

Verifying the Liveness of Eventually-Synchronous BFT Consensus Protocols

A pragmatic approach

Giuliano Losa
Stellar Development Foundation

Motivation: checking that the TetraBFT consensus protocol satisfies its liveness property

TetraBFT (PODC 2024) is a new BFT consensus protocol with a subtle liveness argument, and we would like to mechanically verify it



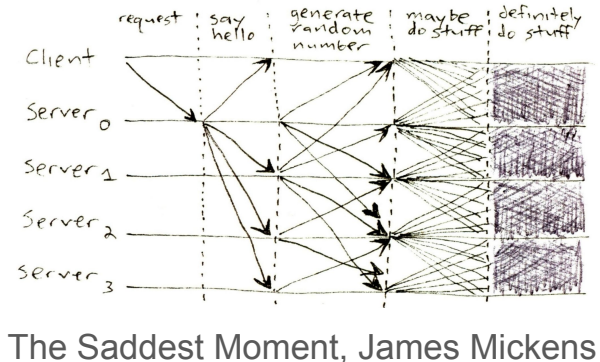
Qianyu Yu
HKUST



Giuliano Losa
SDF



Xuechao Wang
HKUST




Due to Byzantine behavior, state-space exploration with TLC is intractable

TLC explicitly enumerates bounded behaviors

This is a fully-automated, push-button solution

But it is too slow, because Byzantine behavior creates too many successors to each state



Given a user-provided inductive invariant, Apache can efficiently check that a state predicate always holds

An inductive invariant holds initially and that every step preserves it:

$$\wedge \text{Init} \Rightarrow \text{Inv}$$

$$\wedge \text{Inv} \wedge \text{Next} \Rightarrow \text{Inv}'$$

For bounded domains, this can be encoded into SMT formulas, which Apache passes to the Z3 solver for efficient checking.

We obtain a semi-automated proof method for verifying that a state predicate P always holds:

1. The user provides a candidate inductive invariant I
2. Apache checks that I is indeed inductive and that $I \Rightarrow P$
 - a. If not, go back to 1

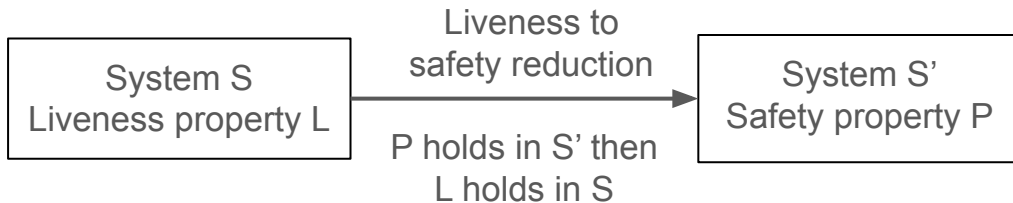
To use inductive invariants for liveness, we propose a simple liveness-to-safety reduction

Safety VS Liveness

- Safety: all reachable states satisfy P
- Liveness: fairness conditions + temporal property of behaviors
 - Example: if every message is eventually delivered, nodes eventually decide

Liveness to safety in practice

- Some tools produce a new safety problem automatically
 - Works best in conjunction with automated checking methods
 - Manually finding an inductive invariants for an automatically-generated safety problem is not fun! (and difficult)
- We give an “easy” reduction for the special case of eventually-synchronous BFT protocols



In this talk, we explain the liveness-proof methodology using the Paxos consensus algorithm

System model

- We have a fixed set of nodes (or “acceptors”) in a message-passing network
- The network is initially asynchronous, and becomes synchronous after time GST
- Local clocks that advance at the same rate
- Less than $\frac{1}{2}$ of the nodes may crash-stop

The Consensus Problem

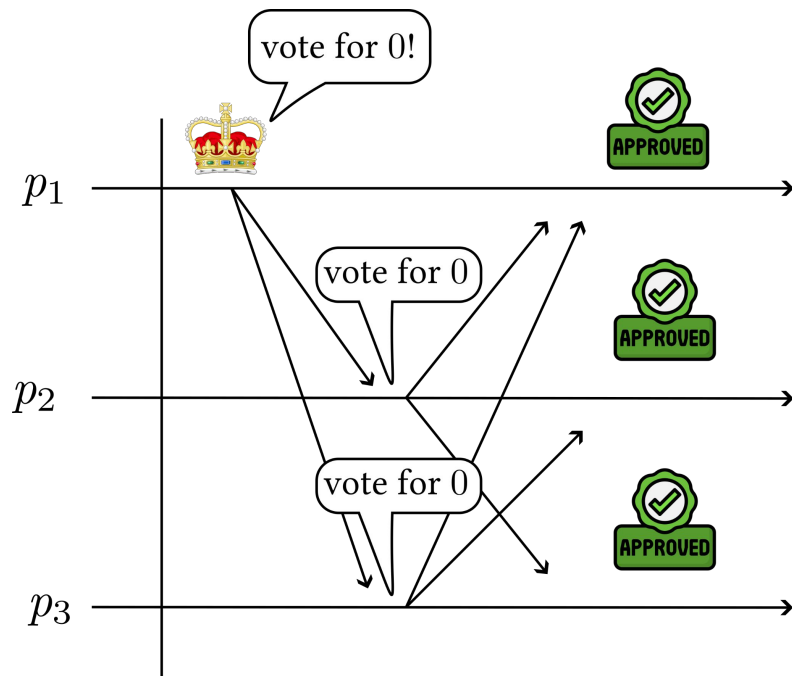
Nodes each start with a private value and must decide on a common value.

- Safety: no two well-behaved nodes decide differently (and, if all nodes are well-behaved, then they decide one of the inputs)
- Liveness: eventually, all well-behaved nodes decide

A leader instruct nodes to vote for a value


Nodes vote for the leader's value

A node decides a value when a majority of nodes vote for it




If the first leader fails to impose its decision, nodes time-out and try again with a new leader

