Safe Permissionless Consensus

Youer Pu
Cornell University

Lorenzo Alvisi
Cornell University

Ittay Eyal
The Technion

August 23, 2022

Abstract

Nakamoto's consensus protocol works in a permissionless model, where nodes can join and leave without notice. However, it guarantees agreement only probabilistically. Is this weaker guarantee a necessary concession to the severe demands of supporting a permissionless model? This paper shows that, at least in a benign failure model, it is not. It presents Sandglass, the first permissionless consensus algorithm that guarantees deterministic agreement and termination with probability 1 under general omission failures. Like Nakamoto, Sandglass adopts a *hybrid synchronous* communication model, where, at all times, a majority of nodes (though their number is unknown) are correct and synchronously connected, and allows nodes to join and leave at any time.

1 Introduction

The publication of Bitcoin's white paper [22], besides jumpstarting an industry whose market is expected to reach over $67B by 2026 [26], presented the distributed computing community with a fundamental question [12]: how should the agreement protocol at the core of Nakamoto's blockchain construction (henceforth, *Nakamoto's Consensus* or *NC*) be understood in light of the combination of consensus and state machine replication that the community has studied for over 30 years? The similarities are striking: in both cases, the goal is to create an append-only distributed ledger that everyone agrees upon, which NC calls a *blockchain*. But so are the differences. Unlike traditional consensus algorithms, where the set of participants n is known and can only be changed by running an explicit reconfiguration protocol, Nakamoto's consensus is *permissionless*: it does not enforce access control and allows the number and identity of participants to change without notice. It only assumes that the computing power of the entire system is bounded, which effectively translates to assuming

the existence of an upper bound $\mathcal{N}$ on the number of participants. [1]

To operate under these much weaker assumptions, NC adopts a new mechanism for reaching agreement: since the precise value of n is unknown, NC forsakes explicit majority voting and relies instead on a *Proof of Work* (PoW) lottery mechanism [22], designed to drive agreement towards the blockchain whose construction required the majority of the computational power of all participants. Finally, whereas traditional consensus protocols guarantee agreement deterministically, NC can do so only probabilistically; furthermore, that probability approaches 1 only as termination time approaches infinity. Is settling for these weaker guarantees the inevitable price of running consensus in a permissionless setting?

In this paper we show that, at least in a benign failure model, one can do much better. We present *Sandglass*, a permissionless consensus algorithm that guarantees *deterministic* agreement and terminates with probability 1. It operates in a model based on Nakamoto's. Our model allows an arbitrary number of participants to join and leave the system at any time and stipulates that at no time the number of participants exceeds an upper bound $\mathcal{N}$ (though the *actual* number n of participants at any given time is unknown). Further, like Nakamoto's, it is *hybrid synchronous*, in that, at all times, a majority of participants are correct and able to communicate synchronously with one another. We call these participants *good*; our protocol's safety and liveness guarantees apply to them. Participants that are not good (whether because they crash, perform omission failures, and/or experience asynchronous network connections) we call *defective*. Sandglass proceeds in asynchronous rounds, with a structure surprisingly reminiscent of Ben-Or's classic consensus protocol [3]. Let's review it. Nodes propose a value by broadcasting it; in the first round, each node proposes its initial value; in subsequent rounds, nodes propose a value chosen among those received in the previous round. Values come with an associated priority, initialized to 0. The priority of v depends on the number of consecutive rounds during which v was the only value received by the node proposing v – whenever a node receives a value other than v, it resets v's priority back to 0. When proposing a value in a given round, node p selects the highest priority value received in the previous round; if multiple values have the same priority, then it selects randomly among them. A node can safely decide a value v after sufficiently many consecutive rounds in which the proposals it receives unanimously endorse v (i.e., when v's priority is sufficiently high); and termination follows from the non-zero probability that the necessary sequence of unanimous, consecutive rounds will actually eventually occur.

Of course, embedding this structure in a permissionless setting introduces unprecedented challenges. Consider, for example, how nodes decide. In Ben-Or, a node decides v after observing two consecutive, unanimous endorsements of v; it can do so safely because any two majority sets of its fixed set of n nodes

[1] The bound can be trivial, e.g., equal to the number of atoms in the Universe, but it needs to exist; otherwise, if it would be possible for a large, unknown group of nodes to be secretly adding blocks onto a different branch of the blockchain, and Nakamoto's decisions would never be, even probabilistically, safe.

intersect in at least one correct node. This approach is clearly no longer feasible in a permissionless setting, where n is unknown and the set of nodes can change at any time.

Instead, Sandglass's approach to establish safety is inspired by one of the key properties of Nakamoto's PoW: whatever the value of n, whatever the identity of the nodes participating in the protocol at any time, the synchronously connected majority of *good nodes* will, in expectation, be faster than the remaining nodes in adding a new block to the blockchain.

Think now of adding a block b at position i of the blockchain as implicitly starting a new round of consensus for all the chain's positions that precede i; for each position, the new round proposes the corresponding block in the hash chain that ends at b. In this light, the greater speed in adding blocks that PoW promises to the majority of connected nodes translates into these nodes moving faster from one asynchronous round to the next in each of the consensus instances.

This insight suggests an alternative avenue for achieving deterministic consensus among good nodes – without relying on quorum intersection. Node p should decide on a value v only after it has seen v unanimously endorsed for sufficiently many rounds that, if p is good, the lead p (and all other good nodes) have gained over any defective node q proposing some other value is so large that q's proposals can no longer affect the proposal of good nodes.

Why can't the same approach be used to achieve deterministic consensus in Nakamoto's original protocol? Because Nakamoto's PoW mechanism, notwithstanding its name, is an indirect and imperfect vehicle for proving work. As evidence of performed work, Nakamoto presents the solution to a puzzle: this solution, however, could just have been produced as a result of a lucky guess. Thus, however unlikely, it is always possible in NC for defective nodes proposing a value other than v to catch up with, or even overtake, good nodes and reverse their decisions.

To avoid this danger, Sandglass relies on a different PoW mechanism, which ties the ability to propose a value to a *deterministic* amount of work. In particular, Sandglass nodes can propose a value in any round other than the first only after they have received a specific threshold of messages from the previous round. Therefore, each proposed value implicitly represents all the work required to generate the messages needed to clear the threshold. The threshold value is chosen as a function of the upper bound $\mathcal{N}$ on the number of nodes that at any time run the protocol, in such a way that, whatever is their actual number n, any node that does not receive messages from good nodes will inevitably take longer than them in moving from round to round.

The full power of this PoW mechanism, however, comes from pairing it with the idea, which we borrow from Ben Or, of associating a priority with the values being proposed. With a fixed set of n nodes, Ben Or leverages priorities and quorum intersection to safely decide a value v once it has reached priority 2, because it can guarantee that henceforth every node executing in the same round as a correct node will propose v. In a permissionless setting, we show that the combination of priorities and our PoW mechanism allows Sandglass to

offer good nodes the same guarantee (though, as we will see, v will be required to reach a significantly higher priority value!). Intuitively, by the time v reaches the priority necessary to decide, any node q that manages not to fall behind (and thus become irrelevant) to the unanimous majority of good nodes who have kept proposing v must have received *some* of the messages proposing v from some good nodes. Furthermore, to keep up, q must have received such messages often enough that, given how the priority of received values determines what a node can propose, it would be impossible for q to propose any value other than v.

In summary, this paper makes the following contributions: (*i*) it formalizes Nakamoto's permissionless model in the vocabulary of traditional consensus analysis; (*ii*) it introduces novel proof strategies suitable for this new model; (*iii*) it exposes the connection between PoW and a voting mechanism that can be implemented by message passing; and (*iv*) it introduces Sandglass, the first protocol that achieves deterministic agreement in a permissionless setting under hybrid synchrony.

2 Related work

The consensus problem has been studied for decades, covering both benign and Byzantine faults under different synchrony assumptions. Common across these classic works is the assumption that the set of nodes that participate in running the protocol is either constant or changes through an agreement among the incumbents (*reconfiguration*). In contrast, Sandglass allows for participants to change arbitrarily and without any coordination, as long as at all times a majority of nodes is correct and synchronously connected. More recent papers also explore models where participants can change dynamically at any time, subject to guarantees of well-behaved majority; unlike Sandglass, those works achieve only probabilistic safety guarantees. We briefly review related prior work in more detail below.

