

UC San Diego

Technical Reports

Title

Weak Leader Election in the receive-omission failure model

Permalink

<https://escholarship.org/uc/item/1r75q8vw>

Authors

Junqueira, Flavio
Marzullo, Keith

Publication Date

2005-06-01

Peer reviewed

Weak Leader Election in the receive-omission failure model

Flavio P. Junqueira
flavio@cs.ucsd.edu

Keith Marzullo
marzullo@cs.ucsd.edu

University of California, San Diego
Department of Computer Science and Engineering
9500 Gilman Drive
La Jolla, CA – USA

1 Introduction

Leader Election is an important primitive in fault-tolerant distributed computing because it enables the solution of problems broadly applicable in real systems such as Consensus, as illustrated by the Paxos algorithm [1], and Primary-Backup, as in [2].

The particular version of the Leader Election problem we develop upon first appeared in the context of Primary-Backup algorithms. In the Primary-Backup approach for fault-tolerant services, clients issue requests that the primary is responsible for handling and replying to. When the primary fails, one of the backup replicas takes over as the new primary. Thus, a Primary-Backup algorithm embeds a Leader Election algorithm that can infinitely often select a primary.

In [2], Budhiraja and Marzullo show a lower bound of $n > \lfloor 3t/2 \rfloor$ for such algorithms when processes can fail to receive messages, where n is the number of processes and t is the maximum number of process failures in an execution. The basic idea of the lower bound proof is that multiple primaries can be elected if fewer than $\lfloor 3t/2 \rfloor + 1$ processes compose the system. In a later section, we repeat this result for exposition purposes.

Still on the early work by Budhiraja and Marzullo on Primary-Backup algorithms, the degree of replication necessary for the algorithm they designed is higher because they require that faulty processes cannot be elected [2]. According to their statement of the problem, if a process does not crash but it commits receive-omission failures, then it cannot be elected. This is due to the assumption that responses to client requests are bounded in time. Failure detection for receive-omission failures, however, requires at least twofold replication.

When implementing a system based on the Primary-Backup approach, servers are often connected by a local area network to bound response time to client requests, which implies bounded fail-over time. For such settings, partitions are unlikely to occur if processors operate at a reasonable speed. Messages, however, can be lost due to, for example, buffer overflows at the receiver. One can imagine using retransmissions to cope with such failures. A retransmission mechanism, however, only guarantees eventual delivery; bounded response is not possible with eventual delivery of messages. Because of the requirements on bounded response and failure-over time, Primary-Backup algorithms are usually synchronous.

In this paper, we describe a synchronous algorithm for Leader Election under receive-omission process failures and prove its correctness. The novelty in this algorithm is fourfold: 1) it proves tight a lower bound that has been known for over 10 years; 2) by permitting faulty (but not crashed) processes to be elected, it requires fewer replicas; 3) it is based on cores and survivor sets which are abstractions that enables one to more expressively represent failure scenarios by considering failures that are not independent or not identically distributed; 4) although it allows for faulty processes to be elected, correct processes are able to detect it, enabling the use of alarms to indicate failures in the system. Relating to our discussion on Primary-Backup protocols, by assuming that faulty processes can be elected, we cannot bound response time for a Primary-Backup algorithm. We can guarantee, however, that there is at most one primary at any time, and that response is bounded whenever a correct process emerges as the primary. We further discuss this and other issues with Primary-Backup protocols later in the paper.

The remainder of this paper is as follows. We detail the system model in Section 2. We then introduce the problem by stating the properties an algorithm must fulfill (Section 3). Still in Section 3, we repeat the lower bound proof for process replication, and generalize this bound to our model of dependent failures. Section 4 describes our FFS-WLE algorithm for Leader Election. As we shall see, the algorithm depends on a primitive that we call *RO Consensus*. The properties for RO Consensus resembles the ones for the traditional Uniform Consensus primitive. The differences, however, are significant enough for naming the problem differently. In Section 4, we also provide an algorithm for RO Consensus. Sections 5 and 6 provide proofs of correctness for FFS-ROC and FFS-WLE, respectively. In Section 7, we strengthen the definition of Leader Election to disable executions in which different leaders are elected infinitely often, and provide a simple modification of the algorithm that enables it. Before concluding, we provide a discussion on the implications of the properties of our algorithm in a Primary-Backup protocol in Section 8. We finally conclude in Section 9.

2 System model

A system is a collection of processes $\Pi = \{p_1, p_2, \dots, p_n\}$ that communicate through messages.¹ For every pair of processes $p_i, p_j \in \Pi$, there is a channel that p_i uses to send messages to p_j .

In such a system, an algorithm alg is a collection of state machines, one for each process. alg then proceeds in steps of processes. In a step, a process p_i executes atomically the following:

- \bigwedge \vee receives a message from a process p_j
- \vee sends a message to a process p_j
- \vee executes a local operation
- \bigwedge undergoes a state transition

We define an execution ϕ of alg as a tuple $\langle F, I, S, T, \mathcal{T} \rangle$, where F is the set of processes that are faulty in ϕ ; I is the set of initial states, one for each process; S is a set of steps; T is a set of real values; $\mathcal{T} : S \rightarrow T$ is a mapping from steps to real values. The real values in T correspond to the global time in which steps execute. $\mathcal{T}(s)$ therefore is the global time in which step $s \in S$ executes. We use global time in

¹We use p_i to denote a process and i to denote the identifier of this process.

proofs, and we do not assume such a virtual clock that produces global time is available to processes. We also use $Correct(\phi)$ for the set of processes that are correct in ϕ : $Correct(\phi) = \Pi \setminus F$. Finally, Φ is the set of executions of alg .

We assume that processes can fail by crashing or by omitting to receive messages. If a process p_i crashes in an execution ϕ , then there is step s of p_i such that p_i executes no further steps after $\mathcal{T}(s)$. We call this step a *crash step*. A faulty process, however, does not necessarily crash: it can selectively fail to receive messages. To characterize failure scenarios, a threshold on the number of process failures is often used. We refer to this characterization using a threshold as the *threshold model*. We instead use our model of dependent process failures based on the abstractions of cores and survivor sets. We define cores and survivor sets as follows:

Definition 2.1 A subset $C \subseteq \Pi$ is a core if and only if: 1) $\forall \phi \in \Phi, Correct(\phi) \cap C \neq \emptyset$; 2) $\forall p_i \in C : \exists \phi \in \Phi : C \setminus \{p_i\} \cap Correct(\phi) = \emptyset$.

Definition 2.2 A subset $S \subseteq \Pi$ is a survivor set if and only if: 1) $\exists \phi \in \Phi, Correct(\phi) = S$; 2) $\forall \phi \in \Phi : \forall p_i \in S : Correct(\phi) \not\subseteq S \setminus \{p_i\}$.

We use the term *system profile* to denote a description of the tolerated failure scenarios. In the threshold model, a system profile is a pair $\langle \Pi, t \rangle$, which means that any subset of t processes in Π can be faulty. In our dependent failure model, the system profile is a triple $\langle \Pi, C_\Pi, S_\Pi \rangle$, where C_Π is the set of cores and S_Π is the set of survivor sets. We assume that each process is a member of at least one survivor set (otherwise, that process can be faulty in each execution, and ignored by the other processes), and that no process is a member of every survivor set (otherwise, that process is never faulty).

We assume that systems are synchronous: the steps of every execution of some algorithm \mathcal{A} can be split into rounds. That is, there is a mapping $Round : S \rightarrow \mathcal{R}$ from steps of processes to round numbers, where $\mathcal{R} = \mathbb{Z}^*$ and round numbers monotonically increase with time. We then have the following properties for rounds:

P-Liveness : If a process executes all the steps of a round r , then every process that does not crash by r executes at least one step of r .

C-Liveness : If a process p_i sends a message m to a correct process p_j at round r and p_i does not crash by round r , then p_j receives m at round r .

Integrity : If p_i receives a message m from p_j , then p_j sent m to p_i .

No duplicates : No message m is received more than once.

3 Problem specification

For the following description of the problem, we assume that each process p_i in Π has a boolean variable $p_i.elected$ that is set to true if the process elects itself, and to false otherwise. We then define the *Weak Leader Election* problem with the following three properties:

Safety $\square \{ \{ p_i \in \Pi : p_i.elected \} < 2. \}$

LE-Liveness $\square \diamond \{ \{ p_i \in \Pi : p_i.elected \} > 0 \}$.

FF-Stability In a failure-free execution, only one process ever has *elected* set to true.

In words, infinitely often some process elects itself, and no more than one process is elected at any time. The third property eliminates the possibility of an algorithm that, for example, elects processes in a round-robin fashion (which can be implemented with *no* communications given that the system is synchronous). It does not rule out, however, executions in which two or more processes are elected infinitely often when there is at least one process failure in the execution. For this reason, we propose another property called *E-Stability* stated as follows:

E-Stability $\exists p_i \in \Pi : \diamond \square (\forall p_j \in \Pi : p_j.elected \Rightarrow (j = i))$

An algorithm satisfying this property eventually elects the same process forever in every execution. Note that with E-Stability only, failure-free executions are allowed to have multiple leaders elected (at different times, to not violate safety), and hence does not render FF-Stability unnecessary.

In the following sections, we first derive an algorithm that satisfies the first three properties. Later we modify this algorithm to also satisfy E-Stability. First, we show a lower bound for this problem.

3.1 Lower bound on process replication

In [2], the following lower bound was shown. The proof was given in the context of showing a lower bound on replication for Primary-Backup protocols.

Lemma 3.1 *Weak Leader Election for receive-omission failures requires $n > \lfloor 3t/2 \rfloor$.*

Proof:

Assume that Weak Leader Election for receive-omission failures can be solved with $n = \lfloor 3t/2 \rfloor$. Partition the processes into three blocks A , B and C , where $|A| = |B| = \lfloor t/2 \rfloor$ and $|C| = \lceil t/2 \rceil$. Consider an execution ϕ_A in which the processes in B and C initially crash. From LE-Liveness and E-Stability, eventually a process in A will be elected infinitely often. Similarly, let ϕ_B be an execution in which the processes in A and C crash. From LE-Liveness and E-Stability, eventually a process in B will be elected infinitely often.

Finally, consider an execution ϕ in which the processes in A fail to receive all messages except those sent by processes in A , and the processes in B fail to receive all messages except those sent by processes in B . This execution is indistinguishable from ϕ_A to the processes in A and is indistinguishable from ϕ_B to the processes in B . Hence, there will eventually be two processes, one in A and one in B , elected infinitely often, violating either Safety or E-Stability.

□

3.2 Replication predicate

To develop the protocol, we first generalize the replication predicate for this problem for the core/survivor set model, where a replication predicate is a predicate that establishes whether there is sufficient replication to enable the solution of a particular problem. From the lower bound proof of Lemma 3.1, we consider any partition of the processes into three blocks. Then, one constructs three executions, where in each execution all of the processes in two of the three subsets are faulty. If we can construct such a partition, then we cannot solve the problem. Before we present the partition property, we introduce a few definitions:

- $\mathcal{P}_k(\Pi)$ is the set of partitions of Π into k blocks;
- $\mathcal{G}_x(A)$ is the set of all the subsets of A of size x ; if $|A| < x$, then $\mathcal{G}_x(A) = \emptyset$

The conclusion of the lower bound proof can be stated, for $k = 3$, as follows:

Property 3.2 (k,k - 1)-Partition, $k > 1, |\Pi| > 2$: $\exists k' \in \{2, \dots, \min(k, |\Pi|)\} : \forall \mathcal{A} \in \mathcal{P}_{k'}(\Pi) : \exists \mathcal{A}' \in \mathcal{G}_{k'-1}(\mathcal{A}) : \exists C \in \mathcal{C}_{\Pi} : C \subseteq \mathcal{A}'$ □

The equivalent intersection property is then:

Property 3.3 (k,k-1)-Intersection, $k > 1$, $|\Pi| > 2$, $|\mathcal{S}_\Pi| > 2$: $\exists k' \in \{2, \dots, \min(k, |\Pi|)\} : \forall \mathcal{T} \in \mathcal{G}_{k'}(\mathcal{S}_\Pi) : \exists T \in \mathcal{G}_2(\mathcal{T}) : (\cap_{S \in T} S) \neq \emptyset \square$

Stated more simply, (k,k-1)-Intersection says that for any set of k' survivor sets, $k' \in \{2, \dots, \min(k, |\Pi|)\}$, at least two of them have a non-empty intersection. (k,k-1)-Intersection and (k,k-1)-Partition generalize replication predicates in the threshold model of the form $n > \lfloor kt/(k-1) \rfloor$.

Consider now an example of a system that satisfies (3,2)-Intersection. It is based on a simple two-cluster system. A process can fail by crashing, and there is a threshold t on the number of crash failures that can occur in a cluster. A cluster can also suffer a catastrophic failure, which causes all of the processes in that cluster to crash. Such a catastrophic failure can result from the failure of a cluster resource such as a disk array or a power supply, or from an administrative fault. We assume that catastrophic failures are rare enough that the probability of both clusters suffering catastrophic failures is negligible. However, processes can crash in one cluster at the same time that the other cluster suffers a catastrophic failure.

Assuming that each cluster has three processes and $t = 1$, we have the following system profile, where processes with identifier a_i belong to one cluster and processes with identifier b_i belong to the other cluster:

Example 3.4 :

$$\begin{aligned} \Pi &= \{p_{a_1}, p_{a_2}, p_{a_3}, p_{b_1}, p_{b_2}, p_{b_3}\} \\ C_\Pi &= \{\{p_{i_1}, p_{i_2}, p_{i_3}, p_{i_4}\} : (i_1, i_2 \in \{a_1, a_2, a_3\}) \\ &\quad \wedge (i_3, i_4 \in \{b_1, b_2, b_3\}) \wedge i_1 \neq i_2 \wedge i_3 \neq i_4\} \\ \mathcal{S}_\Pi &= \{\{p_{i_1}, p_{i_2}\} : ((i_1, i_2 \in \{a_1, a_2, a_3\}) \\ &\quad \vee (i_1, i_2 \in \{b_1, b_2, b_3\})) \wedge i_1 \neq i_2\} \end{aligned}$$

To see why this profile satisfies (3,2)-Intersection, we just have to observe that out of any three survivor sets, at least two intersect.

We now show the equivalence of (k,k-1)-Partition and (k,k-1)-Intersection with the following theorem.