Each such attempt is called a ballot



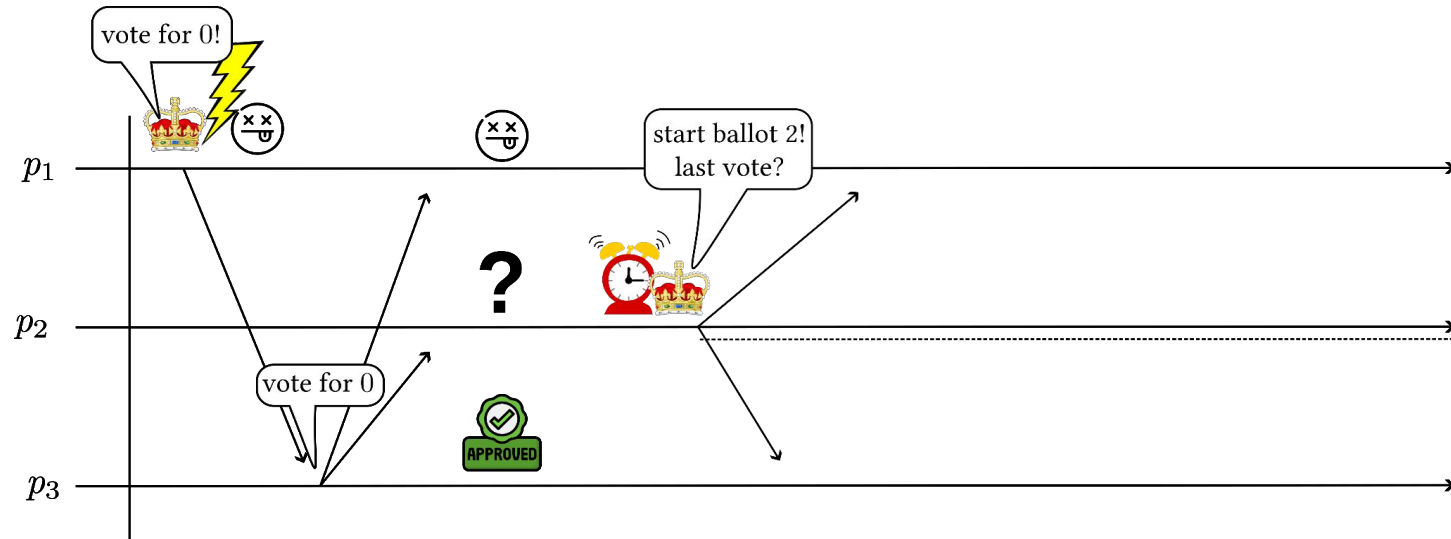
If the first leader fails to impose its decision, nodes time-out and try again with a new leader

Each such attempt is called a ballot



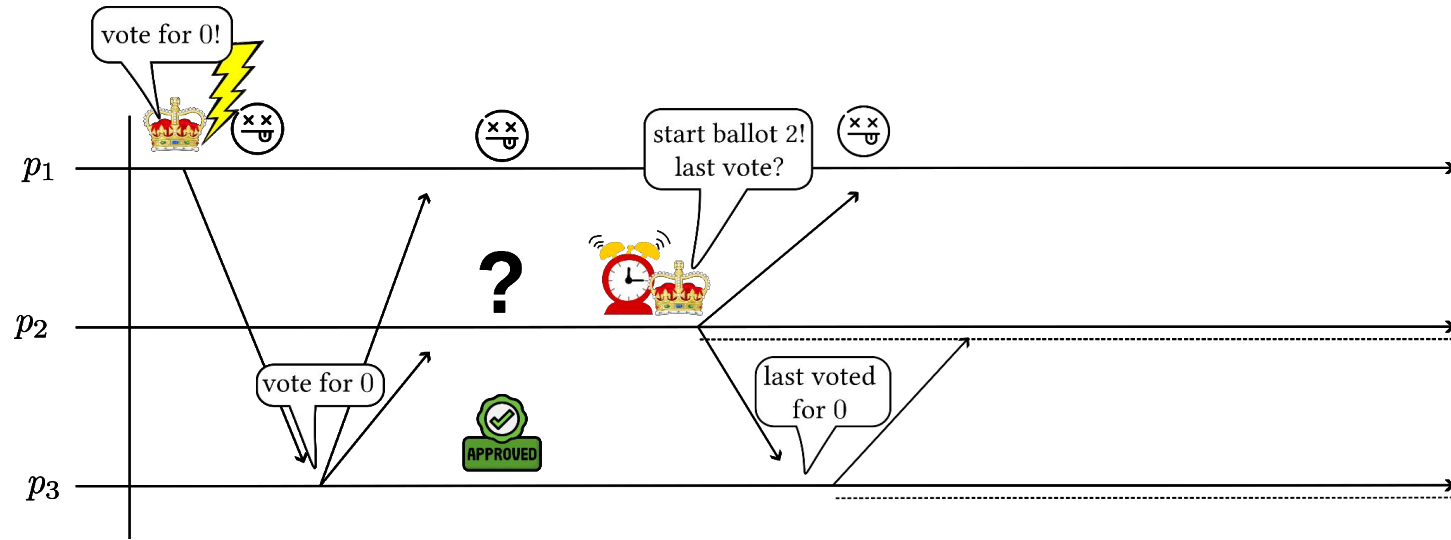
If the first leader fails to impose its decision, nodes time-out and try again with a new leader

Each such attempt is called a ballot



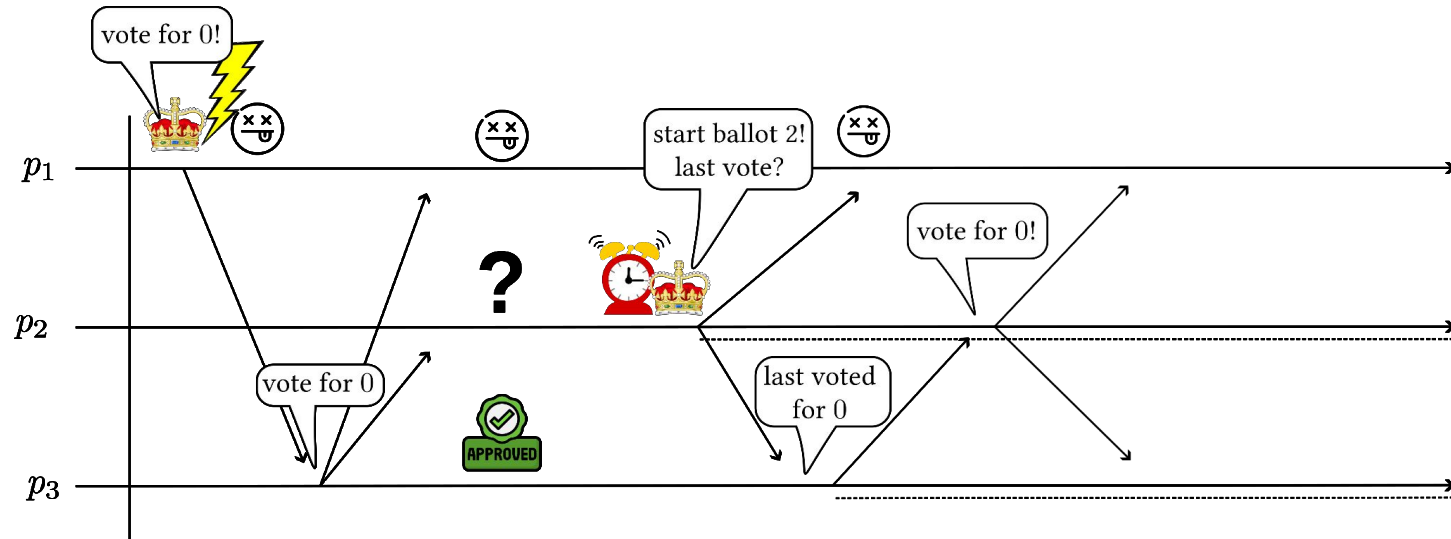
If the first leader fails to impose its decision, nodes time-out and try again with a new leader