The permissionless nature of our model implies that consensus solutions for classical models (e.g., [13]) do not apply. For synchronous networks, previous solutions rely on the fact that the number of failures is bounded in a period of time. They tolerate up to $(n-1)$ benign failures [28] or Byzantine failures with authentication [7, 18]. For an asynchronous network, Fisher, Lynch, and Paterson [8] show that it is impossible to solve consensus with deterministic safety and liveness even with a single crash failure. Various protocols (e.g., [16, 25, 27]) thus either solve asynchronous consensus with weaker liveness guarantees than deterministic termination, or provide deterministic termination after a Global Stabilization Time (GST) (e.g., [4]). They use logical rounds, and for each round collect messages from a sufficient number of (authenticated) nodes, tolerating fewer than $\frac{n}{2}$ failures in a benign failure model [3, 17], and $\frac{n}{3}$ failures with Byzantine failures and authentication [4, 30]. Although our model is not directly comparable, we note that our protocol matches the $(n/2)$ bound of a benign model in an asynchronous network, despite assuming synchrony among good nodes.

Aspnes et al. [2] explore the consensus problem in an asynchronous benign model where an unbounded number of nodes can join and leave [9], but where at least one node is required to live forever, or until termination. It is easy to see that in their model, but without this latter assumption, deterministic safety is impossible. In contrast, Sandglass, in a hybrid synchronous model, guarantees deterministic safety while allowing *all* nodes to freely join and leave.

Consider two groups of good nodes with different initial values running from $t = 0$, with messages within the groups delivered immediately, but messages between the two groups are delayed until at least one in each group decides. By validity of consensus, the two groups will decide on different values, which violates agreement.

A newer line of work, starting with Nakamoto [22], studies systems where principals can unilaterally join or leave without notifying previous participants. These protocols (e.g., [21, 29]) are based on probabilistic assumptions and provide probabilistic guarantees. Specifically, participation is based on probabilistic proofs of work, and the assumption that no minority can find most proofs of work in a long period. They provide safety with high probability, given a sufficiently long running time (latency) [6, 15, 23, 10]. Nonetheless, they are all based on probabilistic techniques and provide probabilistic guarantees, which cannot be directly translated to deterministic guarantees.

Several protocols, inspired by the PoW approach, achieve consensus among a large group of principals while requiring the active participation of only a subset of them. In the Sleepy Model [24] participants join and leave ("sleep"); the assumptions and guarantees of the consensus protocol presented for this model are as probabilistic as those of pure proof of work. Momose and Ren [20] present a consensus protocol in the Sleepy Model with constant latency and deterministic agreement; however, their protocol does not guarantee progress until the participation is stable. Ouroboros [14, 5] forms a chain in the spirit of PoW but using internal tokens for the random choice of participants, again leading to probabilistic guarantees. In Algorand [11], committees elect one another in a series of reconfigurations, with assumptions and guarantees similar to classical consensus, except that participants are chosen at random from a large pool, with a negligible probability of a Byzantine majority – again, providing probabilistic guarantees.

In contrast, despite its permissionless model, our protocol guarantees deterministic safety and terminates with probability 1. We note other differences: Sandglass's failure model assumes only benign failure and is thus stronger than the Byzantine model adopted by many of these works, but its network assumptions are weaker, as defective nodes can experience asynchronous communication, and all nodes can join or leave instantaneously.

Abraham and Malkhi [1] formalize Nakamoto's Consensus within a classical disturbed systems framework, and in particular abstract the PoW primitive as a Pre-Commit, Non-Equivocation, Leader Election (PCNELE) Oracle. However, the power of this probabilistic oracle is similar to that of Nakamoto's PoW and yields a consensus protocol that provides only probabilistic guarantees as Nakamoto.

Lewis-Pye and Roughgarden [19] show that deterministic consensus cannot be achieved in a permissionless synchronous model with Byzantine nodes, let alone in a partially synchronous model (where communication becomes synchronous only after some point unknown to the processes). Sandglass shows, for the first time, that deterministic safety and termination with probability 1 can be achieved in a permissionless model, though the network is hybrid-synchronous rather than synchronous. The exploration of a Byzantine model remains for future work.

3 Model

The system comprises an infinite set of nodes $p_1, p_2, \ldots$. Time progresses in discrete steps; in each step, a subset of the nodes is *active* and the rest are *inactive*. At each step, active nodes are partitioned into *good* and *defective* subsets.

We assue a hybrid synchronous model. *Good* nodes are correct, and the network that connects them to one another is synchronous; at all times, a majority of active nodes are good. *Defective* nodes may suffer from benign failures, such as crashes and omission failures, or simply lack a synchronous connection with some good node.

The system progress is orchestrated by a *scheduler*. In each step, the scheduler can activate any inactive node p_i (we say that p_i has *joined* the system) and deactivate an active node (which then *leaves* the system). The scheduler chooses which nodes to activate and deactivate arbitrarily, subject only to the following three constraints: (i) The upper bound of active nodes in any step is $\mathcal{N}$; (ii) there is at least one active node in every step; and (iii) in every step the majority of active nodes is good.

In each step where it is active, each node p_i executes the stateful protocol shown as procedure *Step* in Sandglass's pseudocode (see Algorithm 1). It can execute computations, update its state variables, and communicate with other nodes with a broadcast network. In particular, since Sandglass assumes benign failures, every active node, whether good or defective, waits for a full step to elapse before sending its next message.

The network allows each active node to *broadcast* and *receive* unauthenticated messages. Node p_i broadcasts a message m with a $Broadcast_i(m)$ instruction and receives messages broadcast by itself and others with a $Receive_i$ instruction. The network does not generate or duplicate messages, *i.e.*, if in step t a node p_i receives message m with $Receive_i$, then m was sent in some step $t' < t$.

The communication model is designed to capture the design of Nakamoto's consensus, which relies on an underlying network layer to propagate and store blocks. Nakamoto's network layer provides a shared storage of data structures, called blocks, and guarantees delivery of published blocks within a bounded time. Each block includes cryptographically secure references to all blocks seen by its creator. This allows a newly joined node to receive and validate the entire

history of published blocks. Thus, in our model, the scheduler determines when each message is delivered to each node under the following constraints.

First, propagation time is bounded between any pair of good nodes. Formally: if a good node p_i calls $Broadcast_i(m)$ in step t, and if a good node p_j calls $Receive_j$ in step $t' > t$, then m is returned, unless it was already received by p_j in an earlier call to $Receive_j$. Thus, a newly activated good node is guaranteed, upon executing its first *Receive*, to receive all messages from other good nodes broadcast in the steps prior to its activation.

Second, the network is reliable, but there is no delivery bound unless both nodes are good. Formally: For any two nodes p_i and p_j, where at least one of p_i and p_j is defective, and for a message m broadcast by p_i, if node p_j calls $Receive_j$ infinitely many times, then m is eventually delivered.

Each node is initiated when joining the system with an initial value $v_i \in \{\mathsf{a}, \mathsf{b}\}$. An active node p_i can decide by calling a $Decide_i(v)$ instruction for some value v. The goal of the nodes is to reach a consensus based on these values:

Definition 1 (Agreement)**.** *If a good node decides a value v, then no good node decides a value other than v.*

Definition 2 (Validity)**.** *If all nodes that ever join the system have initial value v and any node (whether good or defective) decides, then it decides v.*

Definition 3 (Termination)**.** *Every good node that remains active eventually decides.*

4 Protocol

To form an intuition for the mechanics of Sandglass, it is useful to compare and contrast it with Ben-Or. From a distance, the high-level structure of the two protocols is strikingly similar: execution proceeds in asynchronous rounds; progress to the next round depends on collecting a threshold of messages sent in the current round; safety and liveness depend on the correctness of a majority of nodes; and nodes decide a value v when, for sufficiently many consecutive rounds, all the messages they collect propose v. But looking a little closer, the differences are equally striking. On the one hand, Sandglass's notion of node correctness and its hybrid synchronous model are stronger than Ben-Or's. Sandglass assumes a majority of good nodes that are not only free from crashes and omissions, but also synchronously connected to one another. On the other hand, in Sandglass, unlike Ben-Or, the number n of nodes running the protocol is not only unknown, but may be changing all the time. These differences motivate four key aspects that separate the two protocols:

Choosing a threshold In Ben-Or, a node advances to a new round only after having received a message from a majority of nodes. This strict condition for achieving progress is critical to how Ben-Or establishes Agreement. Any node that, from a majority of the nodes in round r, receives a set

of messages that unanimously propose v, can be certain that (i) there cannot exist in r also a unanimous majority proposing a value other than v and (ii) no node can proceed to round $(r+1)$ unaware that v is among the values proposed in round r. Nodes that isolate themselves from a majority simply do not make any progress; and since all majority sets intersect, nodes cannot make contradictory decisions.