Theorem 3.5 $(k,k-1)$ -Partition \equiv $(k,k-1)$ -Intersection

Proof:

\Rightarrow : Proof by contrapositive. Consider a system profile $\langle \Pi, C_\Pi, \mathcal{S}_\Pi \rangle$ such that, for every k' , $k' \in \{2, \dots, \min(k, |\Pi|)\}$, there is a subset $\mathcal{S} =$

$\{S_1, S_2, \dots, S_{k'}\} \subseteq \mathcal{S}_\Pi$ such that no pair of survivor sets S_i, S_j intersects. That is, $\bigcup_{P \in \mathcal{G}_2(\mathcal{S})} \cap P = \emptyset$. We then build a partition $\mathcal{A} = \{A_1, A_2, \dots, A_{k'}\}$ as follows:

$$\begin{aligned} A_1 &= \Pi \setminus (S_2 \cup S_3 \cup \dots \cup S_{k'}) \\ A_2 &= \Pi \setminus (S_1 \cup S_3 \cup \dots \cup S_{k'} \cup A_1) \\ &\vdots \\ A_i &= \Pi \setminus (S_1 \cup S_2 \cup \dots \cup S_{i-1} \cup S_{i+1} \cup \dots \\ &\quad \dots \cup S_{k'} \cup A_1 \cup A_2 \dots \cup A_{i-1}) \\ &\vdots \\ A_{k'} &= \Pi \setminus (S_1 \cup S_2 \cup \dots \cup S_{k'-1} \cup A_1 \dots \cup A_{k'-1}) \end{aligned}$$

It is clear that A_i, A_j are disjoint, $i \neq j$. We now have to show that: 1) $\bigcup \mathcal{A} = \Pi$; 2) For every subset $\mathcal{A}' = \{A_{i_1}, A_{i_2}, \dots, A_{i_{k'-1}}\} \subset \mathcal{A}$, $\bigcup \mathcal{A}'$ does not contain a core. To show 1), let $\psi_i = \bigcup_{(S_j \in \mathcal{S} \setminus S_i)} S_j$, $i \in \{1, \dots, k'\}$. We then have the following derivation:

$$\begin{aligned} \bigcup \mathcal{A} &= (\Pi \setminus \psi_1) \cup (\Pi \setminus \psi_2 \cup A_1) \cup \dots \\ &\quad \dots \cup (\Pi \setminus (\psi_{k'} \cup A_1 \cup A_2 \dots \\ &\quad \dots \cup A_{k'-1})) \end{aligned} \quad (1)$$

$$= \Pi \setminus ((\psi_1 \cap (\psi_2 \cup A_1)) \cap \dots \\ \dots \cap (\psi_{k'} \cup A_1 \cup A_2 \dots \cup A_{k'-1})) \quad (2)$$

$$= \Pi \setminus (\psi_1 \cap \psi_2 \cap \dots \\ \dots \cap (\psi_{k'} \cup A_1 \cup A_2 \dots \cup A_{k'-1})) \quad (3)$$

$$\vdots \\ = \Pi \setminus (\cap_i \psi_i) \quad (4)$$

$$= \Pi \quad (5)$$

- Line 1 to Line 2 follows from the observation that for any subsets A, B of Π , we have that $(\Pi \setminus A) \cup (\Pi \setminus B) = \Pi \setminus (A \cap B)$;
- Line 2 to Line 3: the intersection between ψ_1 and A_1 has to be empty, since ψ_1 contains exactly the elements we removed from Π to form A_1 .
- Line 3 to Line 4: by repeating inductively the process used to derive Line 3, we are able to remove every term A_i present in the equation.
- Line 4 to Line 5: Transforming from a conjunctive form to a disjunctive form, we have that $\bigcap_{P \in \mathcal{G}_{k'-1}(\mathcal{S})} \cup P = \bigcup_{P \in \mathcal{G}_2(\mathcal{S})} \cap P$. To see why this is true, note that for every pair $S_i, S_j \in \mathcal{S}$, $i \neq j$, and $P \in \mathcal{G}_{k'-1}(\mathcal{S})$, we have that $(S_i \in P) \vee (S_j \in P)$. Finally, we have that $\bigcup_{P \in \mathcal{G}_2(\mathcal{S})} (\cap_{S_i \in P} S_i) = \emptyset$ by assumption.

By the construction of the partition and from the assumption that for every $S_i, S_j \in \mathcal{S}$, $S_i \cap S_j = \emptyset$, we have that for every $i \in \{1, \dots, k'\}$, there is $S_i \in \mathcal{S}$ such that $S_i \subseteq A_i$. From this, we conclude that for any $\mathcal{A}' = \{A_{i_1}, A_{i_2}, \dots, A_{i_{k'-1}}\} \subset \mathcal{A}$, $\bigcup \mathcal{A}'$ does not contain elements from some survivor set, and consequently it does not contain a core.

\Leftarrow : Proof also by contrapositive. Suppose a system profile $\langle \Pi, C_\Pi, S_\Pi \rangle$ such that for every k' , $k' \in \{2, \dots, \min(k, |\Pi|)\}$, there is a partition $\mathcal{A} = \{A_1, A_2, \dots, A_{k'}\}$ of Π in which no union of $k' - 1$ blocks of \mathcal{A} contains a core. If a subset of processes does not contain a core, then it contains no elements from some survivor set. The complement of such a set of processes consequently contains a survivor set. Because no union of $k' - 1$ blocks in \mathcal{A} contains a core, for every A_i there is an $S_i \in \mathcal{S}_\Pi$ such that $S_i \subseteq A_i$. Thus, for all $A_i, A_j \in \mathcal{A}$, $i \neq j$, we have by construction that $A_i \cap A_j = \emptyset$, and hence $S_i \cap S_j = \emptyset$. We conclude that no pair $S_i, S_j \in \{S_1, S_2, \dots, S_{k'}\}$ is such that $S_i \cap S_j \neq \emptyset$. \square

4 The algorithm

In this Section, we describe an algorithm FFS-WLE that satisfies Safety, LE-Liveness, and FF-Stability (Figure 2). It assumes a system profile $\langle \Pi, C_\Pi, S_\Pi \rangle$ that satisfies (3,2)-Intersection and uses as a building block an algorithm FFS-ROC that implements a weak version of Uniform Consensus that we call *RO Consensus*. We call it RO Consensus because its definition resembles the one of Consensus. It is tailored, however, to fulfill the requirements of FFS-WLE. Note that use the prefix ‘‘FFS-’’ (FF-Stability) to distinguish the algorithms in this section from the ones of Section 7. Recall that in Section 7 we present an algorithm that also satisfies E-Stability.

We assume that each process p_i has an initial value $p_i.a \in V \cup \{\perp\}$, where V is the set of initial values, and a decision value $p_i.d[1 \dots n]$, where $p_i.d[j] \in V \cup \{\perp\}$. We use $v \in p_i.d$ to denote that there is some $p_\ell \in \Pi$ such that $p_i.d[\ell] = v$. If a process p_i crashes, then we assume that its decision value $p_i.d$ is \mathcal{N} , where \mathcal{N} stands for the n element list $[\perp, \dots, \perp]$. To avoid repetition throughout the discussion of our algorithm, we say that a process p decides in an execution ϕ if $p.d$ is different than \mathcal{N} .

As we describe later, we execute FFS-ROC multiple times in electing a leader. We then have that processes

may crash before starting an execution ϕ of FFS-ROC. Such processes consequently have initial value undefined in ϕ . We therefore use \perp to denote the initial value of crashed processes. That is, if $p_i.a = \perp$, then p_i has crashed.

Let the relation $x \subseteq y$ for x and y lists of n elements be that, for all $i : 1 \leq i \leq n$, $(x[i] \neq \perp) \Rightarrow (x[i] = y[i])$.

The specification of RO Consensus is given by four properties as follows:

Termination: Every process that does not crash eventually decides on some value.

Agreement: If $p_i.d[\ell] \neq \perp$, then for every non-faulty p_c , $p_i.d[\ell] = p_c.d[\ell]$;

RO Uniformity: Let *vals* be $\{d : \exists p_i \in \Pi \text{ s.t. } (p_i.d = d)\} \setminus \mathcal{N}$. Then,

$$\begin{aligned} & \wedge 1 \leq |\text{vals}| \leq 2 \\ & \wedge \forall d, d' \in \text{vals} : d \subseteq d' \vee d' \subseteq d \\ & \wedge \forall d_f, d_c \in \text{vals}, d_f \subseteq d_c : \exists S_f, S_c \in \mathcal{S}_\Pi : \\ & \quad \wedge \forall p \in S_f : \vee p \text{ crashes} \\ & \quad \quad \vee p.d = d_f \\ & \quad \wedge \forall p \in S_c : \wedge p.d = d_c \\ & \quad \quad \wedge p \text{ is not faulty} \end{aligned}$$

That is, there can be no more than two non- \mathcal{N} decision values, and if there are two then one is a subset of the other. Furthermore, if there are two different decision values, then these are the values that processes in two disjoint survivor sets decide upon, one for the processes of each survivor set.

Validity:

$$\begin{aligned} & \wedge \text{If } p_j \in \Pi \text{ does not crash, then for all non-faulty } \\ & \quad p_i, p_i.d[j] = p_j.a \\ & \wedge \text{If } p_j \in \Pi \text{ does crash, then exists } v \in \{\perp, p_j.a\} \\ & \quad \text{such that for all non-faulty } p_i, p_i.d[j] = v; \\ & \wedge \text{If there are survivor sets } S_f, S_c \in \mathcal{S}_\Pi \\ & \quad \text{and values } v_f, v_c \in V, v_f \neq v_c, \text{ such that} \\ & \quad \quad \wedge \forall p \in S_f : p.a \in \{v_f, \perp\} \\ & \quad \quad \wedge \forall p \in S_c : \wedge p.a = v_c \\ & \quad \quad \quad \wedge p \text{ is not faulty} \\ & \quad \quad \wedge \exists p_i, p_\ell \in \Pi : p_i.d[\ell] = v_f \\ & \quad \text{then for all } p_j \text{ that does not crash, } v_f \in p_j.d \end{aligned}$$

That is, if a process p_i is not faulty and $p_i.d[j] \neq \perp$, then the value of $p_i.d[j]$ must be $p_j.a$. The value of $p_i.d[j]$ can be \perp only if p_j crashes. The third case exists because we use the decision values of an execution as the initial values for another execution. From RO Uniformity, there can be two different non- \mathcal{N} values d_f and d_c . If this is the case, then there is

a survivor set S_c containing only correct processes such that all processes in S_c decide upon d_c , and another survivor set S_f containing only faulty processes such that all the processes in S_f either crash or decide upon d_f . Let d_f be v_f and d_c be v_c . By the third case, if some process that decides includes v_f in its decision value, then every process that does not crash also includes v_f in its decision value.

We now describe our algorithm FFS-ROC for RO Consensus. Figure 1 shows the pseudocode for a single process. A TLA+ specification of the algorithm appears in the Appendix. From the figure, the algorithm FFS-ROC executes exactly in $t+1$ rounds, where $t = \min_s \{s = |S_i| \wedge S_i \in \mathcal{S}_\Pi\}$ or alternatively $t+1 = \max_c \{c = |C_i| \wedge C_i \in \mathcal{C}_\Pi\}$.² For the proof of correctness we present in the next section, we assume that the value of t is at least one ($t \geq 1$). Note that for $t = 0$ there is a trivial, much simpler algorithm.

In every round r of FFS-ROC, a process p_i sends its list $p_i.A$ of values to a subset of the processes Π' in Π . If process p_i does not crash or stop³ in round r , then $\Pi = \Pi'$. Otherwise, this subset is arbitrary. Before the end of round r , every process p_i that does not execute a crash step at r receives all the messages sent to it at round r . Note that if a process p_i crashes at round r , but sends a message m_i to process p_j , then p_j does not necessarily receive m_i by C-Liveness. We then use the following, where $0 \leq r \leq t$:

- $p_i.M(r)$ denotes the set of messages of round r that p_i receives;
- $p_i.s(r)$ denotes the set of processes from which process p_i receives messages of round r . That is, $p_i.s(r) = \{p : m \in p_i.M(r) \wedge m.from = p\}$;
- $p_i.sr(r)$ denotes the set $p_i.s(r)$ removed the processes p_i detects to be faulty at round r .

Processes send no messages at the last round. Note that, by the algorithm, messages received by the end of round r are available for processing at the beginning of round $r+1$.

²The first and the last rounds in the algorithms are actually half rounds, and we then consider that both together constitute a single round. Put it another way, we can easily rearrange the order of sending and receiving messages to make it fit into $t+1$. We have chosen the former for expositional convenience.

³A *stop* instruction is equivalent to a crash in that a process does not execute any further steps after executing a stop instruction. A process, however, executes a stop instruction according to its own state machine, and hence it is not in any arbitrary step.

If a process detects that it has failed to receive messages, then it stops by deciding on \mathcal{N} . In the discussion that follows, we treat processes that crash and processes that stop in the same manner. If a distinction is necessary, then we clearly state it. There are two ways a processes p_i can determine that it is faulty: 1) By receiving messages from a set of processes at round r such that $p_i.s(r) \not\subseteq p_i.s(r-1)$, $r > 1$; 2) By determining that in its set of messages of round r , there is no survivor set possibly containing only correct processes. The second form of detection relies on the set of values p_i receives from another process p_j . If p_i notices that p_j did not receive a previous message from p_i , then p_i declares p_j faulty. By removing the obviously faulty processes and looking at the remaining set, if there is no survivor set in the remaining set, then p_i must be faulty as well. More specifically, p_i checks at round $r > 1$ whether $p_i.sr(r)$ contains a survivor set. To decide upon membership for $p_i.sr(r)$, p_i uses the value of $p_i.A$ from round $r-2$. We use $p_i.A_p(r)$ as the value of $p_i.A$ at round r after p_i updates $p_i.A$ with the values received in the messages from round r . Note that the value of $p_i.A_p(r)$, $0 \leq r \leq t-1$, is used only at round $r+2$.

Figure 2 shows the pseudocode for an algorithm that implements Safety, LE-Liveness, and FF-Stability. It proceeds in iterations of an infinite repeat loop. In each iteration, a process executes FFS-ROC twice, and decide if it has to elect itself by the end of the second phase.

In the following sections, we prove the correctness of both FFS-ROC and FFS-WLE.

5 Correctness of FFS-ROC

We provide a proof of correctness for the FFS-ROC algorithm. We say that a process p_i is alive at round r if either one of the following happens:

- if p_i sends at least one message m_i to some process p_j at round r , $0 \leq r \leq t$, and p_j receives m_i by the end of round r ;
- if p_i decides at round r , $r = t+1$.

We use $Live(r)$ to denote the processes that are alive at round r . For an execution $\phi = \langle F, I, S, T, \mathcal{T} \rangle$ of FFS-ROC we define the following to use in the proofs of this section:

- $T_\phi^i(i, r) \in T$ denotes the value of $\mathcal{T}(s_f)$, where $p_i \in Live(r)$ and $s_f \in S$ is the first step p_i executes of round r , $r \in \{0, \dots, t+1\}$. If p_i executes no steps in r , then $T_\phi^i(i, r)$ is undefined;

Algorithm FFS-ROC on input $p_i.a$

round 0:

$p_i.s(0) \leftarrow \Pi$; $p_i.sr(0) \leftarrow p_i.s(0)$
 $p_i.A[i] \leftarrow p_i.a$
for all $p_k \in \Pi$, $p_k \neq p_i$: $p_i.A[i] \leftarrow \perp$
 $p_i.A_p(0) \leftarrow p_i.A$
send $p_i.A$ to all

round 1:

$p_i.sr(1) \leftarrow p_i.s(1)$
if $\nexists S \in \mathcal{S}_\Pi$: $S \subseteq p_i.sr(1)$
then decide $[\perp, \dots, \perp]$
else for each message $m_j \in p_i.M(1)$, $p_k \in \Pi$:
if $(p_i.A[k] = \perp)$ $p_i.A[k] \leftarrow m_j.A[k]$
 $p_i.A_p(1) \leftarrow p_i.A$
send $p_i.A$ to all

round r : $2 \leq r \leq t$:

$p_i.sr(r) \leftarrow p_i.s(r) \setminus \{p_j : \exists m \in p_i.M(r) :$
 $(m.from = p_j) \wedge (p_i.A_p(r-2) \not\subseteq m.A)\}$
if $\forall p_i.s(r) \not\subseteq p_i.s(r-1)$
 $\forall \nexists S \in \mathcal{S}_\Pi$: $S \subseteq p_i.sr(r)$
then decide $[\perp, \dots, \perp]$
else for each message $m \in p_i.M(r)$, $p_k \in \Pi$:
if $(p_i.A[k] = \perp)$ $p_i.A[k] \leftarrow m.A[k]$
 $p_i.A_p(r) \leftarrow p_i.A$
send $p_i.A$ to all

round $t+1$:

$p_i.sr(t+1) \leftarrow p_i.s(t+1) \setminus \{p_j : \exists m \in p_i.M(t+1) :$
 $(m.from = p_j) \wedge (p_i.A_p(t-1) \not\subseteq m.A)\}$
if $\forall p_i.s(t+1) \not\subseteq p_i.s(t)$
 $\forall \nexists S \in \mathcal{S}_\Pi$: $S \subseteq p_i.sr(t+1)$
then decide $[\perp, \dots, \perp]$
else for each message $m \in p_i.M(t+1)$, $p_k \in \Pi$:
if $(p_i.A[k] = \perp)$ $p_i.A[k] \leftarrow m.A[k]$
 $p_i.A_p(t+1) \leftarrow p_i.A$
decide $p_i.A$

Figure 1: Algorithm run by process p_i .