Each such attempt is called a ballot



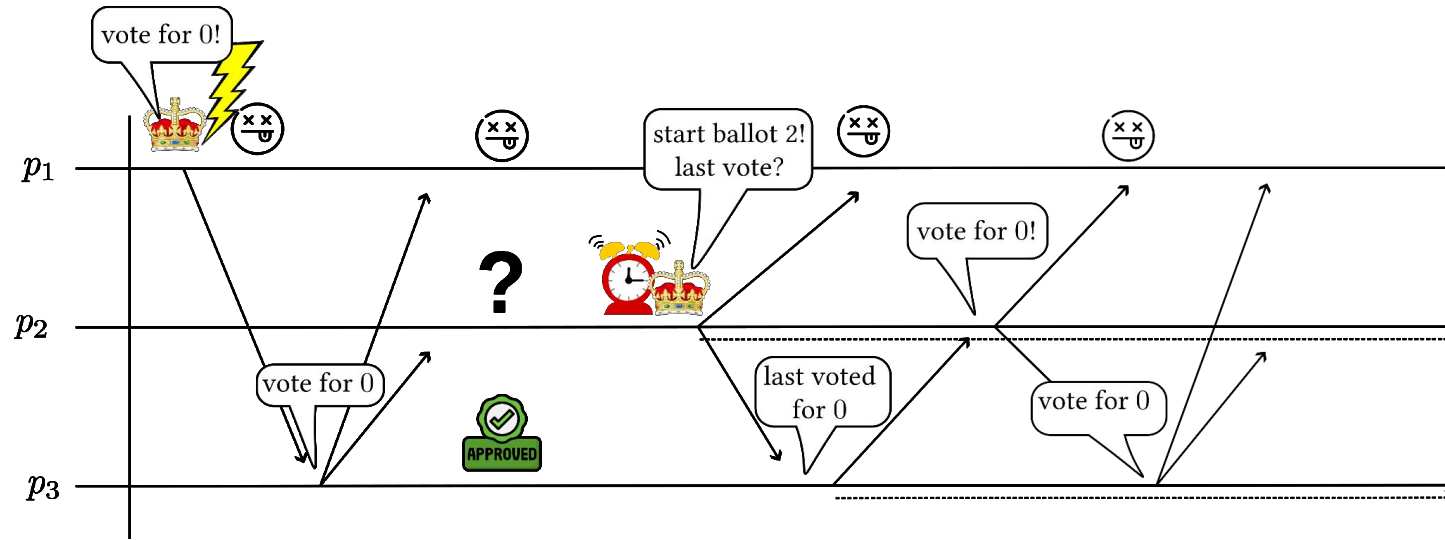
If the first leader fails to impose its decision, nodes time-out and try again with a new leader

Each such attempt is called a ballot



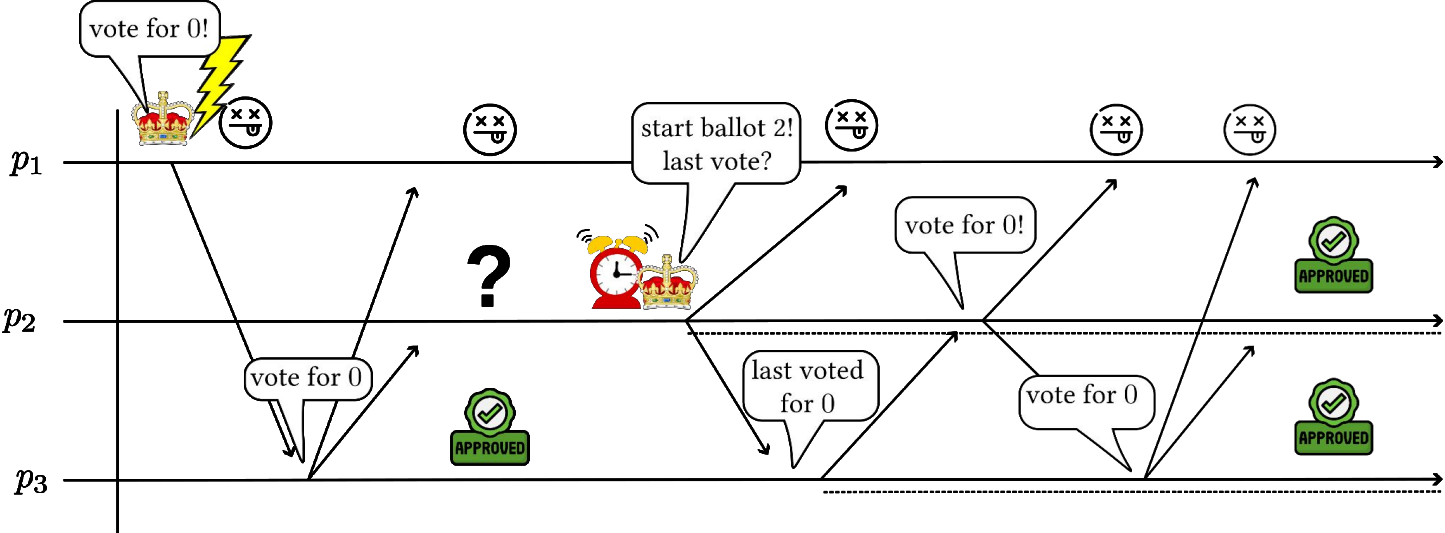
If the first leader fails to impose its decision, nodes time-out and try again with a new leader