Unfortunately, this approach is unworkable in Sandglass: when the cardinality and membership of the majority set can change at any time, receiving messages from a majority can no longer serve as a binary switch to trigger progress. More generally, thresholds based on the cardinality of the set of nodes from which one receives messages become meaningless. Instead, Sandglass allows nodes to broadcast multiple messages during a round, one in each of the round's steps, and lets nodes move to round $(r+1)$ once they have collected a specified threshold of messages sent in round r.

Think of the threshold $\mathcal{T}$ of messages that allows a node to move to a new round as the number of grains of sands in a sandglass: a node (figuratively) flips the sandglass at the beginning of a round, and cannot move to the next until all $\mathcal{T}$ sand grains have moved to the bottom bulb. The value of $\mathcal{T}$ is the same for all nodes; the speed at which messages are collected, however—the width of their sandglass's neck—is not, and can change from step to step: if all nodes broadcast messages at the same rate, the larger the number of nodes that one receives messages from in a timely fashion, the faster it will be to reach the threshold. Thus, while in Sandglass setting a threshold cannot altogether prevent nodes that don't receive messages from a majority from making progress, it ensures that they will progressively fall behind those who do.

Exchanging messages In each step of the protocol, a node currently in round r (i) determines, on the basis of the messages received so far, what is the largest round $r_{max} \geq r$ for which it has received the required threshold of messages and (ii) broadcasts a message for round r_{max}, which includes the node's current proposed value, as well as the critical *metadata* discussed below.

Keeping history Unlike Ben-Or, Sandglass allows nodes to join the system at any time. To bring a newly activated node up to speed, each message broadcast by a node p in round r carries a *message coffer* that includes (i) the set of messages (at least $\mathcal{T}$ of them) p collected in round $r-1$ to advance to round r; (ii) recursively, the set of messages in those messages' coffers; and (iii) the set of messages p collected so far for round r.

Respecting priority In Ben-Or, a node decides v if, for two consecutive rounds, v is the only value it collects from a majority set. To ensure the safety of that decision, Ben-Or assigns a *priority* to the value v that a node p proposes: if v was unanimously proposed by all the messages p

collected in the previous round, it is given priority 1; otherwise, 0. Nodes that collect more than one value in round r, propose for round $r+1$ the one among them with the highest priority, choosing by a coin flip in the event of a tie. Sandglass uses a similar idea, although its different threshold condition requires a much longer streak of consecutive rounds where v is unanimously proposed before v's priority can be increased. To keep track of the length of that streak, every message sent in a given round r carries a *unanimity counter*, which the sender computes upon entering r.

4.1 Selecting the Threshold

Unlike Ben-Or, Sandglass's threshold condition can not altogether prevent nodes from making progress. It is perhaps surprising that, by leveraging only the assumption that at all times a majority of nodes are good (i.e., correct and synchronously connected with each other) without ever knowing precisely how many they actually are, Sandglass retains enough of the disambiguating power of intersecting majorities to ultimately yield deterministic agreement.

In essence, Sandglass succeeds by causing defective nodes that isolate themselves from the majority of nodes in the systen to fall eventually so far behind that they no longer share the same round with good nodes. At the same time, it ensures that, once some good node has decided on a value v, nodes that manage to keep pace with good nodes will never propose anything other than v.

Of course, to obtain this outcome it is critical to set $\mathcal{T}$ appropriately. Consider two nodes, one good and one defective, and suppose they flip their sandglass at the same time—*i.e.*, they enter a new round in the same step. We want that, independent of how the number of active nodes may henceforth vary at each step, if the defective node only receives messages from other defective nodes (*i.e.*, if it fails to hear from a majority of nodes), it will reach the threshold $\mathcal{T}$ at least one step later than the good node will. The following lemma shows that setting $\mathcal{T}$ to $\lceil \frac{\mathcal{N}^2}{2} \rceil$ (where $\mathcal{N}$ is the upper bound on the maximum number of nodes active in any step) does the trick.

Lemma 1. *For any k, consider any time interval comprising $(k+1)$ consecutive steps. Let the number of messages generated by good nodes and defective nodes in each step of the interval be, respectively, $g_0, g_1, ..., g_k$ and $d_0, d_1, ..., d_k$. Setting the threshold $\mathcal{T}$ to $\lceil \frac{\mathcal{N}^2}{2} \rceil$ ensures that, if $\Sigma_{i=1}^{i=k-1} g_i < \mathcal{T}$, then $\Sigma_{i=0}^{i=k} d_i < \mathcal{T}$.*

Proof. Note how the lemma does not count the messages generated by good nodes in the steps at the two ends of the interval. Recall that moving from the current round to the next requires a node to receive at least a threashold $\mathcal{T}$ of messages sent in the current round. Thus, we drop good messages from step 0 because good nodes that in step 0 enter a new round r are unable to count against the threshold for round r messages generated by good node that in step 0 are still in round $r-1$. And we similarly drop step k because good nodes may only need one of the messages sent by good nodes in step k to move to a new round – and have no use for the remaining messages in g_k.

We begin by observing that, when k is either 0 or 1, the lemma trivially holds, since in all steps defective nodes generate fewer than $\mathcal{N}$ messages. For example, when $k = 1$, $d_0 + d_1 < \frac{\mathcal{N}}{2} + \frac{\mathcal{N}}{2} = \mathcal{N} \leq \lceil \frac{\mathcal{N}^2}{2} \rceil$. We then prove the lemma for $k \geq 2$.

Let $\bar{g} = \frac{\Sigma_{i=1}^{i=k-1} g_i}{k-1}$ and $\bar{d} = \frac{\Sigma_{i=1}^{i=k-1} d_i}{k-1}$ denote, respectively, the average number of messages generated by good nodes and by defective nodes during the $k-1$ steps that include all but the interval's first and last step. Expressed in terms of $\bar{g}$ and $\bar{d}$, the lemma requires us to show that, if $\bar{g} \cdot (k-1) < \mathcal{T}$, then $\Sigma_{i=0}^{i=k} d_i = d_0 + \bar{d} \cdot (k-1) + d_k < \mathcal{T}$ when $\mathcal{T}$ is chosen to equal $\lceil \frac{\mathcal{N}^2}{2} \rceil$.

Assume $\bar{g} \cdot (k-1) < \mathcal{T}$; then $k - 1 < \frac{\mathcal{T}}{\bar{g}}$. Substituting for $(k-1)$ in the formula that computes the messages generated by defective nodes, we have:

$$
\begin{aligned}
\Sigma_{i=0}^{i=k} d_i &= \bar{d} \cdot (k-1) + d_0 + d_k \\
&< \bar{d} \cdot \frac{\mathcal{T}}{\bar{g}} + d_0 + d_k \quad (\text{since } (k-1) < \frac{\mathcal{T}}{\bar{g}}) \\
&\leq \bar{d} \cdot \frac{\mathcal{T}}{\bar{g}} + \frac{\mathcal{N}-1}{2} + \frac{\mathcal{N}-1}{2} \quad (\text{since defective nodes are always a minority}) \\
&\leq \bar{d} \cdot \frac{\mathcal{T}}{\bar{g}} + \frac{\mathcal{T}}{\frac{\mathcal{N}^2}{2}}(\mathcal{N}-1) \quad (\text{since } \mathcal{T} = \lceil \frac{\mathcal{N}^2}{2} \rceil \geq \frac{\mathcal{N}^2}{2}) \\
&= \mathcal{T}(\frac{\bar{d}}{\bar{g}} + \frac{2(\mathcal{N}-1)}{\mathcal{N}^2}).
\end{aligned}
$$

Then, to establish that $\Sigma_{i=0}^{i=k} d_i < \mathcal{T}$, it suffices to prove that $\frac{\bar{d}}{\bar{g}} + \frac{2(\mathcal{N}-1)}{\mathcal{N}^2} < 1$.