Algorithm FFS-WLE

repeat {

$p_i.elected \leftarrow \text{FALSE}$

Phase 1:

Run FFS-ROC with

$p_i.a \leftarrow i$.

if $(p_i.d = [\perp, \dots, \perp])$ then stop

Phase 2:

Run FFS-ROC with

$p_i.a \leftarrow p_i.d$ from Phase 1.

if $(p_i.d = [\perp, \dots, \perp])$ then stop

let x be a value of $p_i.d[1 \dots n]$

such that $p_i.d[x] \neq [\perp, \dots, \perp]$

and it has the least number of non- \perp values

if $(p_i$ is the first index of x such that $x[i] \neq \perp$)

then $p_i.elected \leftarrow \text{TRUE}$

}

Figure 2: Algorithm run by process p_i .

- $T_\phi^u(i, r) \in T$ denotes the value of $\mathcal{T}(s_m)$, where $p_i \in \text{Live}(r)$ and $s_m \in S$ is the first step of round r in which p_i sends a message, $0 \leq r \leq t$. If $p_i \notin \text{Live}(r)$, then $T_\phi^u(i, r)$ is undefined;
- $T_\phi^u(i, t+1) \in T$ denotes $\mathcal{T}(s_d)$, where $p_i \in \text{Live}(t+1)$ and $s_d \in S$ is the step in which p_i decides at round $t+1$. If $p_i \notin \text{Live}(t)$, then $T_\phi^u(i, t+1)$ is undefined;
- M_i^r is a list of n values, one for each process in Π such that the following holds:

$$M_i^r[j] \begin{cases} p_j.a, & \exists m \in p_i.M(r) : m.A[j] \neq \perp \\ \perp, & \text{otherwise} \end{cases}$$

Processes that are alive in a round r may send messages to a strict subset of Π . Thus, in executions in which processes fail, we have that processes may have a different knowledge of the initial values. For example, if at round zero, a process p_ℓ sends a message to a non-faulty process p_i , but it crashes before sending a message to a non-faulty process p_j , then $p_i.A[\ell] = p_\ell.a$ and $p_j.A[\ell] = \perp$. For the purpose of analyzing these cases, we define a *process chain* (or simply a chain) $\omega_\ell = (i_0 \circ i_1 \circ \dots \circ i_k)_\ell$, $k \leq t+1$, to be a string over the set of process identifiers. Let $\omega_\ell[x]$ be the process identifier at position x of the chain ω_ℓ . The following holds for a process chain ω_ℓ :

1. $\omega_\ell[r] \neq \omega_\ell[r']$, if $r \neq r'$;

2. If $\omega_\ell[r] = i$, then $(p_i.A_p(r)[\ell] \neq \perp) \wedge (\forall r' \in \{x \in \mathbb{Z}^* : x \leq r - 1\} : p_i.A_p(r')[\ell] = \perp)$;
3. If $\omega_\ell[r] = i$, $r > 0$, then $\exists m \in p_i.M(r) : (m.A[\ell] \neq \perp) \wedge \omega_\ell[r - 1] = j \wedge m.from = p_j$;
4. If $\omega_\ell[0] = i$, then $i = \ell$.

We say that a process p_i is in ω_ℓ ($i \in \omega_\ell$) if and only if there exists an index r such that $\omega_\ell[r] = i$. We use process chains in the proofs below to represent the propagation of knowledge in executions with failures.

Proposition

5.1: {5.20, 5.21, 5.22, 5.23}

Theorems

5.20: {}

5.21: {5.16}

5.22: {5.5, 5.15, 5.17, 5.18}

5.23: {5.2, 5.5, 5.11, 5.16, 5.19}

Lemmas

5.2: {}

5.3: {5.2}

5.4: {5.2, 5.3}

5.5: {5.3, 5.4}

5.6: {5.5}

5.7: {5.4, 5.6}

5.8: {5.6}

5.9: {}

5.10: {5.9}

5.11: {5.2, 5.4, 5.5}

5.12: {5.6, 5.7, 5.10, 5.11}

5.13: {5.2, 5.4, 5.6, 5.7, 5.10, 5.11, 5.12}

5.14: {5.2, 5.4, 5.12}

5.15: {5.13}

5.16: {5.2, 5.4, 5.5, 5.11, 5.12}

5.17: {5.13, 5.16}

5.18: {5.5, 5.11, 5.13, 5.14, 5.15, 5.16, 5.17}

5.19: {5.3, 5.4, 5.13, 5.16}

Figure 3: FFS-ROC proof hierarchy.

We show the correctness of FFS-ROC with Proposition 5.1. We divide the proof of this proposition into several statements (lemmas and theorems). Each of these statements may depend upon others. Thus, we summarize in Figure 5 the structure of the proof. For each statement, we say the other statements it depends upon.

Proposition 5.1 FFS-ROC implements RO Consensus.

We prove Proposition 5.1 with the following lemmas.

Lemma 5.2 Let ϕ be an execution of FFS-ROC, r be a round of ϕ , $0 \leq r \leq t+1$, p_i be a process in $\text{Live}(r)$, and p_j be a process in Π . If $p_i.A_p(r)[j] \neq \perp$, then $p_i.A_p(r) = p_j.a$.

Proof:

We show with an induction on the round numbers ρ , $0 \leq \ell \leq r'$, that for every round $r' \leq r$, if $p_\ell \in \text{Live}(r')$ and $p_\ell.A_p(r')[j] \neq \perp$, then $p_\ell.A_p(r')[j] = p_j.a$. The base case is $\rho = 0$. From the algorithm, at round 0 a process p_ℓ has $p_\ell.A_p(0)[j] = \perp$, if $\ell \neq j$, and $p_\ell.A_p(0)[\ell] = p_\ell.a$. Thus, $p_\ell.A_p(0)[j] \neq \perp$ only if $\ell = j$, and $p_\ell.A_p(0)[\ell] = p_\ell.a$.

Now suppose that for every $p_\ell \in \text{Live}(\rho)$ if $p_\ell.A_p(\rho)[j] \neq \perp$, then $p_\ell.A_p(\rho)[j] = p_j.a$. We show that for every $p_\ell \in \text{Live}(\rho + 1)$, if $p_\ell.A_p(\rho + 1)[j] \neq \perp$, then $p_\ell.A_p(\rho + 1)[j] = p_j.a$. By the algorithm, if $p_\ell \in \text{Live}(\rho)$ is such that $p_\ell.A_p(\rho)[j] = p_j.a$, then for every message m it sends at round ρ , $m.A[j] = p_j.a$. For $p_{\ell'} \in \text{Live}(\rho + 1)$, if $p_{\ell'}.A[j] \neq \perp$ at $T_\phi^i(\ell', \rho + 1)$, then $p_{\ell'}.A_p(\rho)[j] \neq \perp$ and must be equal to $p_j.a$ by the induction hypothesis and the algorithm. Otherwise, for every message m it receives such that $m.A[j] \neq \perp$, $m.A[j] = p_j.a$. Thus, by the algorithm, if $p_{\ell'}$ receives at least one such a message, it sets $p_{\ell'}.A[j]$ to $p_j.a$, and we have that $p_{\ell'}.A_p(\rho + 1)[j] = p_j.a$.

From the previous induction, we conclude that if $p_i.A_p(r)[j] \neq \perp$, then $p_i.A_p(r)[j] = p_j.a$.

□

Lemma 5.3 Let ϕ be an execution of FFS-ROC, r be a round of ϕ , $0 < r \leq t + 1$, and p_i be a process in $\text{Live}(r)$. For every message $m \in p_i.M(r)$ such that $m.A[\ell] \neq \perp$, for some $p_\ell \in \Pi$, $m.A[\ell] = p_\ell.a$.

Proof:

By Lemma 5.2, for every $p_j \in \text{Live}(r - 1)$, if $p_j.A_p(r - 1)[\ell] \neq \perp$, then $p_j.A_p(r - 1)[\ell] = p_\ell.a$. If p_j sends a message m_j to p_i at round $r - 1$, then $m.A[\ell] = p_\ell.a$. We conclude that for every $m \in p_i.M(r)$ such that $m.A[\ell] \neq \perp$, $m.A[\ell] = p_\ell.a$.

□

Lemma 5.4 Let ϕ be an execution of FFS-ROC and r be a round of ϕ , $0 < r \leq t + 1$. For every $p_i \in \text{Live}(r)$, $M_i^r = p_i.A_p(r)$.

Proof:

By the algorithm, if $p_i \in \text{Live}(r)$, then for every $j \in [1 \dots n]$, such that $p_i.A[j] = \perp$ at $T_\phi^i(i, r)$, p_i sets $p_i.A[j]$ to a value $v \in V = \{v : v = m.A[j] \wedge m \in p_i.M(r) \wedge m.A[j] \neq \perp\}$, $j \in [1 \dots n]$ at t if $V \neq \emptyset$, where $T_\phi^i(i, r) \leq t \leq T_\phi^u(i, r)$, $t = \mathcal{T}(s)$, and s is a step of p_i that updates

$p_i.A[j]$ at round r . Otherwise, if $p_i.A[j] \neq \perp$ at $T_\phi^i(i, r)$, then no step of p_i in round r modifies the value of $p_i.A[j]$, and $p_i.A_p(r)[j] = p_i.A_p(r-1)[j]$.

Let p_j be a process of Π . By Lemma 5.3, we have that $V = \{v : v = m.A[j] \wedge m \in p_i.M(r) \wedge m.A[j] \neq \perp\}$, $j \in [1 \dots n]$, is either empty or contains a single value. If $|V| = 1$, then $V = \{p_j.a\}$. There are two cases to consider: 1) $p_i.A_p(r-1)[j] \neq \perp$; 2) $p_i.A_p(r-1) = \perp$.

If $p_i.A_p(r-1)[j] \neq \perp$, then, by the algorithm, p_i does not modify the value of $p_i.A[j]$ in round r . By Lemma 5.2, $p_i.A_p(r-1)[j] = p_j.a$. By the algorithm, p_i sends $p_i.A_p(r-1)$ to itself. Finally, by Lemma 5.3, every message $m \in p_i.M(r)$, $m.A[j] \neq \perp$, is such that $m.A[j] = p_j.a$. We then have that $p_i.A_p(r)[j] = M_i^r[j]$.

If $p_i.A_p(r-1) = \perp$ and $V = \{p_j.a\}$, then $p_i.A_p(r)[j] = p_j.a$. If $V = \emptyset$, then $p_i.A_p(r)[j] = \perp$. In both cases, $p_i.A_p(r)[j] = M_i^r[j]$.

We conclude that $p_i.A_p(r)$ must be equal to M_i^r .

□

Lemma 5.5 *Let ϕ be an execution of FFS-ROC. For every $r \in \{z \in \mathcal{R} : z \leq t+1\}$, $\text{Correct}(\phi) \subseteq \text{Live}(r)$.*

Proof:

By definition, a process is alive at round $t+1$ of ϕ if it neither crashes nor stops before deciding in this round. We show this claim by showing that for every $p_c \in \text{Correct}(\phi)$ and every $r \in \{x \in \mathcal{R} : x \leq t+1\}$, $p_c \in \text{Live}(r)$. Let p_c be a process in $\text{Correct}(\phi)$. By definition, p_c does not crash in any round. It remains to show that, for every $p_c \in \text{Correct}(\phi)$ and every $r \in \{z \in \mathcal{R} : z \leq t+1\}$, p_c does not stop in r . We show this with an induction on the round numbers $\rho, \rho \in \{z \in \mathcal{R} : z \leq t+1\}$.

By the algorithm, no process stops at round 0. At round 1, a process only stops if it does not receive messages from a survivor set. Let p_c be a process in $\text{Correct}(\phi)$. By definition, there is a survivor set S_c containing only correct processes, and every process $p_{c'} \in S_c$ sends a message to p_c at round 0. By C-Liveness, p_c must have messages in $p_c.M(1)$ at least from the processes in S_c . That is, $S_c \subseteq p_c.s(1)$. Consequently, p_c does not stop at round 1.

Now suppose that the claim holds for every $\rho, \rho \in \{z \in \mathcal{R} : 1 \leq z \leq t\}$, and we show for $\rho+1$. Let p_c be a process in $\text{Correct}(\phi)$. By assumption, if p_c does not receive a message from some process p_i in round ρ , then p_i must have crashed by round ρ . This implies that all processes in $\Pi \setminus p_c.s(\rho)$ crashed by round ρ , and $p_c.s(\rho+1)$ therefore cannot contain a process p_i that is not in $p_c.s(\rho)$. Consequently, $p_c.s(\rho+1) \subseteq p_c.s(\rho)$.

For the induction step, it remains to show that $p_c.sr(\rho+1)$ contains some survivor set. From the algorithm, a process p_i is in $p_c.sr(\rho+1)$ if there is a message $m_i \in p_c.M(\rho+1)$ and $p_c.A_p(\rho-1) \subseteq m_i.A$. By assumption, no correct process stops by round ρ . Thus, for every $p_{c'} \in \text{Correct}(\phi)$, there is a message $m_c \in p_{c'}.M(\rho)$ from p_c . By Lemma 5.3, for every $p_\ell \in \Pi$ such that $m_c.A[\ell] \neq \perp$, we have that $m_c.A[\ell] = p_c.A_p(\rho-1)[\ell] = p_\ell.a$ and $M_{c'}^\rho[\ell] = p_\ell.a$. By Lemma 5.4, we then have that for every $p_{c'} \in \text{Correct}(\phi)$, $p_c.A_p(\rho-1) \subseteq p_{c'}.A_p(\rho)$. By the algorithm and by the assumption that no correct process stops at round ρ , for every $p_{c'} \in \text{Correct}(\phi)$, $p_{c'}$ sends a message to p_c . By the observation that for every $p_{c'}$, $p_c.A_p(\rho-1) \subseteq p_{c'}.A_p(\rho)$, we have that $\text{Correct}(\phi) \subseteq p_c.sr(\rho+1)$. By assumption, there is a survivor set $S_c \in \mathcal{S}_\Pi$ such that $S_c \subseteq \text{Correct}(\phi)$. We conclude that $S_c \subseteq p_c.sr(\rho+1)$.

This concludes the proof of the lemma.

□

Lemma 5.6 *Let ϕ be an execution of FFS-ROC such that ω_ℓ is a chain in ϕ , $1 \leq |\omega_\ell| \leq t+1$. If $\omega_\ell[r]$ is the identifier of a correct process for some r , then $M_j^{r+1}[\ell] \neq \perp$ for every process $p_j \in \text{Correct}(\phi)$.*

Proof:

Let r be an index such that $\omega_\ell[r]$ is the identifier of a correct process in ϕ and i be the process identifier in $\omega_\ell[r]$. By the definition of a process chain, we have that $(p_i.A_p(r)[\ell] \neq \perp) \wedge (\forall r' \in \{x \in \mathcal{R} : x \leq r-1\} : p_i.A_p(r')[\ell] = \perp)$. By the algorithm, process p_i sends $p_i.A_p(r)$ to all the processes in Π . By Lemma 5.5, every correct process decides in ϕ ($\text{Correct}(\phi) \subseteq \text{Live}(t+1)$). By C-Liveness, for every $p_j \in \text{Correct}(\phi)$, there is $m_i \in p_j.M(r+1)$ such that $m_i.A[\ell] \neq \perp$. Again by the algorithm, $M_j^{r+1}[\ell]$ must be different than \perp for every correct process p_j in ϕ .

□

Lemma 5.7 *Let ϕ be an execution of FFS-ROC such that there is a chain ω_ℓ of length at least three in ϕ . There are no three correct processes $p_{c_1}, p_{c_2}, p_{c_3}$ such that $c_1, c_2, c_3 \in \omega_\ell$.*

Proof:

Proof by contradiction. Suppose that there are three processes $p_{c_1}, p_{c_2}, p_{c_3} \in \text{Correct}(\phi)$ such that $c_1, c_2, c_3 \in \omega_\ell$, and that r is the smallest index such that $\omega_\ell[r]$ is the identifier of a correct process. Observe that $|\omega_\ell|$ must be at least as large as $r+3$ ($|\omega_\ell| \geq r+3$), otherwise the claim is vacuously true.