Each such attempt is called a ballot

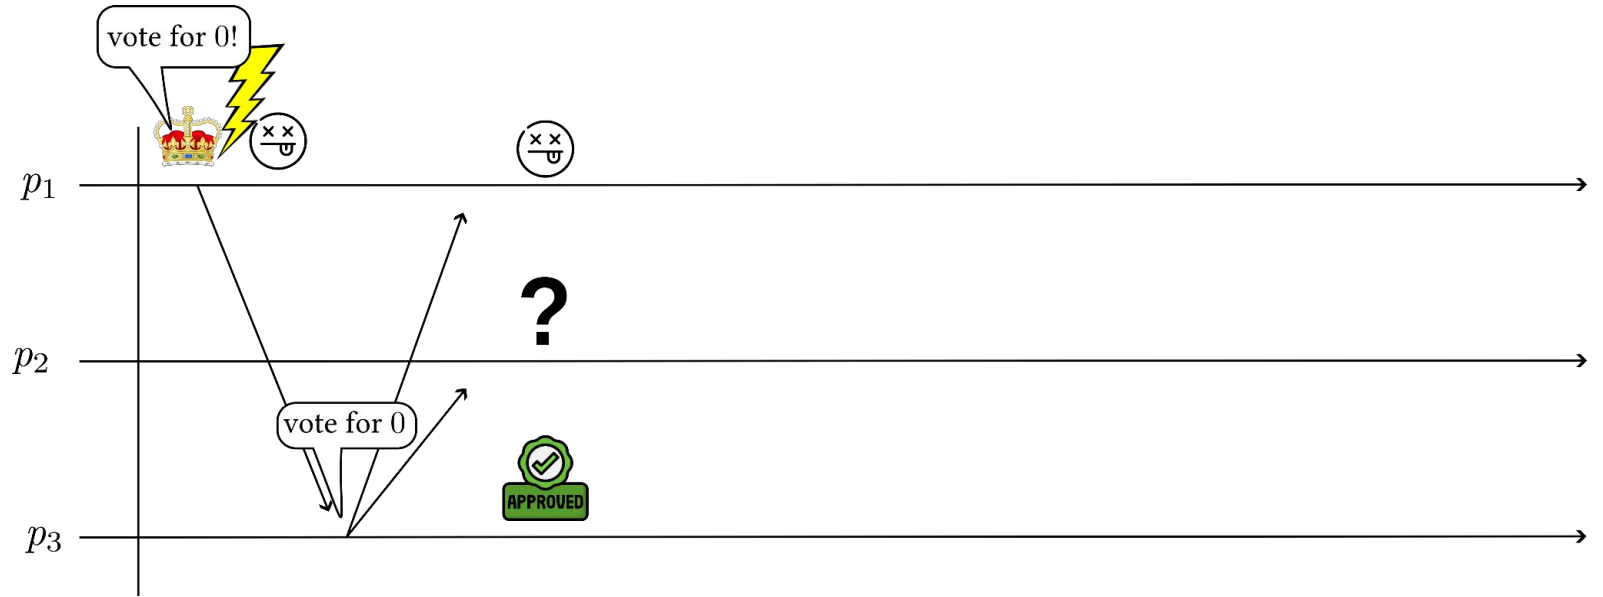


If the first leader fails to impose its decision, nodes time-out and try again with a new leader

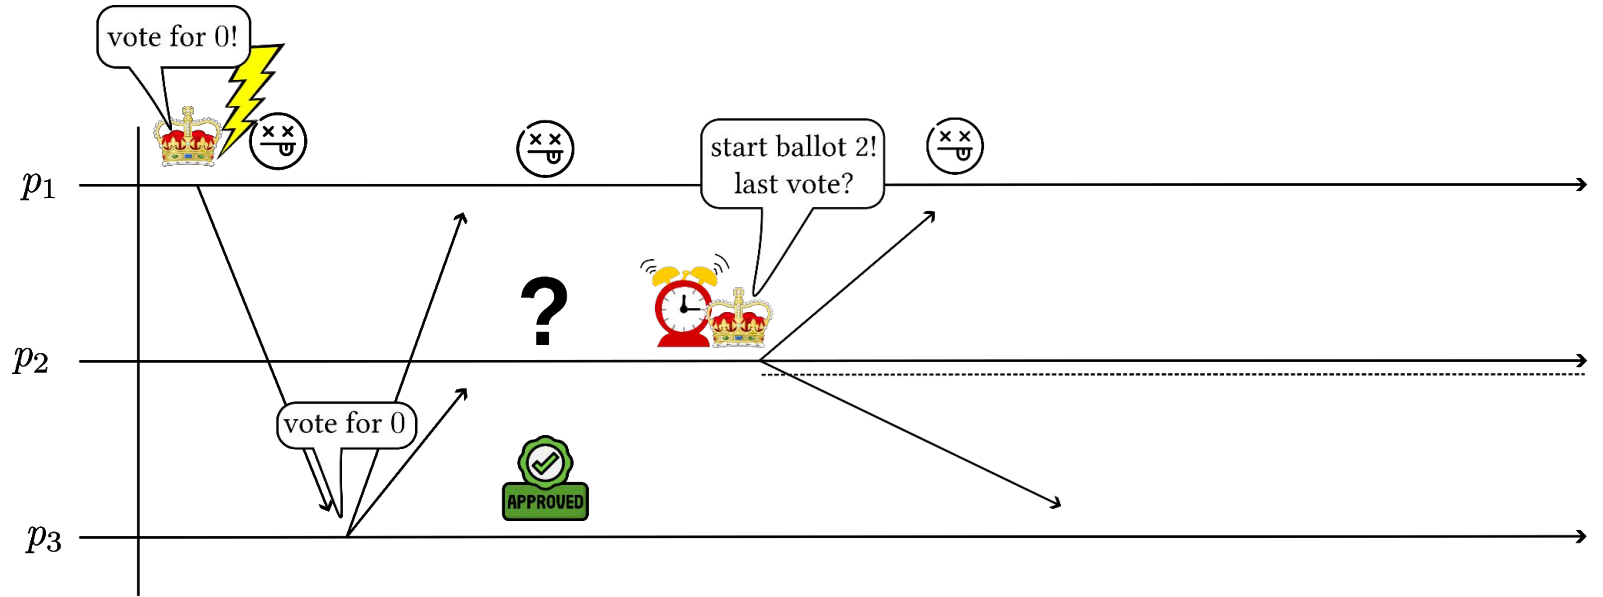
Each such attempt is called a ballot



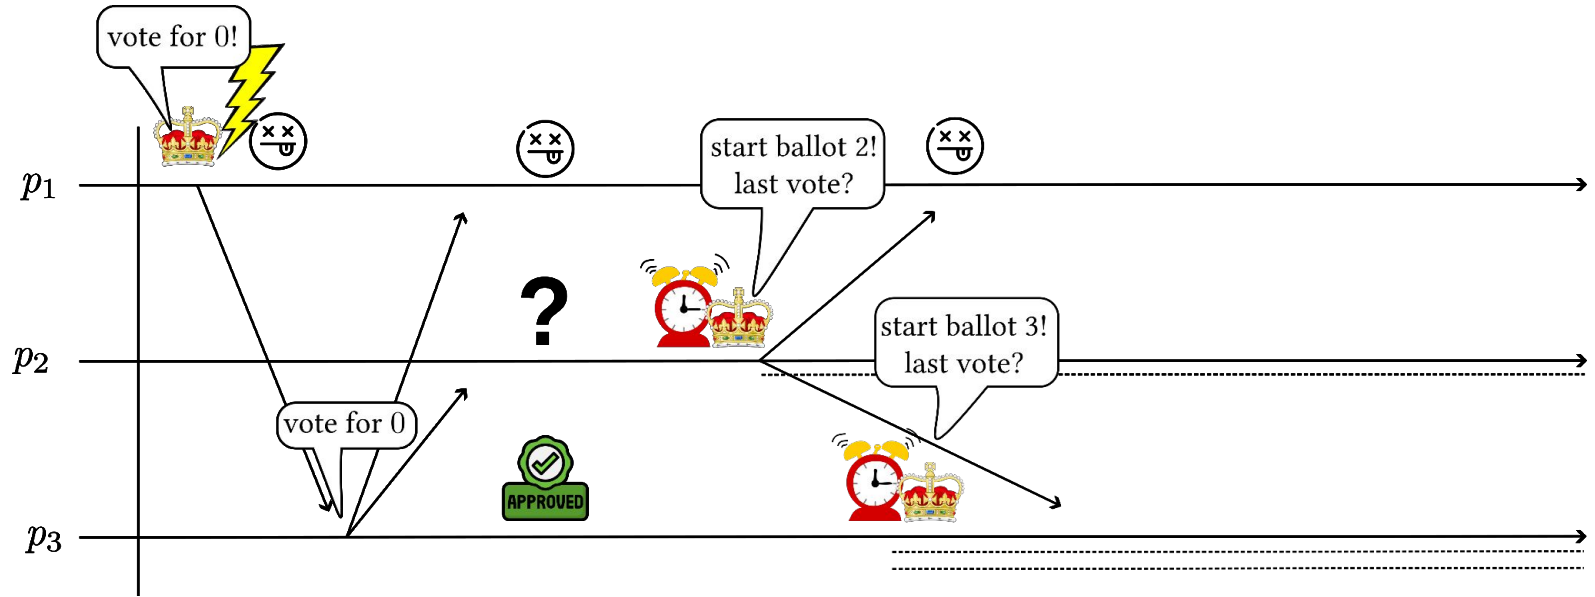
Depending on message delay, node might “step on each other’s toes” by starting new ballots too early



