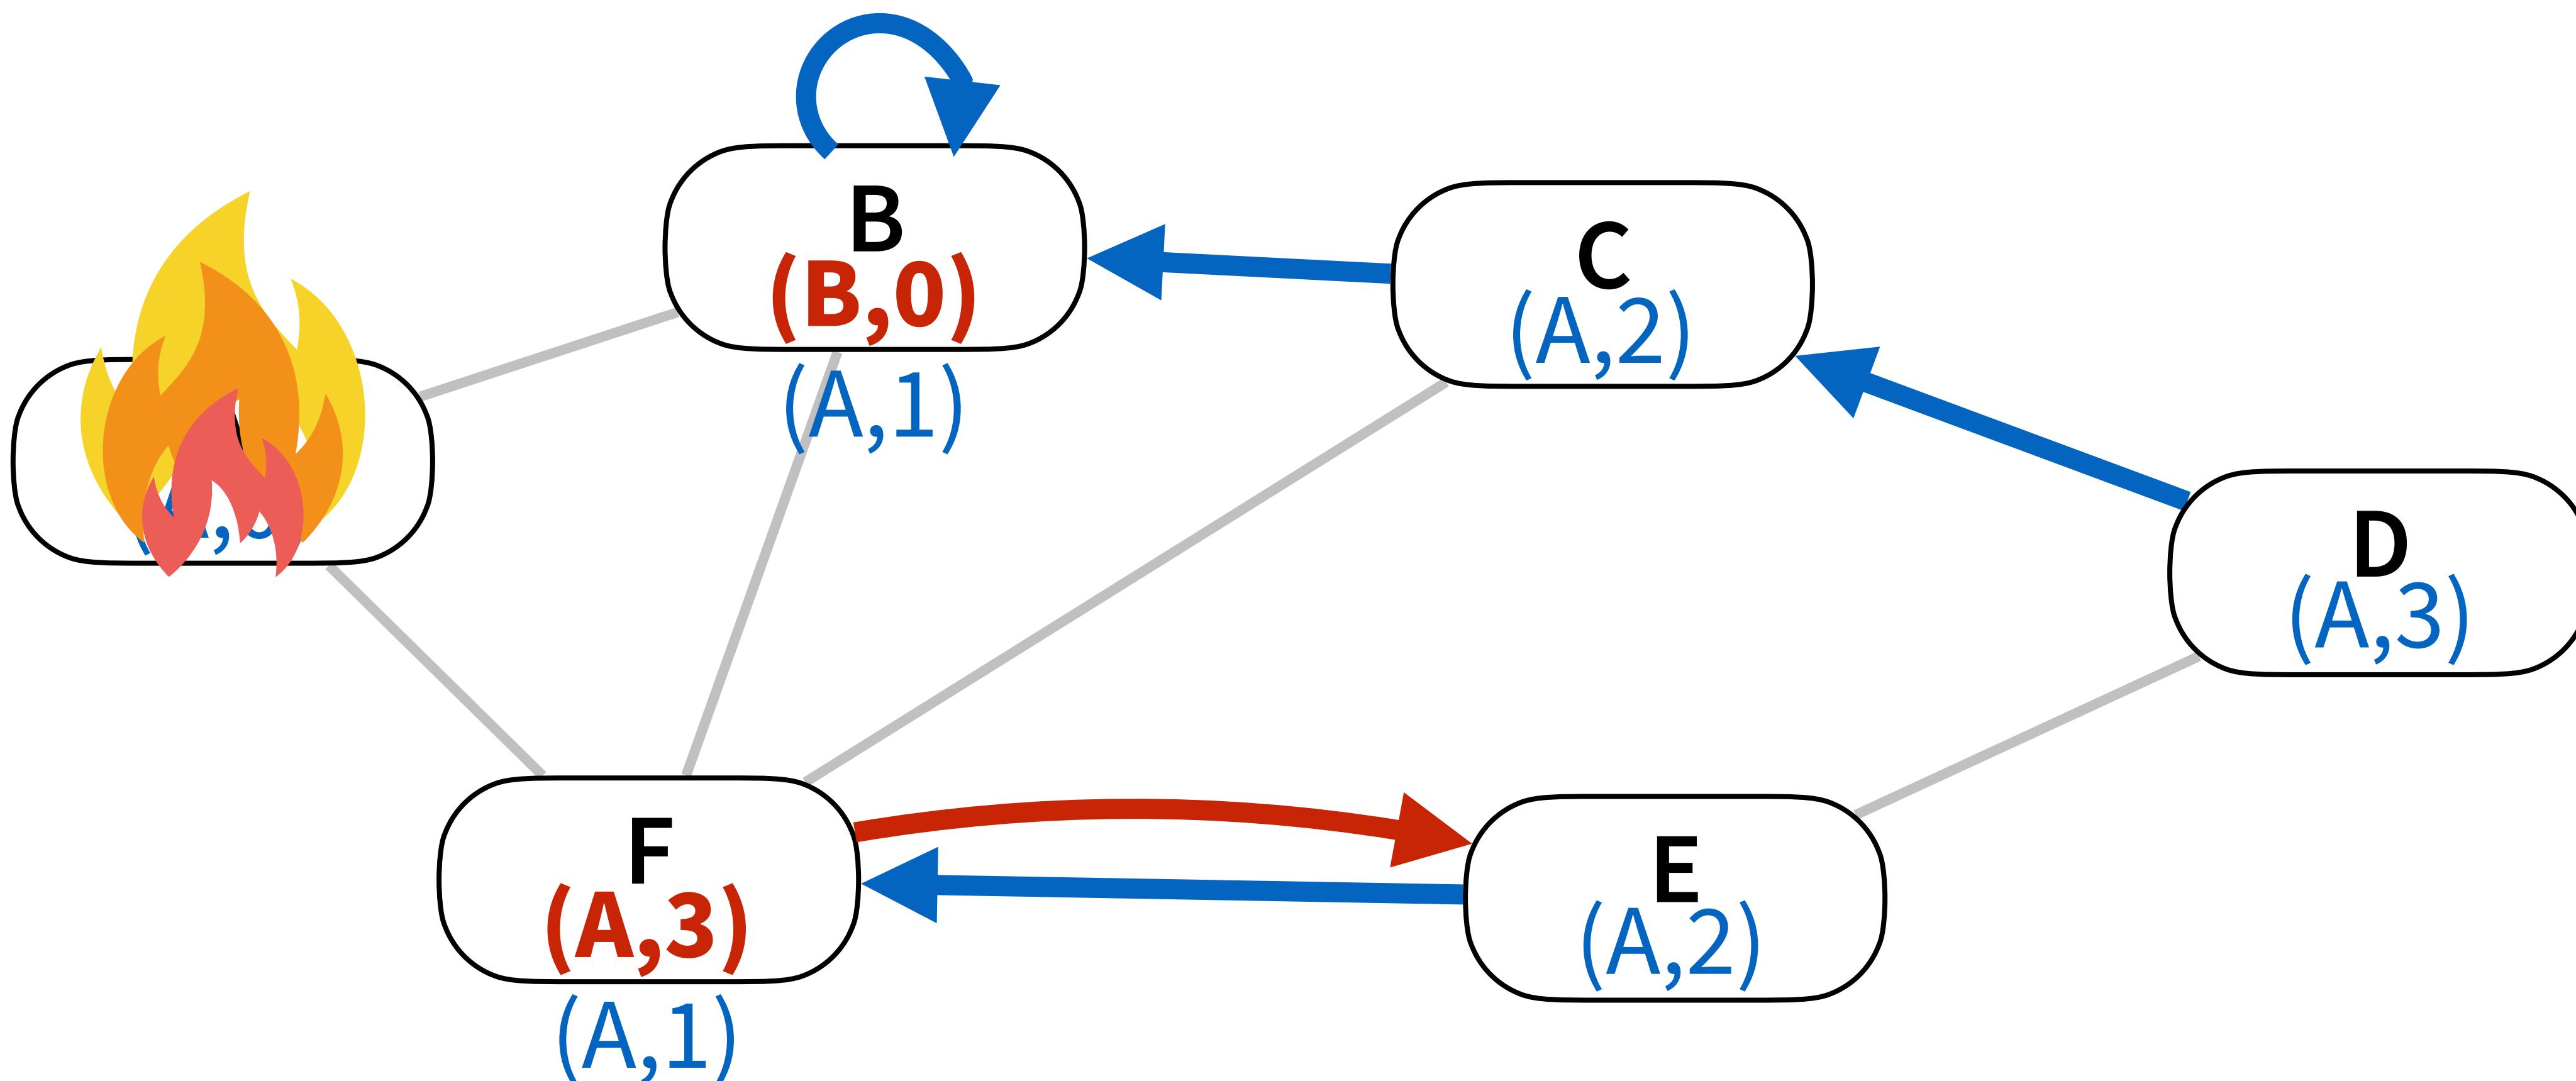


Illustration: BFS tree