Without loss of generality, let $\omega_\ell[r] = c_1$. By Lemma 5.6, for every correct process p_c in $\text{Correct}(\phi)$ we have that $M_c^{r+1}[\ell]$ different than \perp and by Lemma 5.4 $p_c.A_p(r+1)[\ell] = M_c^{r+1}[\ell]$. Consequently, we have that $p_{c_2}.A_p(r+1)[\ell] \neq \perp$ and $p_{c_3}.A_p(r+1)[\ell] \neq \perp$. By the definition of a process chain, c_2 and c_3 cannot be both in ω_ℓ , a contradiction.

□

Lemma 5.8 *Let ϕ be an execution of FFS-ROC such that there is a chain ω_ℓ of length at least three. There is no two correct processes p_i, p_j such that $i, j \in \omega_\ell$ and $\omega_\ell = (\omega' \circ i \circ \omega \circ j \circ \omega'')_\ell$, where $\omega, \omega', \omega''$ are substrings of ω_ℓ and ω is not the empty string.*

Proof:

Proof by contradiction. Suppose that there are two correct processes p_i and p_j in ϕ such that $\omega_\ell[r] = i$ and $\omega_\ell[r'] = j$, $r+1 < r'$. By Lemma 5.6 and by the definition of a process chain, for every correct process p_j in $\text{Correct}(\phi)$, $M_j^{r+1}[\ell] \neq \perp$. We hence have that r' must be equal to $r+1$, and ω must be empty, contradicting out initial assumption that $r+1 < r'$.

□

Lemma 5.9 *Let ϕ be an execution of FFS-ROC and p_i be a process in $\text{Live}(r)$, $r \geq 2$, such that $p_i.A_p(r)[\ell] \neq \perp$ and for all $r' \in \{x \in \mathcal{R} : x \leq r-1\}$, $p_i.A_p(r')[\ell] = \perp$. For every round $\rho \in \{x \in \mathcal{R} : 1 \leq x \leq r\}$, there are processes $p_{j_1} \in \text{Live}(\rho)$ and $p_{j_2} \in \text{Live}(\rho-1)$ such that the following holds:*

1. $p_{j_1}.A_p(\rho)[\ell] \neq \perp$, and for all $\rho' \in \{x \in \mathcal{R} : x \leq \rho-1\}$, $p_{j_1}.A_p(\rho')[\ell] = \perp$;
2. $p_{j_2}.A_p(\rho-1)[\ell] \neq \perp$, and for all $\rho' \in \{x \in \mathcal{R} : x \leq \rho-2\}$, $p_{j_2}.A_p(\rho')[\ell] = \perp$;
3. $\exists m_{j_2} \in p_{j_2}.M(\rho) : (m_{j_2}.A[\ell] \neq \perp)$.

Proof:

We now show with an induction on the values of ψ , $0 \leq \psi \leq r-1$, that for every round $\rho = r-\psi$, the claim holds.

The base case is $\psi = 0, \rho = r$. By assumption, p_i is such that $p_i.A_p(r)[\ell] \neq \perp$ and for all $r' \in \{x \in \mathcal{R} : x \leq r-1\}$, $p_i.A_p(r')[\ell] = \perp$. This implies that $M_i^r[\ell] \neq \perp$. From the algorithm, we have that $p_i.s(\rho) \subseteq p_i.s(\rho-1) \subseteq \dots \subseteq p_i.s(0)$. Consequently, there must be some process $p_j \in \text{Live}(\rho-1)$ such that the following holds: A) $p_j.A_p(r-1) \neq \perp$; B) $p_j.A_p(\rho') = \perp$ for all $\rho' \in \{x \in \mathcal{R} : x \leq r-2\}$;

C) $\exists m \in p_i.M(r) : (m.A[\ell] \neq \perp) \wedge (m.from = p_j)$. If there is no such a process p_j that satisfies both A) and C), then $p_i.A_p(r) = \perp$, contradicting our initial assumption. By the algorithm, once p_j sets the value of $p_j.A[\ell]$ to a value different than \perp , then the value of $p_j.A[\ell]$ does not change in subsequent rounds. This implies that for all $\rho', 0 \leq \rho' < r-1$, $p_j.A_p(\rho')$ must be equal to \perp , because by the algorithm p_i receives a message from p_j in every round ($p_i.s(\rho) \subseteq p_i.s(\rho-1)$, $r \geq \rho > 0$) and $p_i.A_p(\rho')$ is different than \perp otherwise, for some $\rho'' < r$.

Suppose the claim is true for $\psi < r-1$. We show for $\psi+1$. If it is true for ψ , then there is a process p_{j_1} alive in $\rho = r-(\psi+1)$ such that $p_{j_1}.A_p(\rho)[\ell] \neq \perp$, and for all $\rho' \in \{x \in \mathcal{R} : x \leq \rho-1\}$, $p_{j_1}.A_p(\rho')[\ell] = \perp$. From the algorithm, we have that $p_{j_1}.s(\rho) \subseteq p_{j_1}.s(\rho-1) \subseteq \dots \subseteq p_{j_1}.s(0)$. Consequently, there must be some process $p_{j_2} \in \text{Live}(\rho-1)$ such that the following holds: A) $p_{j_2}.A_p(\rho-1) \neq \perp$; B) $p_{j_2}.A_p(\rho') = \perp$ for all $\rho' \in \{x \in \mathcal{R} : x \leq \rho-2\}$; C) $\exists m \in p_{j_1}.M(\rho) : (m.A[\ell] \neq \perp) \wedge (m.from = p_{j_2})$. If there is no such a process p_{j_2} that satisfies both A) and C), then $p_{j_1}.A_p(\rho) = \perp$, contradicting our assumption that the hypothesis hold for ψ . By the algorithm, once p_{j_2} sets the value of $p_{j_2}.A[\ell]$ to a value different than \perp , then the value of $p_{j_2}.A[\ell]$ does not change in subsequent rounds. This implies that for all $\rho', 0 \leq \rho' < \rho-1$, $p_{j_2}.A_p(\rho')$ must be equal to \perp , because by the algorithm p_{j_1} receives a message from p_{j_2} at every round ($p_{j_1}.s(\rho) \subseteq p_{j_1}.s(\rho-1)$, $\rho \geq \rho > 0$) and $p_i.A_p(\rho'')$ is different than \perp otherwise, for some $\rho'' < \rho$.

This concludes the proof of the lemma.

□

Lemma 5.10 *Let ϕ be an execution of FFS-ROC and p_i be a process that is live at round r of ϕ , $r \geq 0$, such that $p_i.A_p(r)[\ell] \neq \perp$ and for all $r' \in \{x \in \mathcal{R} : x \leq r-1\}$, $p_i.A_p(r')[\ell] = \perp$. There is a chain ω_ℓ such that $|\omega_\ell| = r+1$, and $\omega_\ell[r] = i$.*

Proof:

We have to build a chain ω_ℓ such that $|\omega_\ell| = r+1$, and $\omega_\ell[r] = i$.

We build such a chain ω_ℓ as follows:

$$\begin{aligned} \omega_\ell[0] &= \ell \\ \omega_\ell[\rho] &= j \quad , \quad \wedge (0 < \rho < r) \\ &\quad \wedge (p_j.A_p(\rho) \neq \perp) \\ &\quad \wedge (\forall \rho' \in \{x \in \mathcal{R} : x \leq \rho-1\} : p_j.A_p(\rho') = \perp) \\ &\quad \wedge \exists m_j \in p_{\omega_\ell[\rho+1]}.M(\rho+1) \text{ from } p_j \\ \omega_\ell[r] &= i \end{aligned}$$

We can easily verify that ω_ℓ satisfies the properties of a process chain. It remains to show that it is a valid construction.

By the algorithm, we have that $\omega_\ell[0] = \ell$. By Lemma 5.9, we have that for every ρ , $0 < \rho < r$, there is a process p_j that satisfies the properties we stated above. Finally, by assumption, p_i is such that $p_i.A_p(r)[\ell] \neq \perp$ and for all $r' \in \{x \in \mathcal{R} : x \leq r-1\}$, $p_i.A_p(r')[\ell] = \perp$.

This concludes the proof of the lemma.

□

Lemma 5.11 *Let ϕ be an execution of FFS-ROC. If $p_i, p_j \in \text{Correct}(\phi)$, then $p_i.d = p_j.d$ in ϕ .*

Proof:

By Lemma 5.5, every correct process decides in ϕ (no correct process stops). Now let r , $0 \leq r \leq t$, be a round in which no process crashes. Such a round exists in ϕ by assumption (no more than t processes can fail in an execution, where t is $|\Pi|$ subtracted the size of the smallest survivor set).

We first show by induction on the values of ρ , $r+1 \leq \rho \leq t+1$, the following proposition:

$$\begin{aligned} \bigwedge \quad & \forall p_{c_1}, p_{c_2} \in \text{Correct}(\phi) : p_{c_1}.A_p(\rho) = p_{c_2}.A_p(\rho) \\ \bigwedge \quad & \forall p_\ell \in \text{Live}(\rho), p_c \in \text{Correct}(\phi) : p_\ell.A_p(\rho) \subseteq p_c.A_p(\rho) \end{aligned}$$

The base case is $\rho = r+1$. According to the algorithm, every process p_i that is live at round r sends a message containing $p_i.A_p(r)$ to every other process. According to C-Liveness and the assumption that no process crashes in round r , for every process $p_c \in \text{Correct}(\phi)$, $p_c.s(r+1) = \text{Live}(r)$. This implies that for every $p_{c_1}, p_{c_2} \in \text{Correct}(\phi)$, $M_{c_1}^{r+1} = M_{c_2}^{r+1}$. By Lemma 5.4, $M_c^{r+1} = p_c.A_p(r+1)$ for every $p_c \in \text{Correct}(\phi)$. This implies that for every $p_{c_1}, p_{c_2} \in \text{Correct}(\phi)$, $p_{c_1}.A_p(r+1) = p_{c_2}.A_p(r+1)$.

It remains to show the second part of the proposition for the base case. Let p_ℓ be a process in $\text{Live}(r+1)$ and p_c be a process in $\text{Correct}(\phi)$. By the failure assumptions, we have that $p_\ell.s(r+1) \subseteq \text{Live}(r)$. This implies that $p_\ell.s(r+1) \subseteq p_c.s(r+1)$. If $p_\ell.s(r+1) \subseteq p_c.s(r+1)$, then $M_\ell^{r+1} \subseteq M_c^{r+1}$. By Lemma 5.4, $M_\ell^{r+1} = p_\ell.A_p(r+1)$ and $M_c^{r+1} = p_c.A_p(r+1)$, which implies that $p_\ell.A_p(r+1) \subseteq p_c.A_p(r+1)$. This concludes the proof of the base case.

Suppose that the proposition holds for every $\rho < t+1$. We show for $\rho+1$. By the induction hypothesis, the algorithm, and Lemma 5.2, for every process $p_c \in \text{Correct}(\phi)$, $p_c.A_p(\rho) = M_c^{\rho+1}$. By Lemma 5.4, for every $p_c \in \text{Correct}(\phi)$, $p_c.A_p(\rho+1) = M_c^{\rho+1}$. We conclude that for every $p_{c_1}, p_{c_2} \in \text{Correct}(\phi)$, $p_{c_1}.A_p(\rho+1) = p_{c_2}.A_p(\rho+1)$.

By our failure assumptions, a faulty process may receive an arbitrary subset of the messages sent to it in a round. Let p_ℓ be a process in $\text{Live}(\rho+1)$ and p_c be a process in $\text{Correct}(\phi)$. By the induction hypothesis, the algorithm, and Lemma 5.2, $M_\ell^{\rho+1} \subseteq M_c^{\rho+1}$. By Lemma 5.4, $p_\ell.A_p(\rho+1) = M_\ell^{\rho+1}$ and $p_c.A_p(\rho+1) = M_c^{\rho+1}$. We conclude that $p_\ell.A_p(\rho+1) \subseteq p_c.A_p(\rho+1)$. This concludes the proof of the induction step.

From the previous proposition, we have that $p_i.A_p(t+1) = p_j.A_p(t+1)$. By the algorithm, p_i decides upon $p_i.A_p(t+1)$ and p_j decides upon $p_j.A_p(t+1)$. Consequently, $p_i.d = p_j.d$. This concludes the proof of the lemma.

□

Lemma 5.12 *Let ϕ be an execution of FFS-ROC and p_i, p_j be two processes in $\text{Live}(t+1)$, and p_ℓ be a process in Π . If $(p_i.A_p(t+1)[\ell] \neq \perp)$ and for all $r \in \{x \in \mathcal{R} : x \leq t\}$, $(p_i.A_p(r)[\ell] = \perp)$, then $p_j.d[\ell] \neq \perp$.*

Proof:

By the algorithm, once p_j sets the value of $p_j.A[\ell]$ to a value different than \perp , p_j does not change it in subsequent rounds. Thus, we only need to show that there is some round r in which p_j sets $p_j[\ell]$ to a value different than \perp .

Suppose that $((p_i.A_p(t+1)[\ell] \neq \perp) \wedge (\forall r \in \{z \in \mathcal{R} : z \leq t\} : p_i.A_p(r)[\ell] = \perp))$. Assuming that p_i and p_j can be either correct or faulty, there are four possible cases, and we analyze each case separately as follows:

- p_i and p_j are correct in ϕ . By Lemma 5.11, we have that $p_i.d = p_j.d$;
- p_i is faulty and p_j is correct in ϕ . If p_i decides in ϕ , then p_i is in $\text{Live}(t+1)$. By Lemma 5.10, there is a chain ω_ℓ such that $|\omega_\ell| = t+2$, and $\omega_\ell[t+1] = i$. Since there are at most t failures by assumption, there is at least one correct process in ω_ℓ . Moreover, such correct process must be in a position r of the chain such that $0 \leq r \leq t$. Thus, $p_j.A_p(t+1)[\ell]$ must be different than \perp , by Lemma 5.6 and the algorithm;
- p_i is correct and p_j is faulty in ϕ . If p_i is correct, then, by Lemma 5.10, there is a chain ω_ℓ such that $|\omega_\ell| = t+2$, and $\omega_\ell[t+1] = i$. Because p_i is correct, one of the following two must happen: 1) for every r , $0 \leq r \leq t-1$ and $x = \omega_\ell[r]$, p_x crashes at round r of ϕ ; 2) $p_i.A_p(r)[\ell] \neq \perp$ for some $r < t+1$ by Lemma 5.6 (there is a correct process in the chain). Case 1 cannot happen because $p_j \in \text{Live}(t+1)$ by

assumption, and there are at least $t + 1$ faulty processes, violating our assumptions for survivor sets. In case 2, p_i learns the initial value of p_ℓ in an earlier round, contradicting our initial assumption. We conclude that this case is hence not possible;

- p_i and p_j are faulty in ϕ . By Lemma 5.10, there is a chain ω_ℓ such that $|\omega_\ell| = t + 2$, and $\omega_\ell[t + 1] = i$. By Lemma 5.7, there are at most two correct processes in any chain. Thus, ω_ℓ contains t faulty processes. Consequently, there must be an r , $0 \leq r \leq t$, such that $j = \omega_\ell[r]$, and $p_j.A_p(r) \neq \perp$.

From the previous analysis, we have that either $p_i.d = p_j.d$ or $p_j.A_p(t + 1) \neq \perp$. By the algorithm, we have p_j decides upon $p_j.A_p(t + 1)$, and in both cases $p_j.d[\ell] \neq \perp$. This concludes the proof of the lemma.