Since for any i, $d_i \leq g_i - 1$ and $d_i + g_i \leq \mathcal{N}$, we know that $\bar{d} \leq \bar{g} - 1$ and $\bar{d} + \bar{g} \leq \mathcal{N}$. Dividing both inequalities by $\bar{g}$ yields $\frac{\bar{d}}{\bar{g}} \leq \min(1 - \frac{1}{\bar{g}}, \frac{\mathcal{N}}{\bar{g}} - 1)$. Note that the largest value of $\min(1 - \frac{1}{\bar{g}}, \frac{\mathcal{N}}{\bar{g}} - 1)$ occurs when $1 - \frac{1}{\bar{g}} = \frac{\mathcal{N}}{\bar{g}} - 1$; solving for $\bar{g}$ and plugging the solution back in, gives us: $\min(1 - \frac{1}{\bar{g}}, \frac{\mathcal{N}}{\bar{g}} - 1) \leq \frac{\mathcal{N}-1}{\mathcal{N}+1}$.

Therefore, we have that $\frac{\bar{d}}{\bar{g}} + \frac{2(\mathcal{N}-1)}{\mathcal{N}^2} \leq \frac{\mathcal{N}-1}{\mathcal{N}+1} + \frac{2(\mathcal{N}-1)}{\mathcal{N}^2} = \frac{\mathcal{N}^3+\mathcal{N}^2-2}{\mathcal{N}^3+\mathcal{N}^2} < 1$. □

4.2 Protocol Mechanics

Protocol 1, besides showing how Sandglass initializes its key variables, presents the code that node p_i executes to take a step. Every step begins with adding all received messages, as well as the messages in their message coffers, to a single set, Rec_i (lines 4 - 5). Going over the elements of that set, p_i determines the largest round r_{max} for which it has received at least a threshold $\mathcal{T}$ of messages, and, if the condition at line 6 holds, sets the current round to $(r_{max}+1)$ (line 7). Upon entering a new round, p_i does four things. First, after resetting its message coffer M, p_i collects in the coffer all the messages it received from the previous round—as well as the messages stored in the coffers of those messages (lines 8 - 10). Second, p_i chooses the value v that it will propose in the current round (lines 11 - 15): it picks the highest-priority value among those collected in its coffer for the previous round ; if more than one value qualifies, it chooses among

Protocol 1 Sandglass: Code for node p_i

1: **procedure** INIT($input_i$)
2: $v_i \leftarrow input_i$; $priority_i \leftarrow 0$; $uCounter_i \leftarrow 0$; $r_i = 1$; $M_i = \emptyset$; $Rec_i = \emptyset$; $uid_i = 0$
3: **procedure** STEP
4: **for all** $m = (\cdot, \cdot, \cdot, \cdot, \cdot, M)$ received by p_i **do**
5: $Rec_i \leftarrow Rec_i \cup \{m\} \cup M$
6: **if** $\max_{|Rec_i(r)| \geq \mathcal{T}}(r) \geq r_i$ **then**
7: $r_i = \max_{|Rec_i(r)| \geq \mathcal{T}}(r) + 1$
8: $M_i = \emptyset$
9: **for all** $m = (\cdot, r_i - 1, \cdot, \cdot, \cdot, M) \in Rec_i(r_i - 1)$ **do**
10: $M_i \leftarrow M_i \cup \{m\} \cup M$
11: Let C be the multi-set of messages in $M_i(r_i - 1)$ with the largest priority.
12: **if** all messages in C have the same value v **then**
13: $v_i \leftarrow v$
14: **else**
15: $v_i \leftarrow$ one of$\{a, b\}$, chosen uniformly at random
16: **if** all messages in $M_i(r_i - 1)$ have the same value v_i **then**
17: $uCounter_i \leftarrow 1 + \min\{uCounter | (\cdot, r_i - 1, v_i, \cdot, uCounter, \cdot) \in M_i(r_i - 1)\}$
18: **else**
19: $uCounter_i \leftarrow 0$
20: $priority_i \leftarrow \max(0, \lfloor \frac{uCounter_i}{\mathcal{T}} \rfloor - 5)$
21: **if** $priority_i \geq 6\mathcal{T} + 4$ **then**
22: Decide$_i(v_i)$
23: $uid_i \leftarrow uid_i + 1$;
24: $M_i \leftarrow M_i \cup Rec_i(r_i)$
25: broadcast $(p_i, uid_i, r_i, v_i, priority_i, uCounter_i, M_i)$

them uniformly at random. Third, p_i computes the unanimity counter and the priority for all messages that p_i will broadcast during the current round (lines 16 -20). The counter represents, starting from the previous round and going backwards, the longest sequence of rounds for which all corresponding messages in p_i's coffer unanimously proposed v. The priority is simply a direct function of the value of the unanimity counter: we maintain it explicitly because it makes it easier to describe how Sandglass works. Finally, if v's priority is high enough, p_i decides v (lines 21- 22). Whether or not it starts a new round, p_i ends every step by broadcasting a message (line 25): before it is sent, the message is made unique (line 23) and p_i adds to the message's coffer all messages received for the current round (line 24).

5 Correctness: Overview

Sandglass upholds the definitions of Validity, Agreement, and Termination (with probability 1) given in Section 3. We overview the proof below, as its approach differs from proofs of classical, permissioned protocols. We defer the presentation of the full proof to Appendix A, which includes the formal statements of the lemmas we informally state below.

Validity is easily shown by induction on the round number, since if all nodes that join have the same value, there is only one value that can be sent in each round (Lemma 2). Establishing Agreement and Termination is significantly more involved, and hinges on a precise understanding of the kinematics of good and defective nodes—and how that interacts with the ability of good nodes to converge on decision value and on the number of rounds necessary to do so safely. How clustered are good nodes as they move from round to round? At what rate do good nodes gain ground over defective nodes that cut themselves out from receiving messages from good nodes? How often do defective nodes need to receive messages from good nodes to be in turn able to have their messages still be relevant to good nodes?

The answer to these and similar questions constitute the scaffolding of lemmas and corollaries on which the proofs of Agreement and Termination rely. We discuss it in greater detail below, before moving on to the proofs.

5.1 The Scaffolding

The protocol achieves several properties that facilitate the consensus proof.

First, it keeps good nodes close together as they move from round to round. Specifically, in any step two good nodes are at most one round apart (Corollary 2), and if in any step a good node is in round r, then by the next step all good nodes are guaranteed to be at least in round r (Lemma 3). However, note that defective nodes can advance faster than good ones, using a combination of messages from good nodes and messages from defective nodes that do not reach the good nodes. Nonetheless, we show that at any step a defective node is at most one round ahead of any good node (Lemma 5).

Second, the protocol guarantees information sharing among good nodes. This may appear trivial to establish, since good nodes are correct and synchronously connected, but the *laissez-faire* attitude of the permissionless model, with nodes joining and leaving without coordination at any step, complicates matters significantly, making it impossible to prove seemingly basic properties. For example, consider a good node p that, in round r and step T, proposes a value v with a positive *uCounter*. It would feel natural to infer that all good nodes must have proposed v in the previous round—but it would also be wrong. If p just entered r in step T, it would in fact ignore any value proposed by good nodes that newly joined the systems in step T, but are still in round $r - 1$. Fortunately, we show that a much weaker form of information sharing among good nodes is sufficient to carry the day. We say that a node *collects* a message in a round if it receives the message and does not ignore it (messages originated

from a lower round number are ignored). We show that, in any round, a good node collects at least one message from a good node (Lemma 6), and that, for any round, there exists a message from a good node that is collected by all good nodes (Corollary 1).

Third, it allows us to establish the basis for a key insight about the kinematics of Sandglass nodes that will be crucial for proving Agreement and Termination: in the long run, the only values proposed by defective nodes that remain relevant to the outcome of consensus are those that have been, in turn, recently influenced by values proposed by good nodes. This insight stands on a series of intermediate results. We already saw (Lemma 1) that, given any sequence of steps, if good nodes cannot generate enough messages to get into the next round, neither can defective nodes, even if they, unlike good nodes, are allowed to count messages generated in the two steps at the opposite ends of the period. It follows that during the steps that good nodes spent in a round, defective nodes can generate fewer than the $\mathcal{T}$ messages necessary to move to the next round (Lemma 10). It all ultimately leads to Lemma 11, which quantifies the slowdown experienced by defective nodes that don't allow themselves to be contaminated by good nodes: it establishes that defective nodes that do not collect any message from good nodes for $k\mathcal{T}$ consecutive rounds fall behind every good node by at least $(k-1)$ rounds.