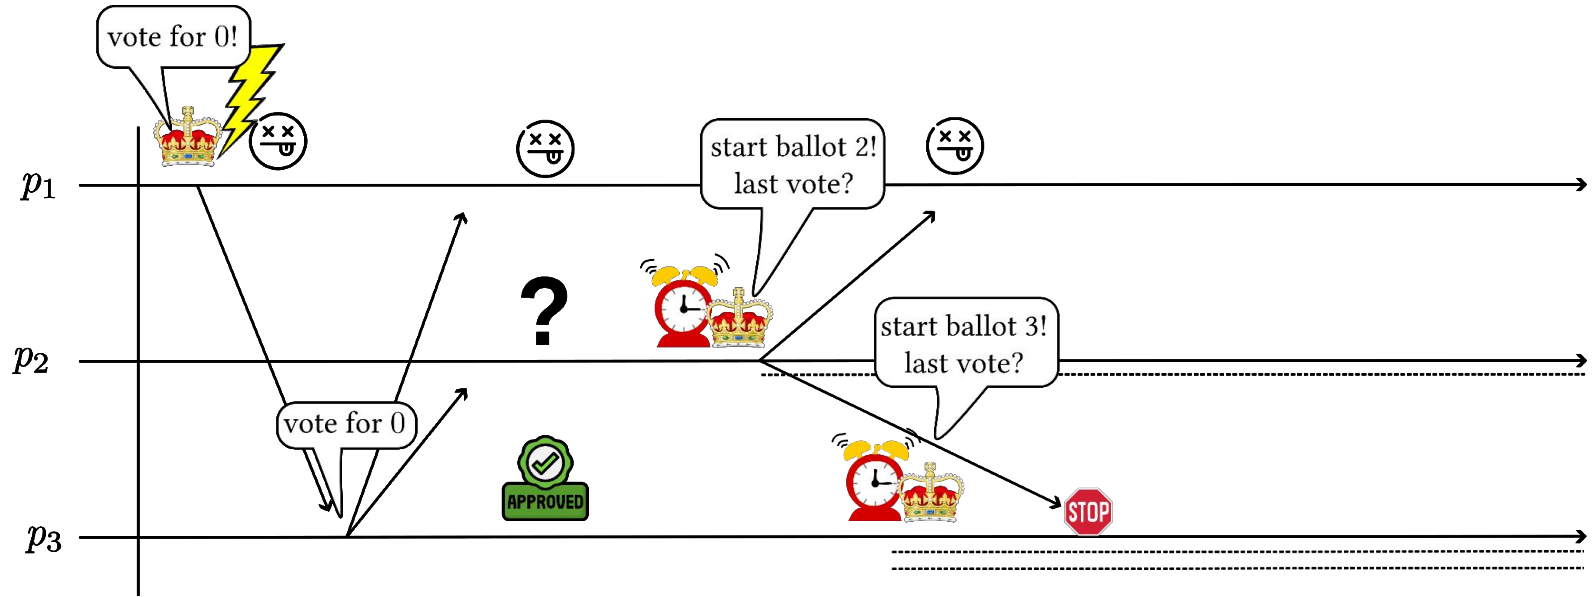
Depending on message delay, node might “step on each other’s toes” by starting new ballots too early



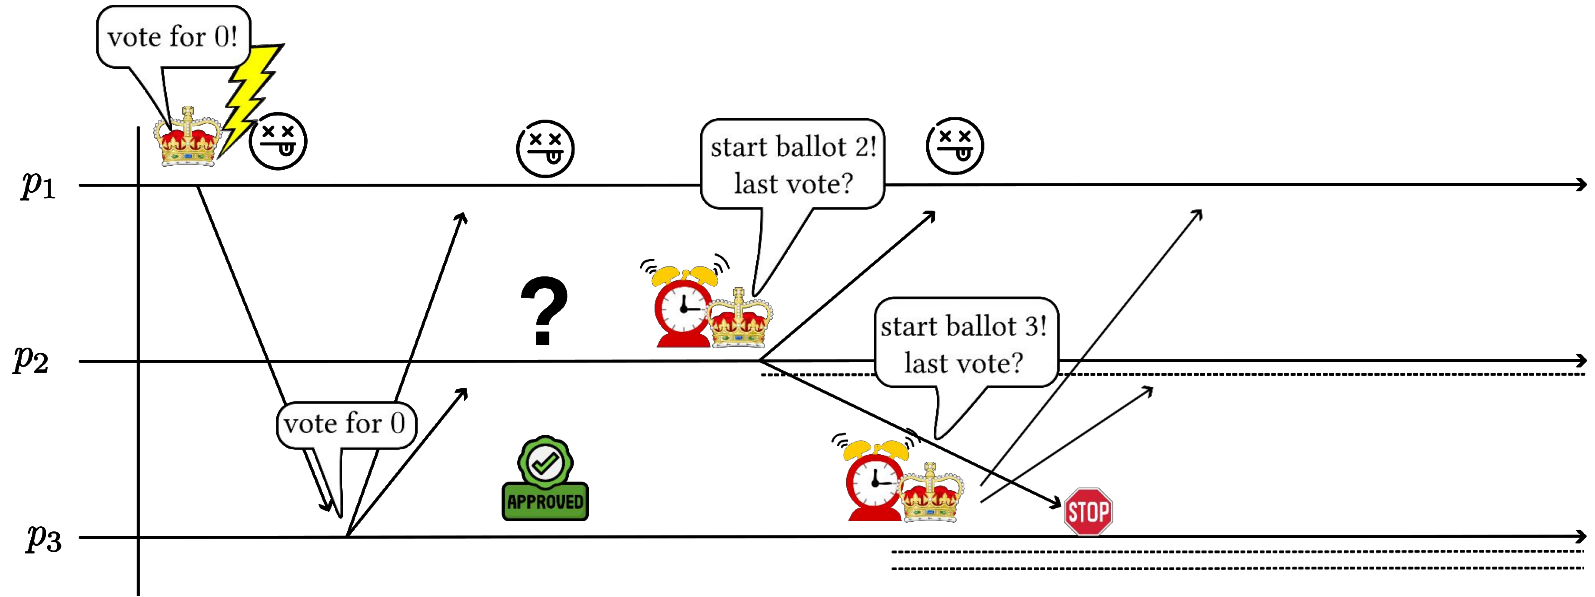
Depending on message delay, node might “step on each other’s toes” by starting new ballots too early