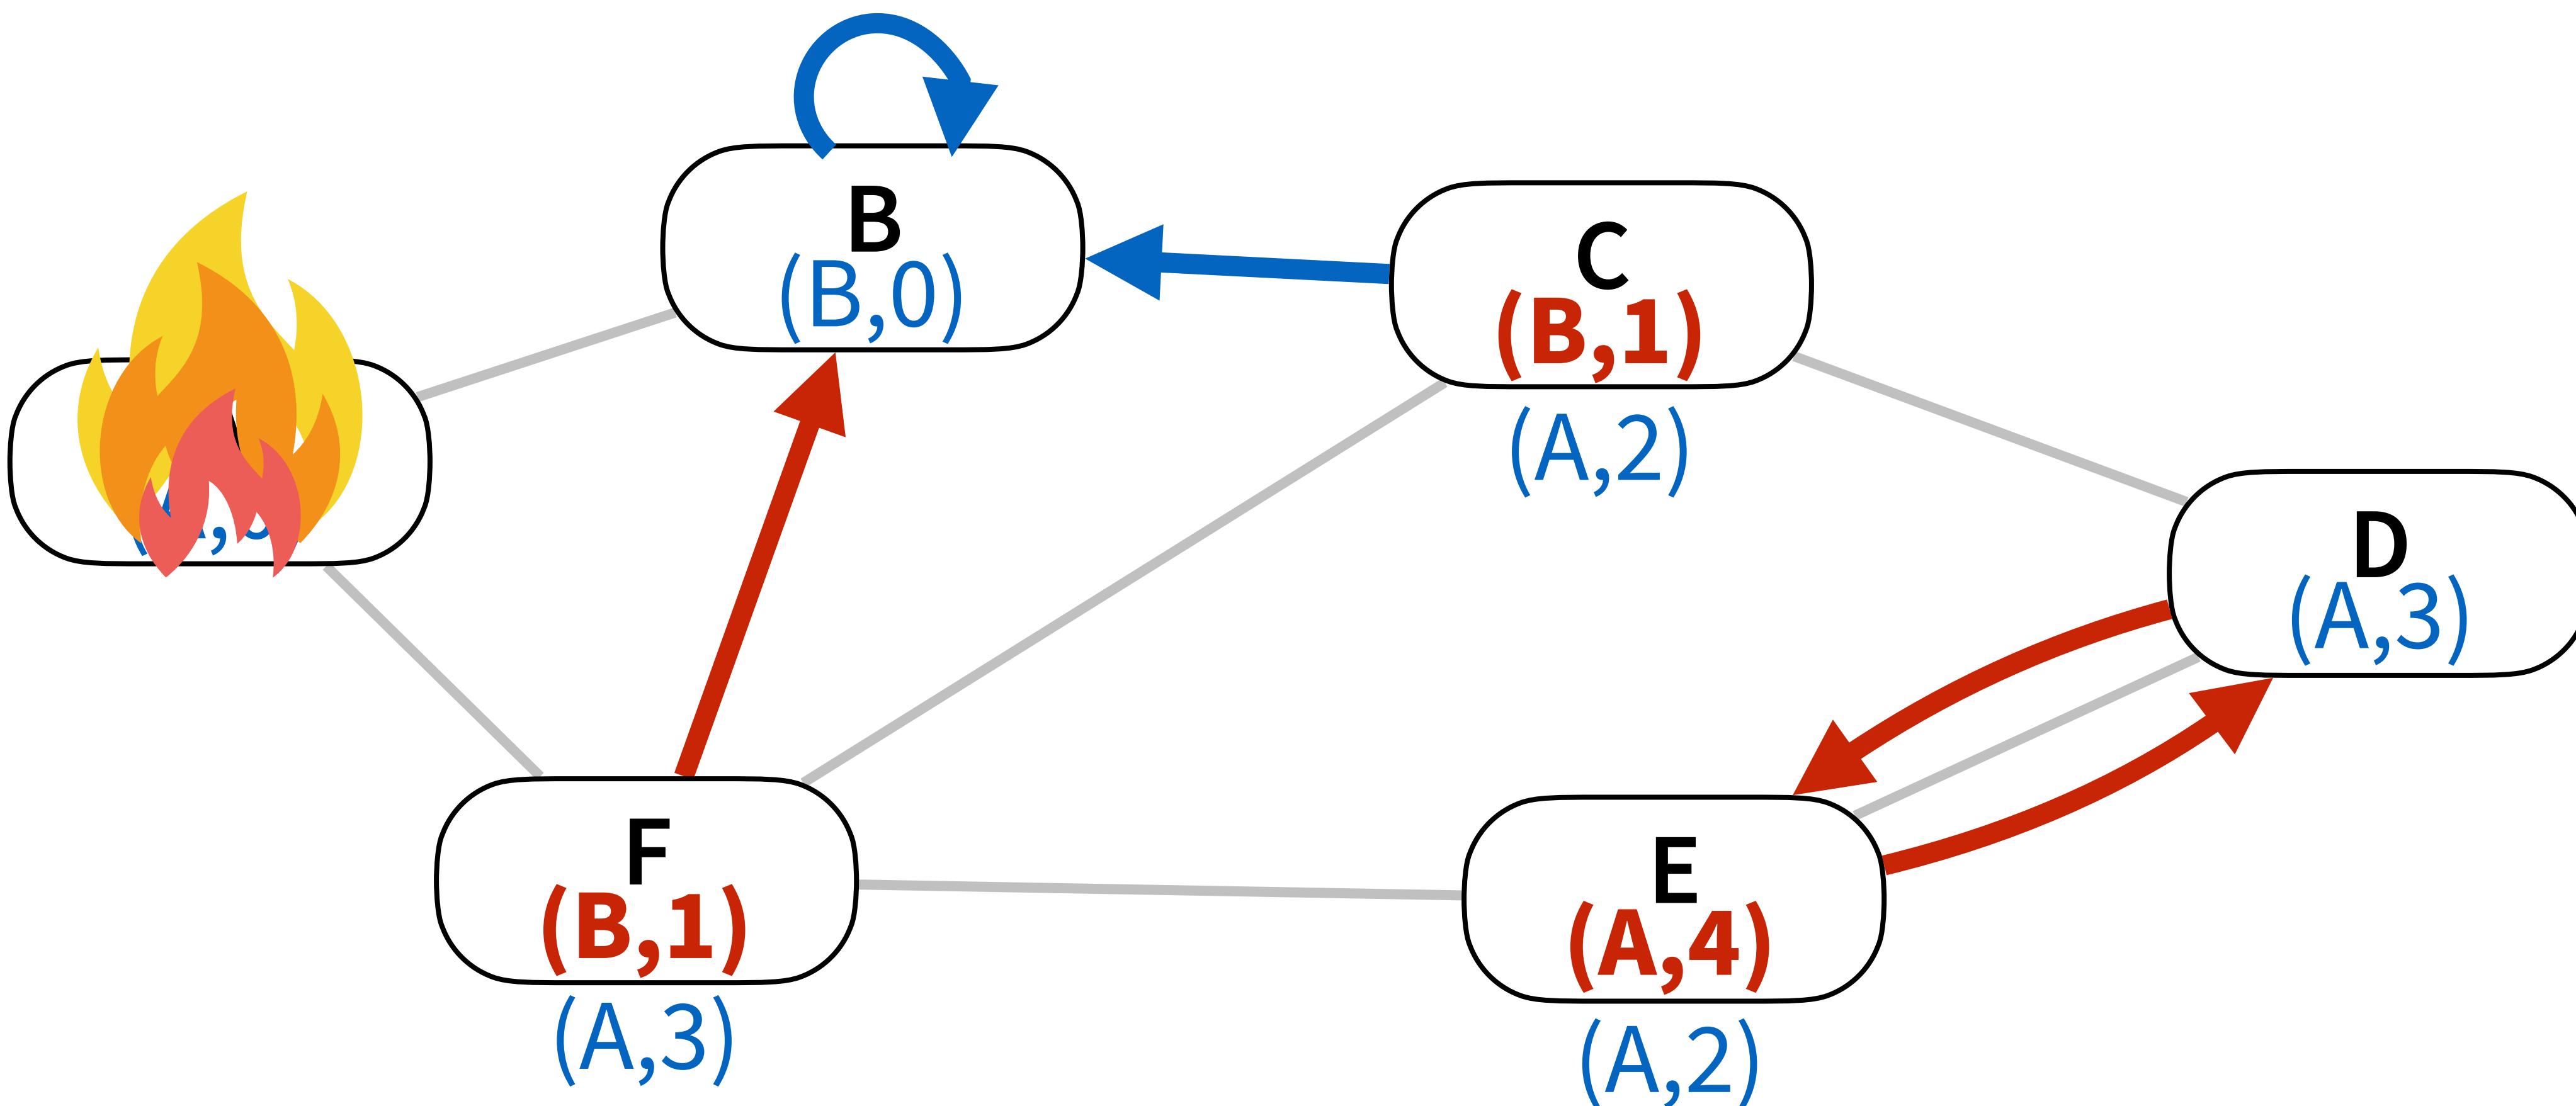


Illustration: BFS tree

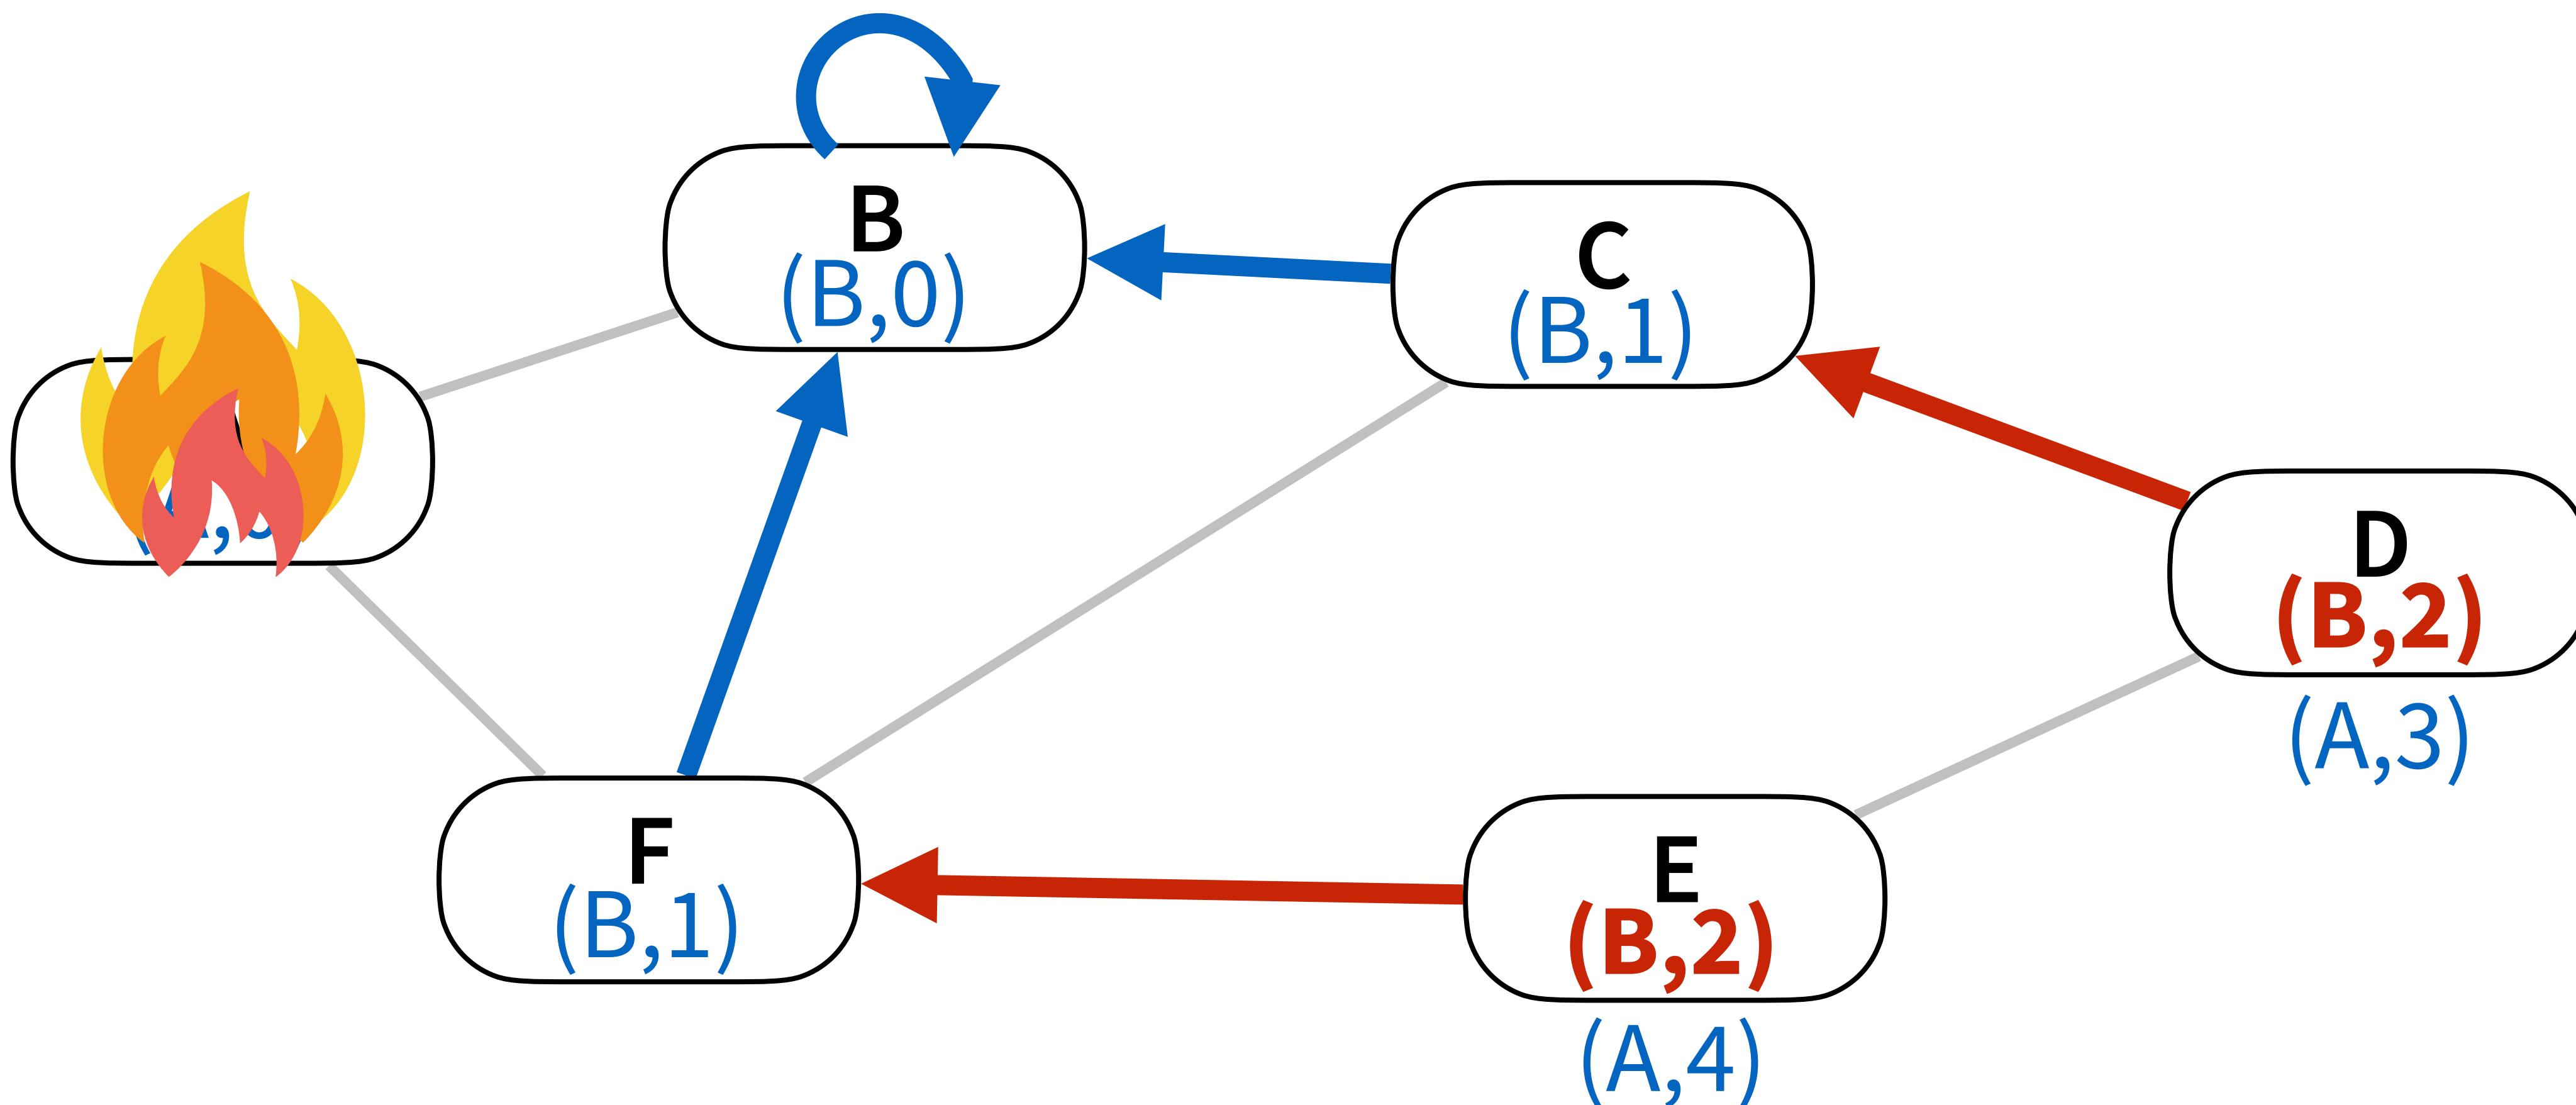