□

Lemma 5.13 *Let ϕ be an execution of FFS-ROC, p_i, p_j be two processes in $\text{Live}(t + 1)$, and S_i, S_j be two survivor sets in \mathcal{S}_Π such that for all $r \in \{z \in \mathcal{R} : z \leq t + 1\}$, $S_i \subseteq p_i.sr(r)$, $S_j \subseteq p_j.sr(r)$, and $S_i \cap S_j \neq \emptyset$. $p_i.d = p_j.d$ in ϕ .*

Proof:

By the algorithm, once a process p_i sets the value of $p_i.A[\ell]$ at round $r \in \{z \in \mathcal{R} : z \leq t\}$ to a value different than \perp , it does not change it in subsequent rounds. We then have to show that if p_i learns about the initial value of p_ℓ at round r (that is, $p_i.A_p(r)[\ell] \neq \perp$ and for all $r' \in \{z \in \mathcal{R} : z \leq r - 1\}$, $p_i.A_p(r')[\ell] \neq \perp$), then there is a round r' such that p_j learns the initial value of p_ℓ at round r' (that is, $p_j.A_p(r')[\ell] \neq \perp$ and for all $r'' \in \{z \in \mathcal{R} : z \leq r' - 1\}$, $p_j.A_p(r'')[\ell] \neq \perp$). By Lemma 5.2, if $p_i.A_p(r)[\ell] = p_j.A_p(r')[\ell] \neq \perp$, then $p_i.A_p(r)[\ell] = p_j.A_p(r)[\ell] = p_\ell.A_p(0)[\ell]$. We now analyze each case separately.

First, suppose that $((p_i.A_p(t + 1)[\ell] \neq \perp) \wedge (\forall r \in \{z \in \mathcal{R} : z \leq t\} : p_i.A_p(r)[\ell] = \perp))$. This follows directly from Lemma 5.12.

Now, suppose that $(p_i.A_p(t)[\ell] \neq \perp) \wedge (\forall r \in \{z \in \mathcal{R} : z \leq t - 1\} : p_i.A_p(r)[\ell] = \perp)$:

- p_i and p_j are correct in ϕ . By Lemma 5.11, we have that $p_i.d = p_j.d$.
- p_i is faulty and p_j is correct in ϕ . From Lemma 5.10, there is a chain ω_ℓ such that $|\omega_\ell| = t + 1$, and $\omega_\ell[t] = i$. Because there are at most t faulty processes by assumption, there must be a correct process in ω_ℓ . That is, there must be some r , $0 \leq r \leq$

$t - 1$, such that $\omega_\ell[r]$ is the identifier of a correct process in ϕ . It follows that $p_j.A_p(r + 1)$ must be different than \perp , by Lemma 5.6.

- p_i is correct and p_j is faulty in ϕ . From Lemma 5.10, there is a chain ω_ℓ such that $|\omega_\ell| = t + 1$, and $\omega_\ell[t] = i$. Because p_j is faulty, either there is some r such that $\omega_\ell[r] = j$ or there are at most $t - 1$ faulty processes in ω_ℓ . If the former holds, then we are done. If the latter holds, then $\omega_\ell[t - 1]$ must be the identifier of a correct process, and there is no $r < t - 1$ such that $\omega_\ell[r]$ is the identifier of a correct process. Otherwise there is some $r \in \{z \in \mathcal{R} : z \leq t - 1\}$ such that $p_i.A_p(r)[\ell] \neq \perp$ (by Lemma 5.6). In addition, because ω_ℓ contains $t - 1$ faulty processes and p_j is faulty, any $p_x \in (S_i \cap S_j)$ is either correct or is in the chain ω_ℓ . Thus, $p_x.A_p(t) \neq \perp$, for every $p_x \in (S_i \cap S_j)$. Since by assumption $S_j \subseteq p_j.sr(r)$ for every $r \in \{z \in \mathcal{R} : z \leq t + 1\}$, we have that $p_j.A_p(t + 1) \neq \perp$, by the algorithm and Lemma 5.4;

- p_i and p_j are faulty in ϕ . From Lemma 5.10, there is a chain ω_ℓ such that $|\omega_\ell| = t + 1$, and $\omega_\ell[t] = i$. Because ω_ℓ contains exactly $t + 1$ process identifiers, at least one must be correct, and by Lemma 5.7, at most two correct processes. We then have that either there is some r such that $\omega_\ell[r] = j$ or $\omega_\ell[r] \neq j$ for every r . In the former case, we have that $p_j.A_p(r)[\ell] \neq \perp$ for some $r < t$. In the latter, ω_ℓ must contain $t - 1$ faulty processes (at most two correct processes and p_j is not in ω_ℓ). Let p_x be a process in $S_i \cap S_j$. p_x is either correct or faulty. Suppose p_x is correct. Because there is some correct process in $\omega_\ell[r]$, for $r \in \{z \in \mathcal{R} : z \leq t - 1\}$, by Lemma 5.6, it must be the case that $p_x.A_p(t - 1)[\ell] \neq \perp$ and consequently $p_j.A_p(t)[\ell] \neq \perp$, by the algorithm and Lemma 5.4.

Now suppose that p_x is faulty. In this case, either there is $r \in \{z \in \mathcal{R} : z \leq t\}$ such that $\omega_\ell[r] = x$ or $x = j$. The case that $\omega_\ell[r] = x$ is straightforward. If $x = j$, then either $\{p_j\} = S_j$ or $\{p_j\} \subset S_j$. Suppose the former. If S_j is a singleton set, then $t = |\Pi| - 1$ and $t + 1 = |\Pi|$. In this case, all the processes in Π must be in the chain ω_ℓ , and hence it must be the case that $p_j \in \omega_\ell$. Thus, S_j must contain at least two processes. Let $p_{j'}$ be a process in S_j such that $j \neq j'$. We then have that either $j' \in \omega_\ell$ or $p_{j'} \in \text{Correct}(\phi)$. In either case, there must be some round $r \in \{z \in \mathcal{R} : z \leq t\}$ such that $p_j.A_p(r)[\ell] \neq \perp$, and $p_j.A_p(t + 1)[\ell] \neq \perp$ by the algorithm.

Finally, suppose that $\exists r \in \{z \in \mathcal{R} : z \leq t-1\} : (p_i.A_p(r)[\ell] \neq \perp) \wedge (\forall r' \in \{z \in \mathcal{R} : z \leq r-1\} : p_i.A_p(r')[\ell] = \perp)$. By assumption $S_i \subseteq p_i.sr(\rho)$ for all $\rho \in \{z \in \mathcal{R} : z \leq t+1\}$. This implies by the algorithm that $p_x.A_p(r+1) \subseteq p_i.A_p(r+2)$, $p_x \in S_i \cap S_j$. Because $S_j \subseteq p_j.sr(\rho)$, for all $\rho \in \{z \in \mathcal{R} : z \leq t+1\}$, and $p_x \in S_j$, we then have by the algorithm and Lemma 5.4 that $p_j.A_p(r+2)[\ell]$ must be different than \perp , and equal to $p_i.A_p(r)[\ell]$ by Lemma 5.2.

From the previous argument, we conclude that $p_i.A_p(t+1) = p_j.A_p(t+1)$. By the algorithm, we have that p_i decides upon $p_i.A_p(t+1)$ and p_j decides upon $p_j.A_p(t+1)$. Again by the algorithm, we have that $p_i.d = p_j.d$.

□

Lemma 5.14 *Let ϕ be an execution of FFS-ROC and p_i, p_j be two processes in $\text{Live}(t+1)$ such that $p_j \in p_i.sr(r)$ for every $r \in \{z \in \mathcal{R} : z \leq t+1\}$. $p_j.d \subseteq p_i.d$ in ϕ .*

Proof:

By the algorithm, once a process p_j sets the value of $p_j.A[\ell]$ to a value different than \perp in a round r , for some $p_\ell \in \Pi$ and some $0 \leq r \leq t+1$, it does not change it in subsequent rounds. If $p_j.A_p(t+1)[\ell] \neq \perp$, then there is some round ρ , $0 \leq \rho \leq t+1$, such that $p_j.A_p(\rho)[\ell] \neq \perp$ and for all $\rho' \in \{z \in \mathcal{R} : z \leq \rho-1\}$, $p_j.A_p(\rho')[\ell] = \perp$. We then have to show that for every $p_\ell \in \Pi$ such that $p_j.A_p(t+1)[\ell] \neq \perp$, there is some ϱ such that $p_i.A_p(\varrho)[\ell] \neq \perp$, $\varrho \in \{z \in \mathcal{R} : z \leq t+1\}$.

Let p_ℓ be a process such that $p_j.A_p(\rho)[\ell] \neq \perp$ and for all $\rho' \in \{x : 0 \leq x < \rho\}$, $p_j.A_p(\rho')[\ell] = \perp$. Suppose that $\rho = t+1$. This case follows directly from Lemma 5.12. Now suppose that $\rho \leq t$. Because p_j sends a message to p_i in every round by assumption, $M_i^{\rho+1}[\ell]$ must be different than \perp , and $p_i.A_p(\rho+1)[\ell] = M_i^{\rho+1}[\ell]$ by Lemma 5.4. We conclude that if $p_j.A_p(t+1)[\ell] \neq \perp$, for some $p_\ell \in \Pi$, then $p_i.A_p(t+1)[\ell] \neq \perp$. By Lemma 5.2, $p_i.A_p(t+1)[\ell] = p_j.A_p(t+1)[\ell] = p_\ell.a$. By the algorithm, p_i decides upon $p_i.A_p(t+1)$ and p_j decides upon $p_j.A_p(t+1)$. Consequently, $p_j.d \subseteq p_i.d$.

□

Lemma 5.15 *Let ϕ be an execution of FFS-ROC. If p_i, p_j , and p_ℓ decide in ϕ , then either $p_i.d = p_j.d$, $p_i.d = p_\ell.d$, or $p_j.d = p_\ell.d$.*

Proof:

If p_i, p_j , and p_ℓ decide in ϕ , then there are survivor sets S_i, S_j , and S_ℓ such that $S_i \subseteq p_i.sr(r)$, $S_j \subseteq p_j.sr(r)$, and

$S_\ell \subseteq p_\ell.sr(r)$, for all r , $0 \leq r \leq t+1$. By the (3,2)-Intersection property, either $S_i \cap S_j \neq \emptyset$, $S_i \cap S_\ell \neq \emptyset$, or $S_j \cap S_\ell \neq \emptyset$. By Lemma 5.13, we then have that either $p_i.d = p_j.d$, $p_i.d = p_\ell.d$, or $p_j.d = p_\ell.d$.

□

Lemma 5.16 *Let ϕ be an execution of FFS-ROC and p_i be a correct process in ϕ . If p_j decides in ϕ , then $p_j.d \subseteq p_i.d$.*

Proof:

By Lemma 5.5, $p_i \in \text{Live}(t+1)$ (p_i decides in ϕ). If p_j is correct, then the lemma follows from Lemma 5.11. Now suppose that p_j commits at least one receive-omission fault in ϕ . By assumption, both p_i and p_j decide in ϕ . By the algorithm, $p_j \in p_i.s(r)$ for every r , $0 \leq r \leq t+1$. Because p_i is correct, we have that there is m from p_j in $p_i.M(r)$ for every r , $0 \leq r \leq t+1$. By Lemma 5.4 and the algorithm, we then have that if $p_j.A_p(r)[\ell] \neq \perp$, for some $p_\ell \in \Pi$ and $0 \leq r \leq t$, then $p_i.A_p(r+1)[\ell] = p_j.A_p(r)[\ell]$. It remains to show that if $p_j.A_p(t+1)[\ell] \neq \perp$, and $p_j.A_p(r)[\ell] = \perp$, $p_\ell \in \Pi$, for every $r \in \{z \in \mathcal{R} : z \leq t\}$, then $p_i.A_p(t+1)[\ell] = p_j.A_p(t+1)[\ell]$. By Lemma 5.12, we have that if $p_j.A_p(t+1)[\ell] \neq \perp$, then $p_i.A_p(t+1)[\ell] \neq \perp$. By Lemma 5.2, $p_i.A_p(t+1)[\ell] = p_j.A_p(t+1)[\ell] = p_\ell.a$. We conclude that $p_j.d \subseteq p_i.d$.

□

Lemma 5.17 *Let ϕ be an execution of FFS-ROC. If there are two processes p_i and p_j , $p_i, p_j \in \text{Live}(t+1)$, then either $p_i.d \subseteq p_j.d$ or $p_j.d \subseteq p_i.d$.*

Proof:

If at least one of p_i and p_j is correct, then the proof follows from Lemma 5.16. Now suppose both p_i and p_j are faulty. Because both p_i and p_j decide in ϕ by assumption, there are survivor sets S_i and S_j such that $(S_i \subseteq p_i.sr(r)) \wedge (S_j \subseteq p_j.sr(r))$ for every r , $0 \leq r \leq t+1$. If $S_i \cap S_j \neq \emptyset$, then the lemma follows because $p_i.d = p_j.d$ by Lemma 5.13. Suppose now the contrary: $S_i \cap S_j = \emptyset$. By assumption, there must be a survivor set S_c containing only correct processes. By the (3,2)-Intersection property, either $S_i \cap S_c \neq \emptyset$ or $S_j \cap S_c \neq \emptyset$. Let p_c be a process in S_c . We then have by Lemma 5.13 that either p_i and p_c decide upon the same value or p_j and p_c decide upon the same value. Suppose without loss of generality that p_i and p_c decide upon the same value. We hence have from Lemma 5.16 that $p_j.d \subseteq p_i.d$. This concludes the proof of the lemma.

□

Lemma 5.18 Let ϕ be an execution and vals be $\{d : \exists p_i \in \Pi \text{ s.t. } (p_i.d = d)\} \setminus \mathcal{N}$. For every $d_f, d_c \in \text{vals}$, $d_f \subseteq d_c$, there are survivor sets $S_f, S_c \in \mathcal{S}_\Pi$ such that the following properties hold:

$$\begin{aligned} \bigwedge \forall p \in S_f : & \quad \vee \quad p \text{ crashes} \\ & \quad \vee \quad p.d = d_f \\ \bigwedge \forall p \in S_c : & \quad \wedge \quad p.d = d_c \\ & \quad \wedge \quad p \text{ is not faulty} \end{aligned}$$

Proof:

By Lemma 5.5, $\text{Correct}(\phi) \subseteq \text{Live}(t+1)$. By the algorithm, every non-faulty process p_c is such that $p_c.d[\ell] = p_c.a$. We then have that vals contains at least one value. By Lemma 5.15, there cannot be three different decision values, and if there are two values d and d' , then either $d \subseteq d'$ or $d' \subseteq d$ by Lemma 5.17. We analyze these two cases separately.

First, suppose that vals contains a single value, say d , and $d_f = d_c = d$. By assumption, there is a survivor set S_i such that S_i contains only non-faulty processes. By Lemma 5.11, every process $p_i \in S_i$ is such that $p_i.d = d$. If we make $S_c = S_f = S_i$, then our claim holds.

Now suppose that vals contains two distinct values d_f and d_c , $d_f \subseteq d_c$. Let p_i be a process such that $p_i \in \text{Live}(t+1)$ and $p_i.d = d_f$. By the algorithm, there is survivor set S_i such that $S_i \subseteq p_i.sr(r)$, for every $r \in \{z \in \mathcal{R} : z \leq t+1\}$. Let p_j be a process in S_i . If $p_j \in \text{Live}(t+1)$, then there is an $S_j \in \mathcal{S}_\Pi$ such that $S_j \subseteq p_j.sr(r)$, for every $r \in \{z \in \mathcal{R} : z \leq t+1\}$. Now let S'_c be a survivor set such that $S'_c \subseteq \text{Correct}(\phi)$. By the (3,2)-Intersection property, either $S_j \cap S'_c \neq \emptyset$ or $S_j \cap S_i \neq \emptyset$. Note that $S_i \cap S'_c$ must be empty, otherwise $p_i.d = d_c$ according to Lemma 5.13, contradicting our initial assumption.