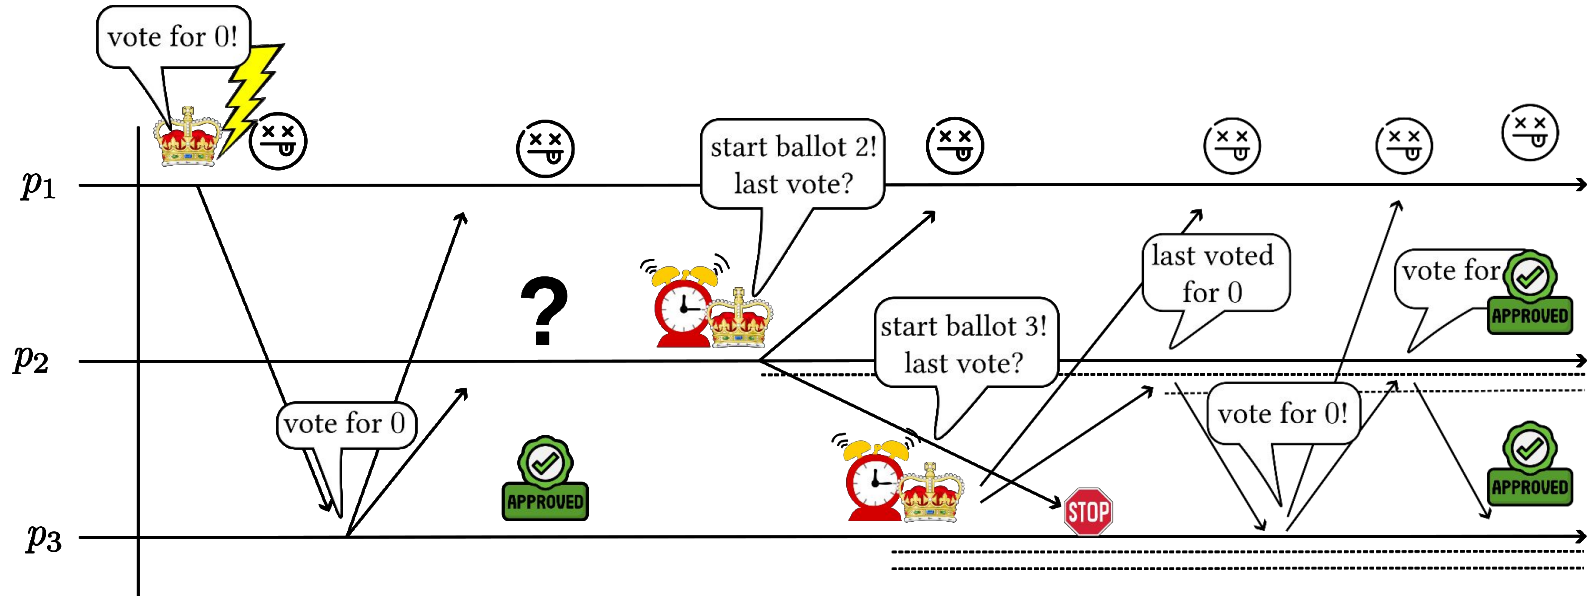
Depending on message delay, node might “step on each other’s toes” by starting new ballots too early



Depending on message delay, node might “step on each other’s toes” by starting new ballots too early



Depending on message delay, node might “step on each other’s toes” by starting new ballots too early



To make things easier, we abstract timers away

Instead of modeling message delay and timers, we prove the following liveness property:

If there is a *good ballot* B such that:

- No node starts a higher ballot
- The leader of B is well-behaved

Then, eventually, all well-behaved nodes decide

In TetraBFT, the general idea is the same but the protocol is more complicated

- There are 6 types of messages per ballot
- The leader may be malicious!
 - Nodes must verify that the leader's proposal is legitimate
 - Main difficulty: despite differing views, the leader must make sure that other nodes will be able to check that it's proposal is legitimate

The Paxos algorithm in TLA+ (excerpts)

VARIABLES *votes*, *currBal*, *proposals*, *crashed*

TypeOK \triangleq

\wedge *votes* \in [*Acceptor* \rightarrow SUBSET (*Ballot* \times *Value*)]

\wedge *currBal* \in [*Acceptor* \rightarrow *Ballot* \cup { - 1 }]

\wedge *proposals* \in SUBSET (*Ballot* \times *Value*)

\wedge *crashed* \in SUBSET *Acceptor*

The Paxos algorithm in TLA+ (excerpts)

VARIABLES *votes*, *currBal*, *proposals*, *crashed*

TypeOK \triangleq

\wedge *votes* \in [*Acceptor* \rightarrow SUBSET (*Ballot* \times *Value*)]

\wedge *currBal* \in [*Acceptor* \rightarrow *Ballot* \cup { - 1 }]

\wedge *proposals* \in SUBSET (*Ballot* \times *Value*)

\wedge *crashed* \in SUBSET *Acceptor*

chosen \triangleq { *v* \in *Value* : \exists *b* \in *Ballot* : \exists *Q* \in *Quorum* :

\forall *a* \in *Q* : \langle *b*, *v* \rangle \in *votes* }

The Paxos algorithm in TLA+ (excerpts)

VARIABLES *votes*, *currBal*, *proposals*, *crashed*

TypeOK \triangleq

\wedge *votes* \in [*Acceptor* \rightarrow SUBSET (*Ballot* \times *Value*)]
 \wedge *currBal* \in [*Acceptor* \rightarrow *Ballot* \cup { - 1 }]
 \wedge *proposals* \in SUBSET (*Ballot* \times *Value*)
 \wedge *crashed* \in SUBSET *Acceptor*

chosen \triangleq { *v* \in *Value* : \exists *b* \in *Ballot* : \exists *Q* \in *Quorum* :
 \forall *a* \in *Q* : \langle *b*, *v* \rangle \in *votes* }

Init \triangleq

\wedge *votes* = [*a* \in *Acceptor* \mapsto {}]
 \wedge *currBal* = [*a* \in *Acceptor* \mapsto - 1]
 \wedge *proposals* = {}
 \wedge *crashed* = {}

The Paxos algorithm in TLA+ (excerpts)

VARIABLES *votes*, *currBal*, *proposals*, *crashed*

TypeOK \triangleq

\wedge *votes* \in [*Acceptor* \rightarrow SUBSET (*Ballot* \times *Value*)]
 \wedge *currBal* \in [*Acceptor* \rightarrow *Ballot* \cup { - 1 }]
 \wedge *proposals* \in SUBSET (*Ballot* \times *Value*)
 \wedge *crashed* \in SUBSET *Acceptor*

chosen \triangleq { *v* \in *Value* : \exists *b* \in *Ballot* : \exists *Q* \in *Quorum* :
 \forall *a* \in *Q* : \langle *b*, *v* \rangle \in *votes* }

Init \triangleq

\wedge *votes* = [*a* \in *Acceptor* \mapsto {}]
 \wedge *currBal* = [*a* \in *Acceptor* \mapsto - 1]
 \wedge *proposals* = {}
 \wedge *crashed* = {}