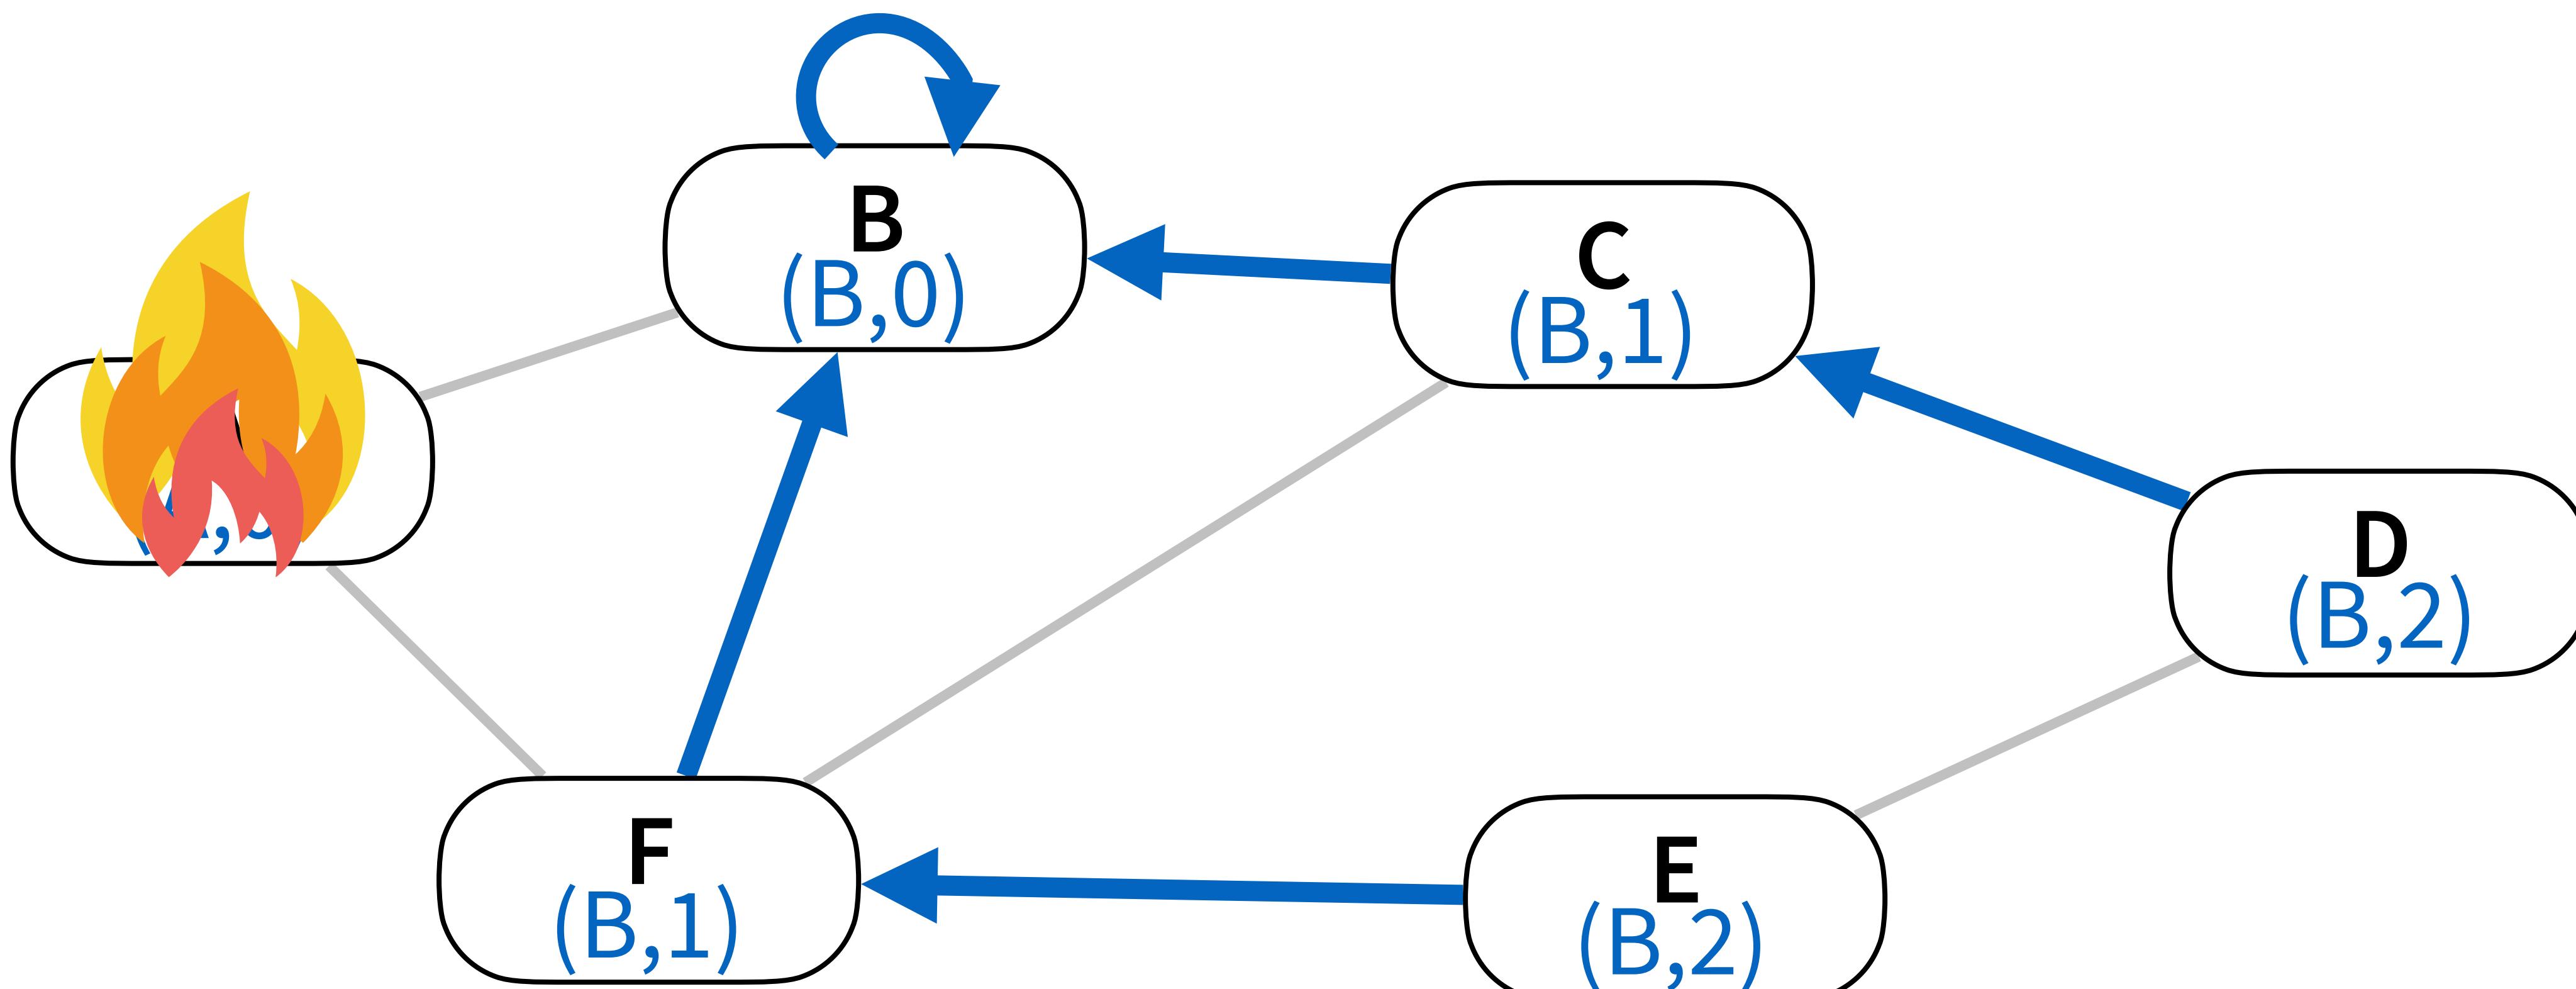


Illustration: BFS tree



Wrapup

Summary

▶ **Intrusion-Tolerance**

- ▶ work in spite of traitors
- ▶ covers for long detection time
- ▶ complementary with other approaches (security in depth)

▶ **Distributed Agreement**

- ▶ cornerstone
- ▶ deal with potential traitors

▶ **Byzantine agreement**

- ▶ robust but expensive

▶ **Blockchain**

- ▶ weak but flexible

▶ **Self-stabilization**

- ▶ attractive property
- ▶ non-masking fault-tolerance