5.2 Agreement

The intuition behind our proof of Agreement is simple. To each value v proposed and collected by Sandglass nodes is associated a *uCounter*, which records the current streak of consecutive rounds for which all the messages collected by the proposer of v were themselves proposing v. Once v's *uCounter* reaches a certain threshold, v's *priority* increases; and once the value v proposed by a node reaches a given priority threshold, then a node decides v (see Algorithm 1, line 21). Since, as we saw, good nodes share information from round to round (recall Corollary 1), proving Agreement hinges on showing that, once a good node decides v, no good node will ever propose a value other than v. To prove that, we must in turn leverage what we learned about the kinematics of Sandglass nodes to identify a priority threshold that makes it safe for good nodes to decide. It should be large enough that, after it is reached, it becomes impossible for a defective node to change the proposal value of any good node.

The technical core of the Agreement proof then consists in establishing the truth of the following (Claim 2):

> Let p_d be the earliest good node to decide, in round r_d at step T_d. Suppose p_d decides v_d. Then, any good node p_g that in any step (whether before, at, or after T_d) finds itself in a round r_g that is at least as large as r_d, proposes v_d for r_g with *priority* at least 1. [2]

[2] Although proving Agreement does not require that v_d be proposed with priority at least 1, it makes proving the claim easier.

It is easy to see that if the above claim holds, then Agreement follows. Say that T_d is the earliest step in which a good node p_d, currently in round r_d, decides v_d. The claim immediately implies that no good node can decide a value other than v_d in a round greater or equal to r_d, since, from r_d on, every good node proposes v_d. Recall that, since good nodes are never more that one round apart at any step (Corollary 2), the earliest round a good node can find itself at T_d is $(r_d - 1)$; and that, by Lemma 3, every good node is guaranteed to be at least in round r_d by step $(T_d + 1)$. All that is left to show then is that no good node p', which at T_d found itself in round $(r_d - 1)$, can decide some value v' other than v_d. To this end, we leverage the information sharing that we proved exists among good nodes.

By Corollary 1, there is at least a message m generated in round $(r_d - 2)$ by a good node that is collected by all the good nodes. Since p_d at T_d has reached the priority threshold required to decide v_d, m must have proposed v_d; but if so, it would be impossible for good node p', which also must have collected m, to have reached the priority threshold required to decide a different value v'.

Proving Claim 2 is non trivial. The core of the proof consists in showing that any node that proposes a value v' other than the decided value v_d must find itself, at T_d, in a much earlier round than the earliest round occupied by any good node. In fact, we show something stronger: we choose a priority threshold large enough that any node, whether good or defective, that at T_d or later is within earshot of a good node (*i.e.*, whose message m can be collected by a good node), not only proposes v_d, but it does so with a *uCounter* large enough that allows whoever collects m to propose v_d with priority at least 1.

To see why those who propose v' are so far behind good nodes, note that the good node p_d that decided v_d at T_d must have received only messages proposing v_d for a long sequence of rounds, so long as to push v_d's priority over the $(6\mathcal{T}+4)$ threshold required for a decision. Let's zoom in on that sequence of rounds. It took $6\mathcal{T}$ unanimous rounds for v_d to reach priority 1 (see Algorithm 1, line 20); after clearing that first hurdle, v_d's priority increased by 1 every $\mathcal{T}$ rounds.

Consider now the set S of messages collected by p_d during the long climb that took v_d's priority from 1 to $(6\mathcal{T} + 4)$. Any node p' that during this climb proposes something other than v_d faces a dilemma. It can either refuse to collect any message in S — but if it does so, it will advance more slowly than good nodes, and, by the time v_d's priority reaches the decision threshold, it will be so far behind that no good node will collect its messages. Or p' can try to keep up by collecting messages from S — but, if it wants to keep proposing $v' \neq v_d$, it can do so in at most one round during the entire climb: since the first message collected from S would reset v' priority to 0, any further message from S collected by p' in later rounds would have higher priority than the one of v', forcing p' to henceforth propose v_d instead of v'.

In short, since p' can collect messages from S in at most one round, to ensure that any node that in round r_d is within earshot of good nodes will propose v_d it suffices to choose a large enough priority threshold for deciding. In particular, setting the threshold to $(6\mathcal{T}+4)$ ensures that (i) all messages collected by good nodes for round $(r_d - 1)$ will propose v_d, and (ii) v_d's *uCounter* in all these

messages is at least $(6\mathcal{T} - 1)$, ensuring that all good nodes in round r_d will propose v_d with *uCounter* at least $6\mathcal{T}$, *i.e.*, with priority at least 1.

Finally, a simple induction argument shows that, if all good nodes propose v_d with priority at least 1 from r_d on, then any node that, from step $(T_d + 1)$ on, continues to propose a value other than v_d, will fall ever more behind good nodes, as it will be allowed to collect messages from good nodes only once every $6\mathcal{T}$ rounds, on pain of being forced to switch its proposed value to v_d.

5.3 Termination

The Termination property requires good nodes that stay active to eventually decide. Sandglass's Termination guarantee is probabilistic: for Termination to hold, Sandglass needs to be lucky, so that it can build a sequence of consecutive rounds during which all messages collected by good nodes propose the same value; long enough that the value will reach the priority required for a node to decide. Luck is required because Sandglass allows some randomness in the values that a node proposes: nodes are required to propose the highest priority value from any message collected in the previous round, but, if they receive multiple values with the same priority, they can choose among them uniformly at random.

To help us prove that luck befalls Sandglass with probability 1, we introduce the interdependent notions of *lucky period*, *lucky value*, and *lucky round*. Intuitively, a lucky period is a sequence of steps that leads to a decision: all nodes that are active in the step that immediately follows the end of the lucky period are guaranteed to decide in that step, if not earlier. A lucky round is simply the first round of a lucky period. What is more interesting is the quality that makes a period lucky: during a lucky period, whenever Sandglass allows nodes to use randomness in picking which value they will propose in the current round, they select the same value — the *lucky value* for that round.

A minimum requirement for a round's lucky value is that it should be a *plausible* value on which good nodes may converge, in the sense that it should not explicitly go counter the value that some good node is required to propose in that round. Concretely, if the messages collected by a good node require it to propose v and all other nodes can randomly choose between v and $\overline{v}$, then the round's lucky value better not be $\overline{v}$. In addition, to encourage the possibility of a lucky period, the lucky value should be *sticky*: we would like random choices to consistently pick the same value, round after round, unless doing so would make the value implausible.

In the end, Sandglass adopts a definition of lucky value (see Appendix A.4) that, in addition to upholding plausibility, has two additional properties that express its stickiness. First, in every round good nodes collect at least one message that proposes the lucky value of the previous round (Observation 2): this guarantees that under no circumstances the previous round's lucky value will simply be forgotten when moving to a new round. The second property, which builds upon the first, establishes that lucky values don't flip easily: (Observation 3): for the lucky value in the current round to change, some good node

must have collected a different value with priority at least 1 from the previous round.

To prove that Sandglass guarantees Termination with probability 1, we then proceed in two steps.

First, we show (Observation 5) that the *uCounter* of all good nodes active in the step that immediately follows the end of the lucky period reaches a value that allows these good nodes to decide. To this end, we begin by proving that, in any lucky period, the lucky value after a while becomes *locked*: specifically, we show that the lucky value v_ℓ at round $6\mathcal{T}$ in the lucky period remains the lucky value until the end of the lucky period, and, further, that after that round all good nodes propose v_ℓ. Then, leveraging techniques similar to those used to prove Agreement, we show that any node p' that proposes a value v' other that v_ℓ must fall behind good nodes during the lucky period. The reason is that, once v_ℓ is locked, p' can collect a message from a good node only every $6\mathcal{T}$ rounds. If it did it more often, p' would collect a message proposing v_ℓ from a good node while v' has priority 0, which would force p' to change its proposal to v_ℓ – even if v' and v_ℓ both had priority 0, and p' could choose randomly among them, it would have to propose v_ℓ in the next round, since v_ℓ is the lucky value. Thus, by choosing a sufficiently long lucky period, we ensure that nodes that propose values other that v_ℓ fall so far behind good nodes that v_ℓ's priority, for any good node that is active in the step right after the end of a lucky period, reaches the threshold necessary for deciding.

Second, we show that lucky periods occur with non-zero probability, since the probability of a certain outcome of random choices for a finite number of nodes during a finite number of steps is non-zero. Since in any infinite execution lucky periods appear infinitely often, it follows that any good node that stays active, no matter when it joins, is guaranteed to eventually decide.