Crash(*a*) \triangleq

\wedge *crashed'* = *crashed* \cup {*a*}
 \wedge \exists *Q* \in *Quorum* : \forall *a2* \in *Q* : *a2* \notin *crashed'*
 \wedge UNCHANGED \langle *votes*, *currBal*, *proposals* \rangle

The Paxos algorithm in TLA+ (excerpts)

VARIABLES *votes*, *currBal*, *proposals*, *crashed*

TypeOK \triangleq

\wedge *votes* \in [*Acceptor* \rightarrow SUBSET (*Ballot* \times *Value*)]
 \wedge *currBal* \in [*Acceptor* \rightarrow *Ballot* \cup { -1 }]
 \wedge *proposals* \in SUBSET (*Ballot* \times *Value*)
 \wedge *crashed* \in SUBSET *Acceptor*

chosen \triangleq { *v* \in *Value* : \exists *b* \in *Ballot* : \exists *Q* \in *Quorum* :
 \forall *a* \in *Q* : \langle *b*, *v* \rangle \in *votes* }

Init \triangleq

\wedge *votes* = [*a* \in *Acceptor* \mapsto {}]
 \wedge *currBal* = [*a* \in *Acceptor* \mapsto -1]
 \wedge *proposals* = {}
 \wedge *crashed* = {}

Crash(*a*) \triangleq

\wedge *crashed'* = *crashed* \cup {*a*}
 \wedge \exists *Q* \in *Quorum* : \forall *a2* \in *Q* : *a2* \notin *crashed'*
 \wedge UNCHANGED \langle *votes*, *currBal*, *proposals* \rangle

Propose(*b*, *v*) \triangleq

\wedge *Leader*(*b*) \notin *crashed*
 \wedge \forall *prop* \in *proposals* : *prop*[1] \neq *b*
 \wedge \exists *Q* \in *Quorum* : *ShowsSafeAt*(*Q*, *b*, *v*)
 \wedge *proposals'* = *proposals* \cup { \langle *b*, *v* \rangle }
 \wedge UNCHANGED \langle *votes*, *currBal*, *crashed* \rangle

IncreaseCurrBal(*a*, *b*) \triangleq

\wedge *a* \notin *crashed*
 \wedge *b* > *currBal*[*a*]
 \wedge *currBal'* = [*currBal* EXCEPT ![*a*] = *b*]
 \wedge UNCHANGED \langle *votes*, *proposals*, *crashed* \rangle

VoteFor(*a*, *b*, *v*) \triangleq

\wedge *a* \notin *crashed*
 \wedge *currBal*[*a*] \leq *b*
 \wedge \forall *vt* \in *votes*[*a*] : *vt*[1] \neq *b*
 \wedge \langle *b*, *v* \rangle \in *proposals*
 \wedge *votes'* = [*votes* EXCEPT ![*a*] = @ \cup { \langle *b*, *v* \rangle }]
 \wedge *currBal'* = [*currBal* EXCEPT ![*a*] = *b*]
 \wedge UNCHANGED \langle *crashed*, *proposals* \rangle

Methodology, step 1: prophesize the “good ballot”

VARIABLES $votes$, $currBal$, $proposals$, $crashed$, $goodBallot$

$TypeOK \triangleq$

$\wedge votes \in [Acceptor \rightarrow \text{SUBSET}(Ballot \times Value)]$
 $\wedge currBal \in [Acceptor \rightarrow Ballot \cup \{-1\}]$
 $\wedge proposals \in \text{SUBSET}(Ballot \times Value)$
 $\wedge crashed \in \text{SUBSET } Acceptor$
 $\wedge goodBallot \in Ballot$

$Init \triangleq$

$\wedge votes = [a \in Acceptor \mapsto \{\}]$
 $\wedge currBal = [a \in Acceptor \mapsto -1]$
 $\wedge proposals = \{\}$
 $\wedge crashed = \{\}$
 $\wedge goodBallot \in Ballot$

$Crash(a) \triangleq$

$\wedge a \neq Leader(goodBallot)$
 $\wedge crashed' = crashed \cup \{a\}$
 $\wedge \exists Q \in Quorum : \forall a2 \in Q : a2 \notin crashed'$
 $\wedge \text{UNCHANGED} \langle votes, currBal, proposals, goodBallot \rangle$

$Propose(b, v) \triangleq$

$\wedge Leader(b) \notin crashed$
 $\wedge \forall prop \in proposals : prop[1] \neq b$
 $\wedge \exists Q \in Quorum : ShowsSafeAt(Q, b, v)$
 $\wedge proposals' = proposals \cup \{\langle b, v \rangle\}$
 $\wedge \text{UNCHANGED} \langle votes, currBal, crashed, goodBallot \rangle$

$IncreaseCurrBal(a, b) \triangleq$

$\wedge a \notin crashed$
 $\wedge b \leq goodBallot$
 $\wedge b > currBal[a]$
 $\wedge currBal' = [currBal \text{ EXCEPT } ![a] = b]$
 $\wedge \text{UNCHANGED} \langle votes, proposals, crashed, goodBallot \rangle$

$VoteFor(a, b, v) \triangleq$

$\wedge a \notin crashed$
 $\wedge currBal[a] \leq b$
 $\wedge \forall vt \in votes[a] : vt[1] \neq b$
 $\wedge \langle b, v \rangle \in proposals$
 $\wedge votes' = [votes \text{ EXCEPT } ![a] = @ \cup \{\langle b, v \rangle\}]$
 $\wedge currBal' = [currBal \text{ EXCEPT } ![a] = b]$
 $\wedge \text{UNCHANGED} \langle crashed, proposals, goodBallot \rangle$

Without Byzantine behavior, we can check the following temporal property with TLC... but, with TetraBFT, we run into scaling issues

$$\begin{aligned} FairSpec &\triangleq \\ &\wedge \textit{Init} \\ &\wedge \Box[Next]_{vars} \\ &\wedge \forall a \in \textit{Acceptor}, b \in \textit{Ballot}, v \in \textit{Value} : \\ &\quad \wedge \text{WF}_{vars}(\textit{IncreaseCurrBal}(a, b)) \\ &\quad \wedge \text{WF}_{vars}(\textit{VoteFor}(a, b, v)) \\ &\quad \wedge \text{WF}_{vars}(\textit{Propose}(b, v)) \\ \\ Liveness &\triangleq \Diamond(\textit{chosen} \neq \{\}) \end{aligned}$$

THEOREM $\textit{LiveSpec} \Rightarrow \textit{Liveness}$

Instead, we will verify that all fairly-scheduled actions are eventually disabled, and that once this is the case we must have a decision

$$\begin{aligned} \textit{Liveness} &\triangleq \\ &\forall a \in \textit{Acceptor}, b \in \textit{Ballot}, v \in \textit{Value} : \\ &\quad \wedge \neg\textit{ENABLED } \textit{IncreaseCurrBal}(a, b) \\ &\quad \wedge \neg\textit{ENABLED } \textit{VoteFor}(a, b, v) \\ &\quad \wedge \neg\textit{ENABLED } \textit{Propose}(b, v) \\ &\Rightarrow \textit{chosen} \neq \{\} \end{aligned}$$