If $S_j \cap S'_c \neq \emptyset$, then by Lemma 5.13 we have that $p_j.d = d_c$ because $p_j.d = p_c.d$ for every $p_c \in S'_c$ (Lemma 5.13) and $p_c.d, p_c \in S'_c$, must be equal to d_c (Lemma 5.16). By Lemma 5.14, however, we have that $p_j.d \subseteq p_i.d$. This implies that $p_i.d = d_c$, again contradicting our initial assumption. It therefore must be the case that $S_i \cap S_j$ is not empty. By Lemma 5.13, we have that $p_i.d = p_j.d = d_f$.

Now suppose that $p_j \notin \text{Live}(t+1)$. We then have that p_j crashes in ϕ , and $p_j.d = \mathcal{N}$. We therefore have that $p_j \in S_i$ either decides upon d_f or crashes in ϕ .

It remains to show the second part of the properties in the statement of the lemma. By Lemma 5.16, every cor-

rect process must decide upon d_c . Thus, every process p_c in S is such that $p_c.d = d_c$ in ϕ .

To conclude, if we make $S_f = S_i$ and $S_c = S'_c$, then our claim holds.

□

Lemma 5.19 Let ϕ be an execution of FFS-ROC such that there are survivor sets $S_f, S_c \in \mathcal{S}_\Pi$ and values $v_f, v_c \in V$, $v_f \neq v_c$, such that the following holds:

$$\begin{aligned} \bigwedge \forall p \in S_f : & \quad p.a \in \{v_f, \perp\} \\ \bigwedge \forall p \in S_c : & \quad p.a = v_c \\ \bigwedge \forall p \in S_c : & \quad p \text{ is not faulty} \end{aligned}$$

If exists $p_i, p_\ell \in \Pi$ such that $p_i.d[\ell] = v_f$, then for all $p_j \in \text{Live}(t+1)$, $v_f \in p_j.d$.

Proof:

Suppose that $p_i.d[\ell] = v_f$, for some $p_i, p_\ell \in \Pi$. By the algorithm, if a process p_j does not crash or stop in an execution of FFS-ROC, then there is some survivor set S_j such that $S_j \subseteq p_j.sr(r)$ for every $r \in \{z \in \mathcal{R} : z \leq t+1\}$. S_f and S_c must be disjoint, otherwise there is some process with two different initial values. By the (3,2)-Intersection property, either $S_j \cap S_f \neq \emptyset$ or $S_j \cap S_c \neq \emptyset$. If $S_j \cap S_f \neq \emptyset$, then $p_j.A_p(r)[\ell] = M'_j[\ell] = p_\ell.a = v_f$ (Lemma 5.4, Lemma 5.3, and the algorithm). We then have by the algorithm that $p_j.A_p(t+1)[\ell] = p_j.d[\ell] = p_\ell.a = v_f$.

If $S_j \cap S_c \neq \emptyset$, then $p_j.d = p_c.d$ by Lemma 5.13, for every $p_c \in S_c$. By Lemma 5.16, $p_i.d \subseteq p_c.d$ for every non-faulty p_c . That is, we have that $p_c.d[\ell] = p_i.d[\ell] = v_f$. We then have that $p_j.d[\ell] = p_c.d[\ell] = p_\ell.a = v_f$. This concludes the proof of the lemma.

□

Theorem 5.20 Algorithm FFS-ROC satisfies Termination.

Proof:

This is straightforward from the algorithm: every process that does not crash in an execution of FFS-ROC decides at round $t+1$.

□

Theorem 5.21 Algorithm FFS-ROC satisfies Agreement.

Proof:

By Lemma 5.16, if a process p_i decides in ϕ , then $p_i.d \subseteq$

$p_c.d$ for every non-faulty p_c . This implies that for every p_ℓ such that $p_i.d[\ell] \neq \perp$, we have that $p_c.d[\ell] = p_i.d[\ell]$ for every non-faulty p_c .

□

Theorem 5.22 *Algorithm FFS-ROC satisfies RO Uniformity.*

Proof:

By Lemma 5.5, every correct process decides in ϕ . By the algorithm, for every non-faulty process p_c , $p_c.d[c] = p_c.a$. Thus, there must be at least one non- \mathcal{N} decision value. By Lemma 5.15, there cannot be three processes in an execution of FFS-ROC such that each process decides upon a different value. This shows the first statement of the property: $1 \leq |\text{vals}| \leq 2$, where $\text{vals} = \{d : \exists p_i \in \Pi \text{ s.t. } (p_i.d = d) \setminus \mathcal{N} \text{ in any execution of FFS-ROC. The second statement follows directly from Lemma 5.17. The third statement follows directly from Lemma 5.18.}$

□

Theorem 5.23 *Algorithm FFS-ROC satisfies Validity.*

Proof:

If $p_i \in \text{Live}(t + 1)$ in some execution ϕ of FFS-ROC, then by Lemmas 5.5 and 5.16 $p_i.d \subseteq p_c.d$, for every $p_c \in \text{Correct}(\phi)$. By the algorithm, $p_i.d[i]$ must be equal to $p_i.a$. We consequently have that $p_c.d[i]$ must be equal to $p_i.a$. This proves the first statement in the specification of Validity.

If p_i crashes in an execution ϕ of FFS-ROC, then by Lemmas 5.2 and 5.11 either $p_c.d[i] = \perp$ or $p_c.d[i] = p_i.a$, for every $p_c \in \text{Correct}(\phi)$. This shows the second statement in the definition of validity. The third statement follows directly from Lemma 5.19.

□

With Theorems 5.20, 5.21, 5.22, and 5.23, we show that FFS-ROC implements the four RO Consensus properties, thereby showing Proposition 5.1.

6 Correctness of FFS-WLE

Algorithm FFS-WLE proceeds in iterations of an infinite repeat loop. In each iteration, processes execute two phases, and in each phase a process participates in the execution of an algorithm that implements RO Consensus. For the following description, we assume that such an algorithm is FFS-ROC. As shown in Figure 2, a process that does not crash in an execution of FFS-WLE executes infinitely many iterations of the repeat loop. According

to our system model, we split an execution of an algorithm into rounds. We further number the iterations of an execution of FFS-WLE and assume that round numbers map to iteration numbers. That is, there is a mapping *Iteration* : $\mathcal{R} \rightarrow \mathcal{I}$, where \mathcal{R} is the set of round numbers as before, and $\mathcal{I} = \mathbb{Z}^*$ is the set of iteration numbers. In addition, we assume that iteration numbers increase monotonically with round numbers, and the number of rounds executed in an iteration is fixed, being a function of the number of rounds in an execution of FFS-ROC.⁴ For the purpose of the proofs that follow, we only need to assume that each phase executes at least two rounds. Note that FFS-ROC requires $t + 1$ rounds, and $t + 1 \geq 2$ if $t \geq 1$. We therefore have that each phase must have at least two rounds, assuming systems in which processes can fail. In fact, because processes can fail by crashing and we assume cores and survivor sets to characterize valid sets of faulty processes, we can use the same argument as in [3] to show that $t + 1$ is a lower bound on the number of rounds.

According to the discussion in the previous paragraph, we associate an iteration number with each iteration of the algorithm in an execution. In the following, we use iteration numbers to refer to iterations of the repeat loop. In proving the correctness of FFS-WLE, we also use the following definitions:

- vals_i , $i \in \{1, 2\}$ is the set $\{d : p_i.d = d \wedge p_i \in \Pi\} \setminus \mathcal{N}$ after executing FFS-ROC at phase i of some iteration ζ , $\zeta \in \mathcal{I}$;
- a process p_i finishes a phase $\rho \in \{1, 2\}$ of some iteration ζ in an execution ϕ if it neither stops nor crashes before executing the last step of that phase;
- A process p_i starts phase $\rho \in \{1, 2\}$ of an iteration ζ at time τ if p_i executes at least one step of phase x of ζ and the first step s of p_i at phase x of ζ is such that $\mathcal{T}(s) = \tau$.

As in Section 4, we assume that FFS-WLE uses a system profile (Π, C_Π, S_Π) and that this profile satisfies (3,2)-Intersection.

Now we prove the following proposition.

Proposition 6.1 FFS-WLE implements Safety, LE-Liveness, and FF-Stability.

⁴Because there are infinitely many executions of FFS-ROC in an execution of FFS-WLE and round numbers monotonically increase with time, the round numbers in the pseudocode for FFS-ROC are relative to the first round in which an execution of FFS-ROC starts.

We show proposition 6.1 with the following set of theorems, each one proving a property of Weak Leader Election.

Theorem 6.2 *Algorithm FFS-WLE satisfies Safety.*

Proof:

Let $\phi = \langle F, I, S, T, \mathcal{T} \rangle$ be an execution of FFS-WLE. We have to show that $|\{p_i \in \Pi : p_i.elected\}| < 2$ for every $\tau \in T$. First, we show that in an iteration of ϕ , at most one process is elected. By the RO Uniformity property of RO Consensus, there is at least one decision value and at most two different decision values as a result of phase 1. That is, $1 \leq |vals_1| \leq 2$. Suppose $|vals_1| = 1$. By the algorithm, every process p_i that finishes phase 2 uses a list x in its decision value $p_i.d$, where x has the minimum number of non- \perp values among all lists in $p_i.d$. x must be the initial value of some processes by Validity. By assumption, there is a single initial value in phase 2, and consequently $p_i.d[j] \in \{x, \perp\}$, for every $p_j \in \Pi$, implying that no two distinct processes that finish phase 2 can be elected.

Now suppose that $|vals_1| = 2$. From RO Uniformity, we have that there are values $d_1, d_2 \in vals_1$ and $S_1, S_2 \in \mathcal{S}_\Pi$ such that:

$$\begin{aligned} \wedge \forall p \in S_1 : & \quad \vee \quad p \text{ crashes} \\ & \quad \vee \quad p.d = d_1 \\ \wedge \forall p \in S_2 : & \quad \wedge \quad p.d = d_2 \\ & \quad \wedge \quad p \text{ is not faulty} \end{aligned}$$

By the algorithm, a process that finishes phase 1 of an iteration executes FFS-ROC once more in phase 2 with its decision value of the previous phase as its initial value. If the above properties hold, then the only processes in S_1 that do not have d_1 as initial value are the ones that crash before phase 2 starts. Let's call this set $Crash_1$. By Validity, if some process p_i decides upon a value $p_i.d$ such that $d_1 \in p_i.d$, then every process p_j that finishes phase 2 is such that $d_1 \in p_j.d$. We then again have that there is a single process that can be elected because every process that finishes phase 2 has d_1 in its decision value and $d_1 \subseteq d_2$ by Agreement.

It remains to show that if p_i is elected in iteration ζ , and p_j is elected in iteration ζ' , and $\zeta > \zeta'$, then there is no $\tau \in T$ such that both $p_i.elected$ and $p_j.elected$ are true at time τ . By the algorithm, every process that starts the execution of phase 1 in an iteration, first sets its flag *elected* to false. If the iteration is the first of ϕ , then p_j

cannot be elected in a previous iteration, and the hypothesis is vacuously true. Now suppose an iteration $\zeta > 0$. By assumption, every process that executes a phase of an iteration ζ executes at least two rounds. By P-Liveness, no process can start a new round $r + 1$, $r \geq 0$, without having every other live process executing at least one step of r . If a process p_i starts phase 2 of an iteration at time τ , then every process that has not crashed by τ must have executed at least one step of round zero of FFS-ROC at phase 1 of ζ . Otherwise, there is a non-crashed process p_j such that p_i executes the first step of a round $r + 1$ of FFS-ROC whereas p_j has not executed any steps of r . This implies that no process can finish phase 2 of an iteration without having all non-crashed processes setting *elected* to false.

This concludes the proof of the theorem.

□

Theorem 6.3 *Algorithm FFS-WLE satisfies LE-Liveness.*

Proof:

We have to show that for every execution $\phi = \langle F, I, S, T, \mathcal{T} \rangle$ of FFS-WLE and for every $\tau \in T$, there is some iteration after τ such that $|\{p_i \in \Pi : p_i.elected\}| > 1$.

Proof by contradiction. Suppose an execution $\phi = \langle F, I, S, T, \mathcal{T} \rangle$ of FFS-WLE and a time $\tau \in T$ such that $p_i.elected$ is false forever after τ for every p_i . By Validity and RO Uniformity, in every iteration ζ of ϕ , $\zeta \in \mathcal{I}$, there is a value $v \in vals_1$ such that v is the list with the least number of non- \perp values, and for every process p_i that finishes phase 2 of iteration ζ , $v \in p_i.d$. Every process that finishes phase 2 of iteration ζ selects the same value i as the first index of v mapping to a value in v with a non- \perp value. If process p_i evaluates the last “if” statement of phase 2, then it sets $p_i.elected$ to true. If p_i crashes, however, then it does not set $p_i.elected$ to true, and no process is elected in iteration ζ . By the assumption that all failures are benign, crashing (or stopping which is equivalent to crashing in our model) is the only possibility for having no process elected in an iteration ζ . By the assumption that t ($|\Pi|$ subtracted the size of the smallest survivor set) is the largest number of processes that can crash in ϕ , there can be at most t iterations after τ such that $|\{p_i \in \Pi : p_i.elected\}| = 0$.

□

Theorem 6.4 *Algorithm FFS-WLE satisfies FF-Stability.*

Proof:

Suppose ϕ is a failure-free execution and ζ is an iteration of ϕ . By Agreement, every process decides upon the same value in both phases of iteration ζ . We then have that every process p_i uses the same value x to determine whether it sets $p_i.elected$ to true in ϕ . Moreover, we have that $x[i] = p_i$ for every i , by Validity. Assuming that $\Pi = \{p_1, p_2, \dots, p_n\}$, we have by the algorithm that p_1 sets $p_1.elected$ to true at phase 2 of ζ .

□

7 Adding E-Stability

FFS-WLE allows for executions in which two or more processes are elected infinitely often. Such behavior, however, is not desirable. As leadership moves to one process to another, the responsibility of accomplishing the tasks of a leader also moves. Recall that the original motivation is to embed such a Leader Election algorithm into a Primary-Backup protocol. This oscillation causes unnecessary overhead such as requests being forwarded to the correct Primary or even system instabilities if changes occur too frequently.

We now show how to modify FFS-WLE (and FFS-ROC) to also satisfy E-Stability. We call WLE the modified version of FFS-WLE, and ROC the modified version of FFS-ROC to distinguish between the previous versions and the modified versions.

First, instead of initializing $p_i.s(0)$ to Π , as in FFS-ROC, $p_i.s(0)$ is initialized to a parameter $p_i.Procs$. We also roll forward the value of $p_i.s(t+1)$ in FFS-WLE instead of having $p_i.s(0)$ constant as in FFS-ROC. That is, in an iteration $\zeta > 0$, $p_i.s(0)$ in Phase 1 is $p_i.s(t+1)$ in Phase 2 of iteration $\zeta - 1$. If $\zeta = 0$, then $p_i.s(0)$ in Phase 1 is Π . For the initial value of $p_i.s(0)$ in Phase 2, we use $p_i.s(t+1)$ of Phase 1 of the same iteration. For clarity, we repeat the pseudocode for these algorithms with the respective modifications in Figures 4 and 5. Note that the main modifications in ROC are: 1) ROC has two parameters instead of one; 2) in round 1, p_i checks whether $p_i.s(1) \subseteq p_i.s(0)$. WLE is different from FFS-WLE by initializing $p_i.Procs$ to Π and by rolling $p_i.s(t+1)$ forward.