6 Conclusion

Sandglass shows, for the first time, that it is possible to obtain consensus with deterministic safety in a permissionless model. This result suggests that it is the probabilistic nature of its PoW mechanism, rather that its permissionless model, that prevents deterministic safety in Nakamoto's consensus. It also opens up several additional interesting questions. First among them is to understand how the interplay between permissionlessness and the hybrid synchronous model shape the boundaries of what is possible, and at what cost. As we noted, Sandglass matches the $(n/2)$ bound of a benign model in an asynchronous network, even though a majority of its nodes are synchronously connected. Perhaps at the root of this result is that in both an asynchronous model and a permissionless hybrid one it is impossible for a node to know when it has received all the messages that were intended for it. Regardless, whether there exists a protocol that achieves deterministic safety and termination in a hybrid synchronous model remains an open question. Another natural question is whether there exists a deterministic solution to consensus in a hybrid-synchronous model with Byzantine failures.

Answering these questions might pave the way to a qualitative improvement of permissionless systems that would provide deterministic guarantees; or, at the very least, give us more insight about the nature of consensus.

Acknowledgments This work was supported in part by the NSF grant CNS-CORE 2106954, BSF and IC3.

References

[1] Ittai Abraham, Dahlia Malkhi, et al. The blockchain consensus layer and bft. *Bulletin of EATCS*, 3(123), 2017.

[2] James Aspnes, Gauri Shah, and Jatin Shah. Wait-free consensus with infinite arrivals. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 524–533, 2002.

[3] Michael Ben-Or. Another advantage of free choice (extended abstract): Completely asynchronous agreement protocols. In *Proceedings of the Second Annual ACM Symposium on Principles of Distributed Computing*, pages 27–30. ACM, 1983.

[4] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.

[5] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. Ouroboros Praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 66–98. Springer, 2018.

[6] Amir Dembo, Sreeram Kannan, Ertem Nusret Tas, David Tse, Pramod Viswanath, Xuechao Wang, and Ofer Zeitouni. Everything is a race and Nakamoto always wins. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 859–878, 2020.

[7] Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.

[8] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. Technical report, MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR COMPUTER SCIENCE, 1982.

[9] Eli Gafni, Michael Merritt, and Gadi Taubenfeld. The concurrency hierarchy, and algorithms for unbounded concurrency. In *Proceedings of the twentieth annual ACM symposium on Principles of distributed computing*, pages 161–169, 2001.

[10] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The Bitcoin backbone protocol: Analysis and applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 281–310. Springer, 2015.

[11] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th symposium on operating systems principles*, pages 51–68, 2017.

[12] Maurice Herlihy. Blockchains and the future of distributed computing. In *Proceedings of the 2017 ACM Symposium on Principles of Distributed Computing (PODC '17)*, page 155, August 2017. Keynote Address.

[13] Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is DAG. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, pages 165–175, 2021.

[14] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual international cryptology conference*, pages 357–388. Springer, 2017.

[15] Lucianna Kiffer, Rajmohan Rajaraman, and Abhi Shelat. A better method to analyze blockchain consistency. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 729–744, 2018.

[16] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. Zyzzyva: speculative Byzantine fault tolerance. *Communications of the ACM*, 51(11):86–95, November 2008.

[17] Leslie Lamport. The part-time parliament. *ACM Transactions on Computer Systems (TOCS)*, 16(2):133–169, 1998.

[18] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3):382–401, 1982.

[19] Andrew Lewis-Pye and Tim Roughgarden. Byzantine generals in the permissionless setting. *arXiv preprint arXiv:2101.07095*, 2021.

[20] Atsuki Momose and Ling Ren. Constant latency in sleepy consensus. *Cryptology ePrint Archive*, 2022.

[21] Tal Moran and Ilan Orlov. Simple proofs of space-time and rational proofs of storage. In *Annual International Cryptology Conference*, pages 381–409. Springer, 2019.

[22] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, Dec 2008. Accessed: 2015-07-01. URL: `https://bitcoin.org/bitcoin.pdf`.

[23] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2017.

[24] Rafael Pass and Elaine Shi. The sleepy model of consensus. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 380–409. Springer, 2017.

[25] Michael O Rabin. Randomized byzantine generals. In *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*, pages 403–409. IEEE, 1983.

[26] Research and Markets. Blockchain market with covid-19 impact analysis, by component (platforms and services), provider (application, middleware, and infrastructure), type (private, public, and hybrid), organization size, application area, and region - global forecast to 2026. https://www.researchandmarkets.com, November 2021.

[27] Yee Jiun Song and Robbert van Renesse. Bosco: One-step Byzantine asynchronous consensus. In *International Symposium on Distributed Computing*, pages 438–450. Springer, 2008.

[28] TK Srikanth and Sam Toueg. Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distributed Computing*, 2(2):80–94, 1987.

[29] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.

[30] Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: BFT consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 347–356, 2019.

A Correctness

We prove that Sandglass satisfies the consensus requirements. Figure 1 illustrates the dependency between the statements proven below. The letters L, O, and C, signify Lemma, Observation, and Corollary, respectively.

Sandglass upholds the definitions of Valdity, Agreement, and Termination (with probability 1) given in Section 3.

A.1 Validity

We show that if all nodes have the same initial value, this is the only value that can be decided.

Lemma 2 (Validity)**.** *If all nodes that ever join the system have initial value v and any node (whether good or defective) decides, then it decides v.*

Proof. By line 22 of Sandglass, if a node p_i decides a value, it decides the value held in its variable v_i. By lines 2 and 12– 15 of Sandglass, v_i is either the initial value of p_i, or one of the values that p_i receives. Therefore, it suffices to show that if all nodes have initial value v, then v is the only value that can be sent by any node.

We prove, by induction on the round number, that any message m sent by any node for round r proposes v.

Base case: $r = 1$ Consider any node p that sends a message in the first round. Since every node's initial value is v, the message that p broadcasts at line 25 proposes v.

Induction hypothesis: Assume that all messages sent by any node up to round $r = k$ propose v.

Induction step: Consider any node p sending a message in round $r = (k + 1)$ at step T. By assumption, all round k messages collected by p must be proposing v. In lines 12–15 of Sandglass,, v_i is randomly selected from among the proposed values with highest priority collected in round k. Since the only collected value is v, v_i can only be set to v. □

A.2 Scaffolding

Before addressing Agreement and Termination, we prove several statements that will serve as scaffolding for our main results.

We start with some terminology. In Sandglass, a node in round r ignores every message it receives that was sent with some round $r' < (r - 1)$ (line 9). We say that a node p *collects* message m if it adds it to M_p (line 10). We say that a node is in round r at step T if it sends a message for round r at step T.

Our first lemma establishes that good nodes are progressing almost together from round to round.

Lemma 3. *If a good node is in round r at step T, then all good nodes will be in round r or larger at step $(T + 1)$.*

Proof. Let p be a good node in round r at step T.

By line 6 of Sandglass, p must have collected at least $\mathcal{T}$ messages for round $(r-1)$ which p will then forward to all nodes in the coffer of the message m that p broadcasts at line 25.

Consider any good node p' that is in a round $r' < r$ at T or joins the system at step $(T + 1)$. Since p and p' are good, p' will receive m by $(T + 1)$ and, by line 5, add to its set $Rec_{p'}$ both m and all the messages p forwarded in m's coffer,

Figure 1: The structure of the proof through the dependencies among its constitutents lemmas, corollaries, and observations. Preparatory results discussed in the Scaffolding section are in black; red and blue denote facts used in the proofs of Agreement and Termination, respectively.