Methodology, step 2: verify that all fairly-scheduled actions are self-disabling

Since we work in a finite domain, the property implies that, if those actions are fairly scheduled, then eventually all will be disabled

This implies that, if Liveness holds and the actions are fairly scheduled, then eventually one value is chosen

Methodology, step 2a: add ghost variable to remember which actions took place already

VARIABLES $votes, currBal, proposals, crashed, goodBallot,$
 $proposeTaken, voteForTaken, increaseCurrBalTaken$

$TypeOK \triangleq$

$\wedge votes \in [Acceptor \rightarrow SUBSET (Ballot \times Value)]$
 $\wedge currBal \in [Acceptor \rightarrow Ballot \cup \{-1\}]$
 $\wedge crashed \in SUBSET Acceptor$
 $\wedge goodBallot \in Ballot$
 $\wedge proposals \in SUBSET (Ballot \times Value)$
 $\wedge proposeTaken \in [Ballot \times Value \rightarrow BOOLEAN]$
 $\wedge voteForTaken \in [Acceptor \times Ballot \times Value \rightarrow BOOLEAN]$
 $\wedge increaseCurrBalTaken \in [Acceptor \times Ballot \rightarrow BOOLEAN]$

$Init \triangleq$

$\wedge votes = [a \in Acceptor \mapsto \{\}]$
 $\wedge currBal = [a \in Acceptor \mapsto -1]$
 $\wedge crashed = \{\}$
 $\wedge proposals = \{\}$
 $\wedge goodBallot \in Ballot$
 $\wedge proposeTaken = [x \in Ballot \times Value \mapsto FALSE]$
 $\wedge voteForTaken = [x \in Acceptor \times Ballot \times Value \mapsto FALSE]$
 $\wedge increaseCurrBalTaken = [x \in Acceptor \times Ballot \mapsto FALSE]$

$Propose(b, v) \triangleq$

$\wedge Leader(b) \notin crashed$
 $\wedge \forall prop \in proposals : prop[1] \neq b$
 $\wedge \exists Q \in Quorum : ShowsSafeAt(Q, b, v)$
 $\wedge proposals' = proposals \cup \{ \langle b, v \rangle \}$
 $\wedge proposeTaken' = [proposeTaken \text{ EXCEPT } ![\langle b, v \rangle] = TRUE]$
 $\wedge UNCHANGED \langle votes, currBal, crashed, goodBallot,$
 $voteForTaken, increaseCurrBalTaken \rangle$

$IncreaseCurrBal(a, b) \triangleq$

$\wedge a \notin crashed$
 $\wedge goodBallot > -1 \Rightarrow b \leq goodBallot$
 $\wedge b > currBal[a]$
 $\wedge currBal' = [currBal \text{ EXCEPT } ![a] = b]$
 $\wedge increaseCurrBalTaken' =$
 $[increaseCurrBalTaken \text{ EXCEPT } ![\langle a, b \rangle] = TRUE]$
 $\wedge UNCHANGED \langle votes, crashed, goodBallot, proposals,$
 $proposeTaken, voteForTaken \rangle$

$VoteFor(a, b, v) \triangleq$

...

Methodology, step 2b: assert and verify with an inductive invariant that all fair actions are self-disabling

$SelfDisabling \triangleq$

$\forall a \in Acceptor, b \in Ballot, v \in Value :$

$\wedge voteForTaken[\langle a, b, v \rangle] \Rightarrow$

$\neg \text{ENABLED } VoteFor(a, b, v)$

$\wedge increaseCurrBalTaken[\langle a, b \rangle] \Rightarrow$

$\neg \text{ENABLED } IncreaseCurrBal(a, b)$

$\wedge proposeTaken[\langle b, v \rangle] \Rightarrow$

$\neg \text{ENABLED } Propose(b, v)$

Methodology, step 3: verify with an inductive invariant that, once all actions are disabled, we have a decision

$$\begin{aligned} \text{GoodBallotIsMaxBallot} &\triangleq \\ \forall a \in \text{Acceptor}, b \in \text{Ballot}, v \in \text{Value} : \\ &\langle b, v \rangle \in \text{proposals} \vee b = \text{currBal}[a] \\ &\Rightarrow b \leq \text{goodBallot} \end{aligned}$$

$$\text{QuorumNotCrashed} \triangleq \exists Q \in \text{Quorum} : Q \cap \text{crashed} = \{\}$$

$$\begin{aligned} \text{VoteForProposal} &\triangleq \forall a \in \text{Acceptor}, b \in \text{Ballot}, v \in \text{Value} : \\ &\text{VotedFor}(a, b, v) \Rightarrow \langle b, v \rangle \in \text{proposals} \end{aligned}$$

$$\begin{aligned} \text{ProposalsUnique} &\triangleq \\ \forall a1, a2 \in \text{Acceptor}, b \in \text{Ballot}, v1, v2 \in \text{Value} : \\ &\langle b, v1 \rangle \in \text{proposals} \wedge \langle b, v2 \rangle \in \text{proposals} \Rightarrow v1 = v2 \end{aligned}$$

$$\text{GoodBallotLeaderNotCrashed} \triangleq \neg \text{Leader}(\text{goodBallot}) \in \text{crashed}$$

$$\begin{aligned} \text{LivenessInvariant} &\triangleq \\ &\wedge \text{TypeOK} \\ &\wedge \text{VoteForProposal} \\ &\wedge \text{ProposalsUnique} \\ &\wedge \text{GoodBallotIsMaxBallot} \\ &\wedge \text{QuorumNotCrashed} \\ &\wedge \text{GoodBallotLeaderNotCrashed} \end{aligned}$$

THEOREM $\text{Spec} \Rightarrow \square \text{LivenessInvariant}$

THEOREM $\text{LivenessInvariant} \Rightarrow \text{Liveness}$

Verifying liveness of eventually-synchronous consensus protocols, in a nutshell

1. We modify the specification to non-deterministically pick a good ballot B and enforce that no higher ballots are started and that the leader of B is well-behaved
2. We verify that each fair action is self-disabling. Under fair scheduling, this means that all (the finitely many) actions will eventually be disabled.
3. We verify that, once all actions are disabled, we have a decision

In steps 2 and 3 we verify safety properties
We can use an inductive invariant verified with Apache

We are able to check the liveness of TetraBFT for interesting problem sizes

Problem size

- 3 rounds (with 5 sub-phases per round)
- 3 possible decision values
- 3 nodes n_1, n_2, n_3
- Either n_2 or n_3 are malicious (but not both)
- Quorums are $\{n_1, n_2\}$ and $\{n_1, n_3\}$

Resources used by Apache

- ~ few hours
- ~ 40 GB of RAM
- 16 cores on a recent (2022) desktop machine

Specifications are available at
<https://github.com/nano-o/tetrabft-tla>



Future work

- Prove what we assumed about timer management, i.e. that eventually, there is a ballot led by a well-behaved node such that no node ever starts a higher ballot (Paxos)
 - Requires a real-time model of the behavior after GST
- Same thing for TetraBFT, where this is more complex: round progression must be governed by a so-called BFT synchronizer algorithm