It is straightforward to see that the proof of Section 5 is valid for ROC if the following constraint holds for every execution $\phi = \langle F, I, S, T, \mathcal{T} \rangle$ of ROC: if $p_c \in Correct(\phi)$ and $p_i \in \Pi$ is a process in $p_c.s(1)$, then $p_i \in p_c.Procs$. That is, $p_c.Procs$ must contain all the processes that send messages to p_c at round zero if p_c is not faulty. Other-

Algorithm ROC on input $p_i.a, p_i.Procs$

round 0:

```

 $p_i.s(0) \leftarrow p_i.Procs; p_i.sr(0) \leftarrow p_i.s(0)$ 
 $p_i.A[i] \leftarrow p_i.a$ 
for all  $p_k \in \Pi, p_k \neq p_i : p_i.A[i] \leftarrow \perp$ 
 $p_i.A_p(0) \leftarrow p_i.A$ 
send  $p_i.A$  to all

```

round 1:

```

 $p_i.sr(1) \leftarrow p_i.s(1)$ 
if  $\forall p_i.s(1) \not\subseteq p_i.s(0)$ 
   $\forall \#S \in \mathcal{S}_\Pi : S \subseteq p_i.sr(1)$ 
  then decide  $[\perp, \dots, \perp]$ 
else for each message  $m_j \in p_i.M(1), p_k \in \Pi$ :
  if  $(p_i.A[k] = \perp) p_i.A[k] \leftarrow m_j.A[k]$ 
 $p_i.A_p(1) \leftarrow p_i.A$ 
send  $p_i.A$  to all

```

round $r: 2 \leq r \leq t$:

```

 $p_i.sr(r) \leftarrow p_i.s(r) \setminus \{p_j : \exists m \in p_i.M(r) : p_i.A_p(r-2) \not\subseteq m.A \wedge m.from = p_j\}$ 
if  $\forall p_i.s(r) \not\subseteq p_i.s(r-1)$ 
   $\forall \#S \in \mathcal{S}_\Pi : S \subseteq p_i.sr(r)$ 
  then decide  $[\perp, \dots, \perp]$ 
else for each message  $m \in p_i.M(r), p_k \in \Pi$ :
  if  $(p_i.A[k] = \perp) p_i.A[k] \leftarrow m.A[k]$ 
 $p_i.A_p(r) \leftarrow p_i.A$ 
send  $p_i.A$  to all

```

round $t+1$:

```

 $p_i.sr(t+1) \leftarrow p_i.s(t+1) \setminus \{p_j : \exists m \in p_i.M(t+1) : p_i.A_p(t-1) \not\subseteq m.A \wedge m.from = p_j\}$ 
if  $\forall p_i.s(t+1) \not\subseteq p_i.s(t)$ 
   $\forall \#S \in \mathcal{S}_\Pi : S \subseteq p_i.sr(t+1)$ 
  then decide  $[\perp, \dots, \perp]$ 
else for each message  $m \in p_i.M(t+1), p_k \in \Pi$ :
  if  $(p_i.A[k] = \perp) p_i.A[k] \leftarrow m.A[k]$ 
 $p_i.A_p(t+1) \leftarrow p_i.A$ 
decide  $p_i.A$ 

```

Figure 4: Algorithm run by process p_i .

Algorithm WLE

```

P ← Π
repeat {
  pi.elected ← FALSE
  Phase 1:
    Run ROC with
      pi.a ← i; pi.Procs ← P.
    P ← pi.s(t + 1)
    if (pi.d = [⊥, . . . , ⊥]) then stop
  Phase 2:
    Run ROC with
      pi.a ← pi.d from Phase 1; pi.Procs ← P.
    P ← pi.s(t + 1)
    if (pi.d = [⊥, . . . , ⊥]) then stop
    let x be a value of pi.d [1 . . . n]
      such that pi.d[x] ≠ [⊥, . . . , ⊥]
      and it has the least number of non-⊥ values
    if (pi is the first index of x such that x[i] ≠ ⊥)
      then pi.elected ← TRUE
}

```

Figure 5: Algorithm run by process p_i .

wise, p_c can falsely detect that it is faulty, and stop, violating Validity. Lemma 5.5 states that correct processes do not stop, and hence the proof changes with the modification to the algorithm. The change is small, however. It consists in modifying the base case of the inductive argument to show that no correct process p_c stops in round 1 if $p_c.s(0)$ contains all the processes that send messages to p_c in round zero (0).

If p_f is faulty, then there are no restrictions on the input $p_f.Procs$. Intuitively, a faulty process $p_f.Procs$ can receive any subset of processes sent to it. Consequently, it is not possible to impose a similar constraint as we did for correct processes. Different from correct processes, if a faulty process stops, it does not violate any of the RO Consensus properties.

According to the modifications described previously, $p_i.Procs$ is the set of processes from which p_i receives a message in the last round of the previous execution of ROC (Π if it is the execution of ROC in Phase 1 of iteration 0). By assumption and by the algorithm, once a process crashes (or stops) it never sends messages again in an execution of WLE. Thus, if p_c is a correct process, then $p_c.Procs$ must contain all the processes that p_c receives messages from in an execution of ROC, satisfying our constraint on $p_c.Procs$ for correct processes.

Because the proofs of Theorems 6.2 and 6.3 rely solely on the properties of RO Consensus, we also have that these proofs hold for WLE. It remains to show that WLE satisfies E-Stability. First, we present a few definitions to be used in the proof of E-Stability. By Theorem 6.2, in every iteration ζ of an execution ϕ of WLE it is the case that either one process p_i is such that $p_i.elected$ evaluates to true at the end of ζ or no process p_i is such that $p_i.elected$ evaluates to true at the end of ζ . We then use $Leader(\zeta, \phi)$ to denote the process p_i , $p_i.elected$ evaluates to true at the end of iteration ζ of ϕ , or \perp if no such process exists. Finally, we need some terminology to refer to values that processes decide across iterations and in the two different phases of an iteration. We then use $\mathcal{D}_\zeta^\rho(i)$ to denote $p_i.d$ at the end of phase $\rho \in \{1, 2\}$ of iteration ζ .

Proposition 7.1 WLE implements Weak Leader Election.

With Theorems 6.2, 6.3, and 6.4, we showed that WLE satisfies Stability, LE-Liveness, and FF-Stability. It remains to show E-Stability. Before we show our main result of this section, we state and prove a few preliminary lemmas.

Lemma 7.2 Let ϕ be an execution of WLE. For every iteration ζ of WLE, if p_i finishes phase 1 of both ζ and $\zeta + 1$, then $\mathcal{D}_{\zeta+1}^1(i) \subseteq \mathcal{D}_\zeta^1(i)$.

Proof:

Proof by contradiction. Suppose that there is an iteration ζ such that the assertion $\mathcal{D}_{\zeta+1}^1(i) \subseteq \mathcal{D}_\zeta^1(i)$ is false. This implies that there is some process p_ℓ , $\ell \neq i$, for which $\mathcal{D}_{\zeta+1}^1(i)[\ell] \neq \perp$ and $\mathcal{D}_\zeta^1(i)[\ell] = \perp$. By Lemma 5.10, in the execution of ROC in phase 1 of iteration $\zeta + 1$, there is a chain ω_ℓ such that $\omega_\ell[r] = i$, $r > 0$. By assumption, for every process p_j such that $\omega_\ell[\rho] = j$, $\rho \in \{z \in \mathcal{R} : 1 \leq z \leq r\}$, p_j must be in $p_j.Procs$, where $\omega_\ell[\rho-1] = j'$, otherwise the process with identifier $\omega_\ell[\rho]$, $\rho \leq r$, stops at round ρ .

Now suppose that $\omega_\ell[1] = j$. By the algorithm, $p_j.sr(r)$ contains some survivor set S_j , for every round $r \in \{z \in \mathcal{R} : z \leq t + 1\}$ of the execution of ROC in iteration ζ . Also, there is such a survivor set S_i for p_i , and there is some survivor set S_c such that $S_c \subseteq Correct(\phi)$. By (3,2)-Intersection, S_i either intersects S_j or intersects S_c . If S_i intersects S_j , then by Lemma 5.13, $\mathcal{D}_\zeta^1(i)[\ell] \neq \perp$, contradicting our initial assumption. If S_i intersects S_c , then by Lemma 5.14 $\mathcal{D}_\zeta^2(c) \subseteq \mathcal{D}_\zeta^2(i)$, for some non-faulty

$p_c \in \text{Correct}(\phi)$. By Agreement, $\mathcal{D}_\zeta^2(c)[\ell] \neq \perp$, and consequently $\mathcal{D}_\zeta^1(i)[\ell] \neq \perp$, again contradicting our initial assumption.

This completes the proof of the lemma.

□

Lemma 7.3 *Let ϕ be an execution of WLE and ζ be an iteration of ϕ . No process is elected in ζ only if some process crashes or stops in ζ .*

Proof:

By RO Uniformity, we have that $1 \leq \text{vals}_i \leq 2$, $i \in \{1, 2\}$. By Validity and RO Uniformity, there is some value d in vals_1 such that $d \in p_i.d$ for every process p_i that finishes phase 2 of iteration ζ , and d contains the least number of non- \perp values. Because $d \neq \mathcal{N}$, there must be a process p_e such that e is the smallest index of d such that $d[e] \neq \perp$. We then have that p_e sets $p_e.\text{elected}$ to true unless p_e crashes. Thus, if $\text{Leader}(\zeta, \phi) = \perp$, then p_e must crash or stop in ζ .

□

Lemma 7.4 *Let ϕ be an execution of WLE, ζ and ζ' be iterations of ϕ , $\zeta + 1 < \zeta'$, such that $\text{Leader}(\zeta, \phi) = \text{Leader}(\zeta', \phi) = p_e$. If $\text{Leader}(\zeta + 1, \phi) = p_{e'}$, $e \neq e'$, then there is a process p_i that crashes or stops in some iteration ζ'' , $\zeta \leq \zeta'' \leq \zeta'$.*

Proof:

If p_e is elected in iteration ζ of ϕ , then there is a value d in vals_1 of ζ such that e is the smallest index of d with a non- \perp value, and $d \in \mathcal{D}_\zeta^2(e)$. Now if $p_{e'}$ is elected in iteration $\zeta + 1$, then there is a value d' in vals_1 of $\zeta + 1$ such that e' is the smallest index of d' with a non- \perp value, and $d' \in \mathcal{D}_{\zeta+1}^2(e')$. By assumption, p_e is elected again in iteration ζ' . As before, there must be a value d'' in vals_1 of ζ' such that e is the smallest index of d'' with a non- \perp value.

There are two possibilities regarding the identifiers e and e' : 1) $e' < e$; 2) $e < e'$. If $e' < e$, then there must be a second value d_c in vals_1 of ζ such that $d \subseteq d_c$ (by assumption, p_e has not crashed by iteration $\zeta' > \zeta + 1$; by Validity, every non-faulty process p_c is such that $p_e.a \in p_c.d$). Let p_i be a process such that $\mathcal{D}_\zeta^1(i) = d$. Suppose by way of contradiction that p_i finishes phase 2 of $\zeta + 1$. By Lemma 7.2, $\mathcal{D}_{\zeta+1}^1(i) \subseteq \mathcal{D}_\zeta^1(i)$, and $\mathcal{D}_{\zeta+1}^1(i)[e'] = \perp$. By Validity, there is some $p_c \in \text{Correct}(\phi)$ such that $\mathcal{D}_{\zeta+1}^1(c)[e'] \neq \perp$. Note that $p_{e'}$ does not crash or stop in iteration $\zeta + 1$ or in a previous iteration. By RO Uniformity, $\mathcal{D}_{\zeta+1}^1(i) \subseteq \mathcal{D}_{\zeta+1}^1(c) = d'$. By RO Uniformity and Validity, $\mathcal{D}_{\zeta+1}^1(i)$ must be the value used by every process that

completes phase 2 of iteration $\zeta + 1$ to determine whether it elects itself or not. Since e' is not the smallest index of $\mathcal{D}_{\zeta+1}^1(i)$ that evaluates to a non- \perp value, $p_{e'}$ is not elected in $\zeta + 1$. This contradicts our initial assumption. We hence have that p_i must crash or stop by iteration $\zeta + 1$.

Now if $e < e'$, then $d'[e] = \perp$ by assumption (e' is the smallest index in d' with a non- \perp value). We use a similar argument as in the first case. Suppose by way of contradiction that $p_{e'}$ finishes phase 2 of ζ' . By Lemma 7.2, $\mathcal{D}_{\zeta'}^1(e') \subseteq \mathcal{D}_{\zeta+1}^1(e')$, and $\mathcal{D}_{\zeta'}^1(e')[e] = \perp$. By Validity, there is some $p_c \in \text{Correct}(\phi)$ such that $\mathcal{D}_{\zeta'}^1(c)[e] \neq \perp$. Note that p_e does not crash or stop in iteration ζ' or in a previous iteration. By RO Uniformity, $\mathcal{D}_{\zeta'}^1(e') \subseteq \mathcal{D}_{\zeta'}^1(c) = d''$. By RO Uniformity and Validity, $\mathcal{D}_{\zeta'}^1(e')$ must be the value used by every process that completes phase 2 of iteration ζ' to determine whether it elects itself or not. Since e is not the smallest index of $\mathcal{D}_{\zeta'}^1(i)$ that evaluates to a non- \perp value, p_e is not elected in ζ' . This contradicts our initial assumption. We conclude that p_e is not elected in ζ' unless $p_{e'}$ crashes or stops by iteration ζ' .

Finally, we have that either $p_{e'}$ crashes or stops by iteration ζ' or some faulty process p_i crashes or stops by iteration $\zeta + 1$. This concludes the proof of the lemma.

□

Theorem 7.5 *WLE satisfies E-Stability.*

Proof:

Let ϕ be an execution of WLE. By LE-Liveness, infinitely often some process is elected in ϕ . By Lemma 7.3, an iteration ζ has no leader elected only if some process crashes in ζ , and by assumption there is a finite number of processes that crash or stop. Thus, there is a bounded number of iterations that have no leader elected.

Let t be a time such that every iteration that starts after t has a leader elected. Such a t exists by the previous argument. We then have that every iteration that starts after t has a leader elected, and it remains to show that there is some $t' \geq t$ and some process p_e such that for every iteration ζ that starts after t' , p_e is elected in both ζ and $\zeta + 1$. Suppose by way of contradiction that there is no such t' in ϕ . Let p_e be a process that is elected infinitely often after t' . Such a process must exist because the set of processes is finite. By assumption, there is an infinite sequence of iterations $\zeta_1 < \zeta_2 < \zeta_3 \dots$, which are not necessarily consecutive, such that p_e is elected in ζ_i but not in $\zeta_i + 1$. By Lemma 7.4, for every $i \in \mathbb{Z}^+$, there is an iteration ζ , $\zeta_i \leq \zeta \leq \zeta_{i+1}$, such that some process crashes or stops in ζ . By assumption, the number of processes crashing or stopping is bounded. Consequently, there cannot be such an

infinite sequence. We conclude that there must be some $t' \geq t$ and some process p_e such that for every iteration ζ that starts after t' , p_e is elected in both ζ and $\zeta + 1$.

□

8 Discussion

Developing a Primary-Backup protocol that uses WLE is future work. We can, however, make a few observations regarding the use of an algorithm as WLE to elect primaries in a Primary-Backup protocol. As mentioned previously, WLE enables faulty processes to be elected. In a Primary-Backup system, this feature impacts on liveness, although not on correctness. Often, there is a time bound in the replies to client requests, and it is impossible to meet such bounds if the primary can be faulty. An immediate consequence of electing faulty processes is that service time is not bounded during the period of time a faulty process remains as the primary. As discussed before, processes that commit failures (but do not stop or crash) are detected. In practice, we rely on an off-line mechanism to detect these anomalies and take the appropriate measures that can be, for example, to remove faulty processes from the system.

It is possible, however, that faulty processes go through an iteration of WLE undetected as such, and fail to reply to client requests due to receive-omission failures. To solve this problem, we can require clients to broadcast requests to all the replicas and the primary to broadcast replies to all the backup replicas as well. Correct processes are also capable of detecting failures in such cases, although they may not be able to "warn" the faulty primary that it is actually faulty. Recall that failure detection for omission failures requires twofold replication.

Finally, the iterations of the repeat loop of WLE are consecutive without any delay in between for expositional purposes. In practice, iterations should be delayed until failures are detected, they are manually triggered, or if none of these are desirable or possible, some time threshold is reached.