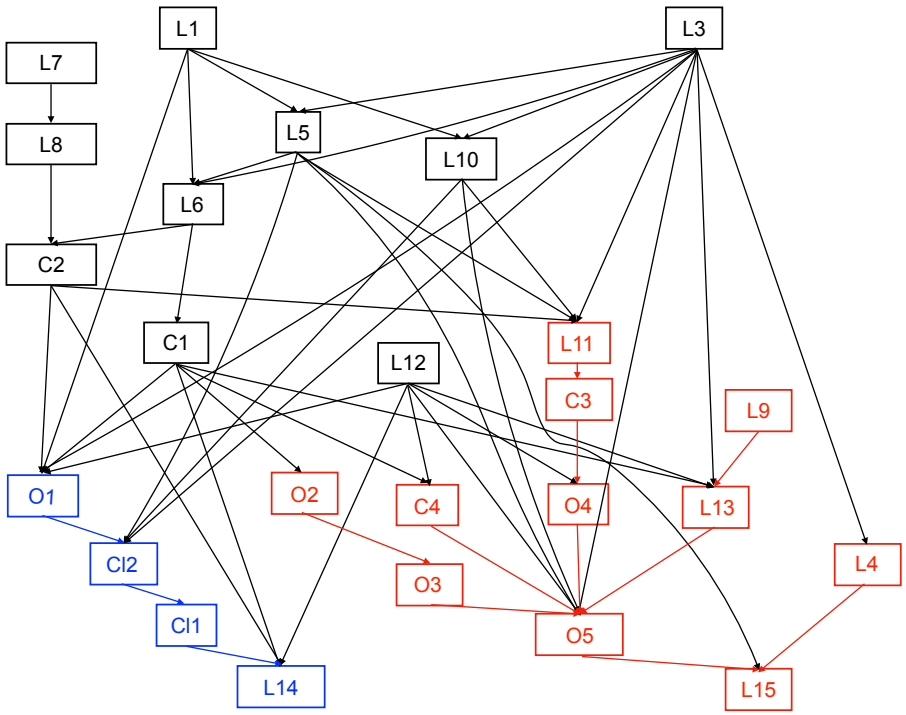

including at least $\mathcal{T}$ messages for round $(r-1)$. Then, computing at line 6 the largest round for which p' has received $\mathcal{T}$ messages or more will return at least $(r-1)$, and, at line 7, p' will update its round number, if it was smaller, to be at least r.

□

Good nodes progress by at least one round every $\mathcal{T}$ steps.

Lemma 4. *If, at step T, r is the earliest round that any good node is in, then at step $(T+\mathcal{T})$ all good nodes are at least in round $(r+1)$.*

Proof. Let p be the good node that, by hypothesis, is in round r at step T; and all the good nodes are at least in round r at step T. By Lemma 3, all good nodes are in round r or larger at step $(T+1)$; indeed, by a similar argument, all good nodes are in round r or larger for any step $T' > T$. Further, in each of these steps the system contains at least one good node, since, by assumption,

the system contains at least one node in every step and a majority of its nodes are good.

For the time interval from T to $(T + \mathcal{T} - 1)$, consider all the good nodes in each of the steps of the interval. There are two cases:

- In some step of the interval, some good node is in some round $r' > r$.

  If so, by the same reasoning used above, all good nodes will be in round $r' \geq (r + 1)$ or larger at step $(T + \mathcal{T})$.

- In all steps of the interval, all good nodes are in round r.

  If so, in each of these steps there exists at least one good node that broadcasts a message for round r (by line 25 of Sandglass). Consider any good node p_g at step $(T + \mathcal{T})$. Again, there are two cases:

  - p_g receives some message m' for round $r' \geq (r + 1)$, possibly from a defective node.

    If so, by line 5, Rec_{p_g} will include at least the $\mathcal{T}$ messages for round $(r' - 1) \geq r$ forwarded on m'.

  - p_g only receives messages for round r or smaller.

    If so, p_g receives at least $\mathcal{T}$ messages for round r.

  In both cases, computing at line 6 the largest round for which p_g has received $\mathcal{T}$ messages or more will return at least r, and, at line 7, p_g will update its round number, if it was smaller, to be at least $(r + 1)$.

□

Lemma 5. *At any step T, any defective node is at most one round ahead of any good node.*

Proof. By contradiction. Assume that there exists an earliest step, T, where some defective node p is more than one round ahead of a good node p_g, *i.e.,* at T node p is in some round r and node p_g is in round $r_{p_g} \leq (r - 2)$.

Note that no good node is in round $(r - 1)$ or larger before T; otherwise, by Lemma 3, all good nodes would be in round $(r - 1)$ or larger at T, contradicting $r_{p_g} \leq (r - 2)$. Therefore, defective node p received no messages from good nodes for round $(r - 1)$ by T.

Consider the earliest step $T' \leq (T - 1)$ where some defective node is in round $(r - 1)$. Since T is the first step where some defective node is more than a round ahead of a good node, all good nodes must be in round $(r - 2)$ or larger at T'; but, as we just showed, no good node is in round $(r - 1)$ or larger before T. Therefore, all good nodes must be in round $(r - 2)$ from T' until T.

Consider the k consecutive steps from T' to $(T - 1)$. Let the number of messages generated by good nodes and defective nodes in each step be, respectively, $g_1, ..., g_k$ and $d_1, ..., d_k$. Since by step T node p has received for round $(r - 1)$ only messages from defective nodes, and yet p is in round r at T, by line 6 of Sandglass, $\Sigma_{i=1}^{i=k} d_i \geq \mathcal{T}$ and thus, by Lemma 1, $\Sigma_{i=2}^{i=k-1} g_i \geq \mathcal{T}$. Since

by assumption every step includes at least one good node (*i.e.,* $g_1 > 0$), we have that $\Sigma_{i=1}^{i=k-1} g_i > \mathcal{T}$. Recall that during these k steps all good nodes are in round $(r-2)$; then, all messages $g_1, ..., g_k$ are for round $(r-2)$ and will all be received by all good nodes by T. By line 6 and line 7, then, *all* good nodes (including p_g) must be in round $(r-1)$ at T. This contradicts our assumption and completes the proof. □

Lemma 6. *For any r, a good node that enters round $(r+1)$ collects at least one message from a good node for round r.*

Proof. By contradiction. Let T be the first step where some good node p_g enters round $(r+1)$ without collecting any messages from any good node for round r.

Since, by line 9 of Sandglass, p_g collects all the messages it receives for round r, and yet it collects no messages from good nodes for round r, p_g must have received no messages for round r from good nodes by T.

By our model's assumptions about good nodes, this implies that no good node has sent a message for round r (and hence that no good node was in round r) before step T. Therefore, both of the following statements must be true:

S1: *Defective nodes generated at least $\mathcal{T}$ messages for round r before step T.* By line 6, a node must receive at least $\mathcal{T}$ messages for round r to be in round $(r+1)$. Since p_g received no messages for round r from good nodes before T, all the messages p_g received for round r must be from defective nodes.

S2: *No good node moved past round $(r-1)$ before T.* We have showed above that no good node is in round r before T; further, since p_g is in round $(r+1)$ at T, by Lemma 3, no good node is in a round larger than $(r+1)$ before T. Finally, no good node p'_g can be in round $(r+1)$ at $T' < T$: otherwise, since good nodes send no messages for round r before T, p'_g would not have collected any message from a good node for round r at T', contradicting our assumption that T is the first step where some good node enters round $(r+1)$ without collecting any message from good nodes for round r. Hence, before T no good node can be in round r or larger: the largest round any good node can be in is round $(r-1)$.

Let T' be the earliest step when some defective node is in round r. By Lemma 5, the earliest round that any good node can be in at T' is round $(r-1)$. Combining this observation with S2, we conclude that all good nodes are in round $(r-1)$ from T' until T.

Denote by k the number of consecutive steps from T' to $(T-1)$. Let the number of messages generated by good nodes and defective nodes in each step be, respectively, $g_1, ..., g_k$ and $d_1, ..., d_k$. Since by T node p_g has received only messages from defective nodes for round r, and yet it is in round $(r+1)$ at T, then, by line 6 of Sandglass, $\Sigma_{i=1}^{i=k} d_i \geq \mathcal{T}$ and thus, by Lemma 1, $\Sigma_{i=1}^{i=k-1} g_i \geq \mathcal{T}$. Recall that during these k steps all good nodes are in round $(r-1)$; then, all messages $g_1, ..., g_{k-1}$ are for round $(r-1)$ and will all be received by all good

nodes by $(T-1)$. By lines 6 and 7 of Sandglass, then, *all* good nodes (including p_g) must be in round r at $(T-1)$, contradicting S2. □

It follows that all good nodes collect a message from a single good node for each round.

Corollary 1. *For any round r, there exists a message from a good node for round r that is collected by all good nodes that are in round $r' \geq (r+1)$.*

Proof. By induction on r'.

Base case: $r' = (r+1)$ We are going to prove that, for any round r, there exists a message from a good node for round r that is collected by all good nodes that are in round $(r+1)$.

Consider the earliest step T when some good node p_g reaches some round $(r+1)$. By Lemma 6, p_g collects by T at least one message, m, from a good node for round r. Since m is sent by a good node, all good nodes must have received m by T. Since p_g is the earliest good node who reaches round $(r+1)$, any good node who reaches round $(r+1)$ at the same step or later must have collected m at line 10 of Sandglass.