9 Conclusions

We described in this paper a weaker version of the Leader Election problem and an algorithm that solves this problem. This version of the problem, unlike the traditional definition of Leader Election, enables faulty processes to

be elected. The main advantage of enabling it is requiring a lower degree of replication.

There are other interesting features of the WLE algorithm. First, it uses cores and survivor sets instead of a threshold. This enables more flexible characterizations of systems with an heterogeneous set of processes. Second, it uses an unusual type of Intersection property, *i.e.*, (3,2)-Intersection. This property generalizes a degree of replication of the form $n > \lfloor 3t/2 \rfloor$, where t is the threshold on the number of failures in any execution. Finally, correct processes are able to detect faulty processes. By Lemma 5.16, non-crashed faulty processes decide upon lists with fewer values, and one can build an alarm system by collecting decision values by the end of every iteration.

Although we have not thoroughly investigated using WLE to implement Primary-Backup protocols, we believe our algorithm provides practical benefits compared to previous solutions.

References

- [1] L. Lamport, "The Part-Time Parliament," *ACM Transactions on Computer Systems*, vol. 16, pp. 133–169, May 1998.
- [2] N. Budhiraja, K. Marzullo, F. Schneider, and S. Toueg, "Optimal primary-backup protocols," in *6th International Workshop on Distributed Algorithms (WDAG)*, pp. 362–378, Nov 1992.
- [3] F. Junqueira and K. Marzullo, "Lower Bound on the Number of Rounds for Synchronous Consensus with Dependent Process Failures," Tech. Rep. CS2003-0734, UCSD, 2001.

A Specification of ROC in TLA+

ROConsensus

Naturals, FiniteSets, TLC

P ,	Set of processes
$Vals$,	Domain of d values
$NULL$,	A null value
$SurvivorSet$,	All survivor sets
I	Initial values
A ,	Accumulated values of a process
$Aprime$,	Accumulated values of a process (previous round)
d ,	Decision value of a process
$messages$,	Sent messages
$crashed$,	Crashed processes
$faulty$,	R-O faulty processes
$round$,	Current round
$pRound$,	Round that p is in
$MyTurn$,	
$recdFrom$	$recdFrom[p][r]$ is the set of processes p received a message from in round r .

$N \triangleq Cardinality(P)$

$P \subseteq Nat$

SurvivorSet satisfies (3, 2)-Intersection

$$\begin{aligned} & \wedge \forall s \in SurvivorSet : s \subseteq P \\ & \wedge \forall s1, s2, s3 \in SurvivorSet : \\ & \quad \vee s1 \cap s2 \neq \{\} \\ & \quad \vee s2 \cap s3 \neq \{\} \\ & \quad \vee s1 \cap s3 \neq \{\} \\ & \wedge \forall p \in P : \exists s \in SurvivorSet : p \in s \end{aligned}$$

$I \in [P \rightarrow Vals]$

Used for number of rounds

$$\begin{aligned} t \triangleq \quad x \triangleq \quad & s \in SurvivorSet : \\ & \forall s2 \in SurvivorSet : \\ & \quad Cardinality(s) \leq Cardinality(s2) \\ N - Cardinality(x) \end{aligned}$$

$$\begin{aligned} ROETypeOK \triangleq & \wedge A \in [P \rightarrow [P \rightarrow Vals \cup \{NULL\}]] \\ & \wedge Aprime \in [P \rightarrow [P \rightarrow Vals \cup \{NULL\}]] \\ & \wedge d \in [P \rightarrow [P \rightarrow Vals \cup \{NULL\}] \cup \{NULL\}] \\ & \wedge messages \subseteq [\\ & \quad from : P, \\ & \quad to : P, \\ & \quad round : 0 .. t, \end{aligned}$$

$$\begin{aligned}
& val : [P \rightarrow Vals \cup \{NULL\}] \\
& \wedge faulty \subseteq P \\
& \wedge crashed \subseteq P \\
& \wedge round \in 0 .. (t + 1) \\
& \wedge pRound \in [P \rightarrow 0 .. (t + 1)] \\
& \wedge Cardinality(faulty) + Cardinality(crashed) \leq t \\
& \wedge faulty \cap crashed = \{\} \\
& \wedge MyTurn \in P
\end{aligned}$$

$$\begin{aligned}
& \wedge A \in [P \rightarrow [P \rightarrow Vals \cup \{NULL\}]] \\
& \wedge \forall p, q \in P: \quad p = q \quad A[p][q] = I[p] \\
& \quad \quad \quad \quad \quad \quad A[p][q] = NULL
\end{aligned}$$

$$\begin{aligned}
ROEInit \triangleq & \wedge A = [p \in P \mapsto [q \in P \mapsto \quad p = q \quad I[p] \quad NULL]] \\
& \wedge Aprime = A \\
& \wedge d = [p \in P \mapsto NULL] \\
& \wedge messages = \{\} \\
& \wedge faulty = \{\} \\
& \wedge crashed = \{\} \\
& \wedge round = 0 \\
& \wedge pRound = [p \in P \mapsto 0] \\
& \wedge recdFrom = [p \in P \mapsto [r \in 0 .. (t + 1) \mapsto \\
& \quad \quad \quad r = 0 \quad P \quad \{\}]] \\
& \wedge MyTurn = \quad p \in P : 1 = 1
\end{aligned}$$

Send message v in round r

$$\begin{aligned}
Send(from, to, r, v) \triangleq & \\
& \wedge from \notin crashed \\
& \wedge messages' = messages \cup \\
& \quad \{[from \mapsto from, to \mapsto to, round \mapsto r, val \mapsto v]\}
\end{aligned}$$

Process p has decided

$$Decided(p) \triangleq d[p] \neq NULL$$

Consensus has terminated

$$Terminated \triangleq \forall p \in P \setminus crashed : Decided(p)$$

Sources of messages

$$From(msgs) \triangleq \{p \in P : \exists m \in msgs : m.from = p\}$$

$$\begin{aligned}
LastStep(p) \triangleq & \wedge p \notin crashed \\
& \wedge round = t + 1 \\
& \wedge pRound[p] = t + 1 \\
& \wedge \neg Decided(p) \\
& \wedge d' = [d \quad ![p] = A[p]] \\
& \wedge \quad \langle A, Aprime, messages, crashed, \\
& \quad \quad faulty, round, pRound, recdFrom, \\
& \quad \quad MyTurn \rangle
\end{aligned}$$

$$\begin{aligned}
Fail(p) \triangleq & \wedge \exists s \in SurvivorSet : s \subseteq (P \setminus (faulty \cup crashed \cup \{p\})) \\
& \wedge \forall p \notin crashed \quad \cup faulty
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{faulty}' = \text{faulty} \cup \{p\} \\
& \wedge \text{crashed} \\
\vee & \wedge p \notin \text{crashed} \\
& \wedge \text{crashed}' = \text{crashed} \cup \{p\} \\
& \wedge \text{faulty}' = \text{faulty} \setminus \{p\} \\
\wedge & \langle A, A_{\text{prime}}, d, \text{messages}, \text{round}, p_{\text{Round}}, \\
& \text{recdFrom}, \text{MyTurn} \rangle
\end{aligned}$$

Each round (except round 0) starts by receiving messages sent in the previous round and terminates (except round $t + 1$) by sending $A[p]$ to all processes.

$$\begin{aligned}
\text{RoundDone}(r) \triangleq & \wedge \text{round} = r \\
& \wedge \forall p \in P \setminus \text{crashed} : \\
& \quad \vee \text{Decided}(p) \\
& \quad \vee \forall q \in P : \\
& \quad \quad \exists m \in \text{messages} : \wedge m.\text{from} = p \\
& \quad \quad \quad \wedge m.\text{to} = q \\
& \quad \quad \quad \wedge m.\text{round} = r \\
& \wedge \text{round}' = \text{round} + 1 \\
& \wedge \langle A, A_{\text{prime}}, d, \text{messages}, \text{crashed}, \text{faulty}, \\
& \quad p_{\text{Round}}, \text{recdFrom}, \text{MyTurn} \rangle
\end{aligned}$$

$$\begin{aligned}
\text{RecvRound1}(p) \triangleq & \text{msgsSentToMe} \triangleq \{m \in \text{messages} : (m.\text{to} = p \wedge m.\text{round} = 0)\} \\
& \text{Am faulty if didn't receive messages from a survivor set} \\
& \text{IMustBeFaulty}(mset) \triangleq \\
& \quad \vee \neg(\text{From}(mset) \subseteq \text{recdFrom}[p][0]) \\
& \quad \vee \neg(\exists s \in \text{SurvivorSet} : s \subseteq \text{From}(mset))
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{MyTurn} = p \\
& \wedge \text{round} = 1 \\
& \wedge p \notin \text{crashed} \\
& \wedge p_{\text{Round}}[p] = 0 \\
& \wedge \neg \text{Decided}(p) \\
& \wedge p_{\text{Round}}' = [p_{\text{Round}} \quad ![p] = 1] \\
& \wedge \exists \text{msgsRecd} \in \text{msgsSentToMe} : \\
& \quad \wedge (p \in \text{faulty} \vee \text{msgsRecd} = \text{msgsSentToMe}) \\
& \quad \wedge p \in \text{From}(\text{msgsRecd}) \\
& \quad \wedge \text{recdFrom}' = [\text{recdFrom} \quad ![p][1] = \text{From}(\text{msgsRecd})] \\
& \quad \wedge \text{IMustBeFaulty}(\text{msgsRecd}) \\
& \quad \quad \wedge d' = [d \quad ![p] = [q \in P \mapsto \text{NULL}]] \\
& \quad \quad \quad \wedge A \\
& \quad \quad \wedge A' = [A \quad ![p] = [q \in P \mapsto \\
& \quad \quad \quad A[p][q] \neq \text{NULL} \rightarrow A[p][q] \\
& \quad \quad \quad \square \wedge (A[p][q] = \text{NULL}) \\
& \quad \quad \quad \quad \wedge (\forall m \in \text{msgsRecd} : m.\text{val}[q] = \text{NULL}) \\
& \quad \quad \quad \quad \rightarrow A[p][q] \\
& \quad \quad \quad \square \wedge (A[p][q] = \text{NULL}) \\
& \quad \quad \quad \quad \wedge (\exists m \in \text{msgsRecd} : m.\text{val}[q] \neq \text{NULL}) \\
& \quad \quad \quad \quad \rightarrow v \in \text{Vals} : (\exists m \in \text{msgsRecd} : \\
& \quad \quad \quad \quad \quad m.\text{val}[q] = v)
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{round} = r \\
& \wedge \text{pRound}[p] = r \\
& \wedge \exists q \in P : \wedge q \notin \text{sentTo} \\
& \quad \wedge \text{Send}(p, q, r, A[p]) \\
& \wedge \langle A, \text{Aprime}, d, \text{crashed}, \text{faulty}, \text{round}, \text{pRound}, \\
& \quad \text{recdFrom}, \text{MyTurn} \rangle
\end{aligned}$$

$$\text{NextP}(r) \triangleq \text{sentTo} \triangleq \{q \in P : \exists m \in \text{messages} : \\
(m.\text{from} = \text{MyTurn} \wedge m.\text{to} = q \wedge m.\text{round} = r)\}$$

$$\begin{aligned}
& \wedge \vee \forall q \in P : q \in \text{sentTo} \\
& \quad \vee \text{MyTurn} \in \text{crashed} \\
& \wedge \text{MyTurn}' = \quad p \in P : p \neq \text{MyTurn} \wedge p \notin \text{crashed} \\
& \wedge \langle A, \text{Aprime}, d, \text{messages}, \text{crashed}, \text{faulty}, \text{round}, \\
& \quad \text{pRound}, \text{recdFrom} \rangle
\end{aligned}$$

$$\begin{aligned}
\text{ROENext} \triangleq & \vee \exists p \in P : \text{Fail}(p) \\
& \vee \exists p \in P : \vee \text{RecvRound1}(p) \\
& \quad \vee \text{LastStep}(p) \\
& \quad \vee \exists r \in 0 \dots t : \vee \text{RoundDone}(r) \\
& \quad \quad \vee \text{SendRound}(p, r) \\
& \quad \vee \exists r \in 2 \dots (t+1) : \text{RecvRoundR}(p, r) \\
& \quad \vee \exists r \in 0 \dots (t+1) : \text{NextP}(r)
\end{aligned}$$

$$\text{vars} \triangleq \langle A, \text{Aprime}, d, \text{messages}, \text{crashed}, \text{faulty}, \text{round}, \text{pRound}, \\
\text{recdFrom}, \text{MyTurn} \rangle$$

Since behaviors are finite ignoring stuttering, ignore liveness

$$\text{ROESpec} \triangleq \text{ROEInit} \wedge \square[\text{ROENext}]_{\text{vars}}$$

Some sets of processes

$$\text{NotCrashed} \triangleq P \setminus \text{crashed}$$

$$\text{NotFaulty} \triangleq (P \setminus \text{crashed}) \setminus \text{faulty}$$

$$\text{DecidedSomething} \triangleq \{p \in P : \text{Decided}(p)\}$$

$$\begin{aligned}
\text{DecidedSomethingInteresting} \triangleq & \{p \in \text{DecidedSomething} : \\
& \exists q \in P : d[p][q] \neq \text{NULL}\}
\end{aligned}$$

Relations on decisions

$$d\text{vSubsetEq}(d1, d2) \triangleq \forall r \in P : (d1[r] \neq \text{NULL}) \Rightarrow (d1[r] = d2[r])$$

$$d\text{SubsetEq}(p, q) \triangleq d\text{vSubsetEq}(d[p], d[q])$$

$$d\text{Subset}(p, q) \triangleq d\text{SubsetEq}(p, q) \wedge (d[p] \neq d[q])$$

RO-Consensus safety properties

$$\begin{aligned}
\text{GoodDecision} \triangleq & \wedge \forall p \in \text{NotFaulty}, q \in P : \\
& q \in \text{NotCrashed} \quad d[p][q] = I[q]
\end{aligned}$$

$$\begin{aligned}
& d[p][q] \in \{NULL, I[q]\} \\
\wedge \forall v1, v2 \in Vals : \\
& \quad \forall S1, S2 \in SurvivorSet : \\
& \quad (\forall \exists p \in S1 : I[p] \neq v1 \\
& \quad \quad \forall \exists p \in S2 : (\forall I[p] \neq v2 \\
& \quad \quad \quad \vee p \in faulty) \\
& \quad \quad \forall \forall p1, p2 \in P : (\forall d[p1] = NULL \\
& \quad \quad \quad \quad \vee d[p1][p2] \neq v1) \\
& \quad \quad \forall \forall p \in NotCrashed : \exists q \in P : d[p][q] = v1) \\
\wedge \forall p \in NotCrashed : \forall q \in P : (\forall d[p][q] = NULL \\
& \quad \quad \quad \forall \forall c \in NotFaulty : \\
& \quad \quad \quad \quad d[c][q] = d[p][q]) \\
\wedge \quad vals \stackrel{\Delta}{=} \{d[p] : p \in DecidedSomethingInteresting\} \\
& \quad \wedge Cardinality(vals) \leq 2 \\
& \quad \wedge Cardinality(vals) \geq 1 \\
& \quad \wedge \forall d1, d2 \in vals : dvSubsetEq(d1, d2) \vee dvSubsetEq(d2, d1) \\
& \quad \wedge \forall d1, d2 \in vals : \\
& \quad \quad \vee d1 = d2 \\
& \quad \quad \vee \exists S1, S2 \in SurvivorSet : \\
& \quad \quad \quad (\wedge \forall p \in S1 : (\vee p \in crashed \\
& \quad \quad \quad \quad \vee d[p] = d1) \\
& \quad \quad \quad \wedge \forall p \in S2 : (\wedge d[p] = d2 \\
& \quad \quad \quad \quad \wedge p \in NotFaulty))
\end{aligned}$$

$OnlyGoodDecisions \stackrel{\Delta}{=} Terminated \Rightarrow GoodDecision$

$ROESpec \Rightarrow \square ROETypeOK$

$ROESpec \Rightarrow \square OnlyGoodDecisions$