Induction hypothesis Assume the lemma holds for $r' = k \geq (r+1)$.

Induction step We are going to prove that the lemma holds for $r' = (k+1)$. By induction hypothesis, there exists a message, $m_{r,all}$, for round r from a good node that is collected by all good nodes that are in round k; i.e., $\forall m = (\cdot, \cdot, k, \cdot, \cdot, M)$ sent by a good node for round k, $m_{r,all} \in M$.

Consider the earliest step T when some good node p_g reaches round $(k+1)$. By Lemma 6, p_g collects by T at least one message, $m_k = (\cdot, \cdot, k, \cdot, \cdot, M_k)$, from a good node for round k. As we argued above, $m_{r,all}$ must be included in M_k. Since m_k is sent by a good node, all good nodes must have received m by T. Since p_g is the earliest good node to reach round $(k+1)$, any good node that reaches round $(k+1)$ at the same step or later must have collected both m_k and all the messages in M_k, including $m_{r,all}$ at line 10 of Sandglass. Therefore, there exists a message from a good node for round r, namely $m_{r,all}$, that is collected by all the good nodes that are in round $(k+1)$. □

Lemma 7. *For any message $m = (\cdot, \cdot, r \geq 2, \cdot, \cdot, \cdot, M)$, M contains at least $\mathcal{T}$ messages generated for round $(r-1)$.*

Proof. Consider a message $m = (p, \cdot, r, \cdot, \cdot, \cdot, M_p)$ for any round $r \geq 2$.

Let T be the earliest step when p is in round r (i.e., the earliest step when p broadcasts at line 25 of Sandglass a message for round r). Independent of whether p has just been activated at T, or was already active in a round smaller than r at $(T-1)$, p's round number r_p must have been updated to r in line 7 of Sandglass. Therefore, the condition on line 6 must be satisfied, i.e., Rec_p must contain at least $\mathcal{T}$ messages for round $(r-1)$ at T. Then by lines 9-10 , M_p contains at least $\mathcal{T}$ messages for round $(r-1)$ at T. By line 24, any message p generates and broadcasts while in round r, either at T or later, contains at least $\mathcal{T}$ messages generated in round $(r-1)$. □

Lemma 8. *If a node p receives a message for round r at step T, then p will be in at least round r at step T.*

Proof. When p receives a message m generated in round r, it adds to Rec_p all the messages contained in the set M included in m (line 5 of Sandglass). By Lemma 7, M contains at least $\mathcal{T}$ messages for round $(r-1)$; thus, if r_p is smaller than r at line 6, r_p will be set to at least r at line 7. □

Corollary 2. *At any step T, if a good node is in round $r \geq 2$, then any good node is at least in round $(r-1)$.*

Proof. Consider a good node p that is in round r at T. By Lemma 6, p must have collected at least one message, m, from a good node p_{r-1} for round $(r-1)$. Consider any other good node p', it must also have received m by T. By Lemma 8, p' is at least in round $(r-1)$ at step T. □

Lemma 9. *The round number of a node never decreases.*

Proof. The lemma follows trivially since line 5 of Sandglass only adds new messages to Rec_i; thus, the set of received messages used to compute the current round number at line 7 never shrinks. □

Lemma 10. *Let T_r and T_{r+1} be the earliest steps where all good nodes are, respectively, at least in rounds r and $(r+1)$. Let g_i and d_i denote, respectively, the number of messages generated by good nodes and defective nodes in the i-th step of the sequence of k steps starting from T_r and up to step $(T_{r+1}-1)$. Then, $\Sigma_{i=1}^{k} d_i < \mathcal{T}$.*

Proof. First of all, by lemma 3, no good node is in a round smaller than r after T_r; and no good node is in a round smaller than $(r+1)$ after T_{r+1}.

If $k = 0$, i.e., $T_r = T_{r+1}$, the lemma trivially holds.

If $k \geq 1$, it suffices to establish that $\Sigma_{i=1}^{k-1} g_i < \mathcal{T}$; then, by Lemma 1, we can conclude that $\Sigma_{i=1}^{k} d_i < \mathcal{T}$ and proves the lemma.

All that is left to prove then is that $\Sigma_{i=1}^{k-1} g_i < \mathcal{T}$ holds. To do so, we begin by observing that no good node is in round $(r+1)$ or later at step $(T_{r+1}-2)$; otherwise, by Lemma 3, all good nodes would already be in round $(r+1)$ at step $(T_{r+1}-1)$, i.e., before T_{r+1}, which, by definition, is the earliest step where all good nodes are at least in round $(r+1)$.

Therefore, since all $\Sigma_{i=1}^{k-1} g_i$ messages sent by good nodes from T_r and up to $(T_{r+1}-2)$ must be at least for round r, they must be exactly for round r.

From this, it immediately follows that $\Sigma_{i=1}^{k-1} g_i$ must be less that $\mathcal{T}$ (proving the lemma): if $\Sigma_{i=1}^{k-1} g_i$ equaled or exceeded $\mathcal{T}$, then all good nodes would have received at least $\mathcal{T}$ messages for round r by step $(T_{r+1}-1)$ and thus would all be in round $(r+1)$ or larger at step $(T_{r+1}-1)$, contradicting the definition of T_{r+1}. □

The following important corollary characterizes the rate of progress experienced by defective nodes that do not collect messages from good nodes. In

particular, it establishes that defective nodes that do not collect any message from good nodes for $k\mathcal{T}$ consecutive rounds fall behind every good node by at least $(k-1)$ rounds.

Lemma 11. *Suppose a good node p_g is in round r at step T, and a node p_d is in round r_d at step $T' \leq T$. If p_d does not collect any messages from good nodes in any round $(r-i)$, where $0 \leq i < k\mathcal{T}$, then $r_d \leq (r-(k-1))$.*

Proof. To prove the corollary, we compute the maximum number of messages D_{max} that a defective node p_d can collect during the $k\mathcal{T}$ rounds when it does not collect any message from good nodes. To help us count these messages, for any $1 \leq i \leq k\mathcal{T}$, denote by $T_{(r-k\mathcal{T}+i)}$ the earliest step for which all good nodes are at least in round $(r-k\mathcal{T}+i)$.

Recall that, to be collected by p_d at step T', a message must have been generated no later than step $(T'-1) \leq (T-1)$. Then, we partition the execution of the system up to step $T-1$ into three time intervals, and compute, for each interval, the maximum number of messages generated during these intervals that p_d could have collected for rounds $(r-k\mathcal{T}+2)$ or larger.

I1: Up to step $(T_{(r-k\mathcal{T}+1)}-1)$.

By definition of $T_{(r-k\mathcal{T}+1)}$, some good node is in some round $r' < r-k\mathcal{T}+1$ at step $(T_{(r-k\mathcal{T}+1)}-1)$. Therefore, neither a defective node nor a good node can be in some round $r'' > r-k\mathcal{T}+1$ at step $(T_{(r-k\mathcal{T}+1)}-1)$, respectively because of Lemma 5 and Corollary 2. Therefore, during this interval no messages were generated for rounds $(r-k\mathcal{T}+2)$ or larger.

I2: From $T_{(r-k\mathcal{T}+1)}$ up to (T_r-1).

By assumption, p_d only collects messages generated by defective nodes throughout interval I2. We further split I2 into i consecutive subintervals, each going from $T_{(r-k\mathcal{T}+i)}$ up to $(T_{(r-k\mathcal{T}+i+1)}-1)$ for $1 \leq i \leq (k\mathcal{T}-1)$. By Lemma 10, in each of these sub-intervals defective nodes can generate at most $(\mathcal{T}-1)$ messages. Therefore, the number of messages generated by defective nodes during I2 is at most $(\mathcal{T}-1)\cdot(k\mathcal{T}-1)$.

I3: From T_r to $T-1$.

Once again, by assumption p_d only collects messages generated by defective nodes throughout interval I3. There are two cases:

- $T-1$ precedes T_r.
  If so, defective nodes trivially generate no messages during I3.
- $T-1$ does not precede T_r.
  By assumption, some good node p_g is in round r at T, where it collects all messages generated by good nodes before T; further, since p_g is still in round r, the messages for round r sent by good nodes before T must be fewer than $\mathcal{T}$. Finally, since p_g is in round r at T, by Lemma 3, in all steps preceding T no good node can be in

round $(r+1)$ or higher. We then conclude that from step T_r and up to $(T-1)$ good nodes generated at most $(\mathcal{T}-1)$ messages, all for round r. Thus, since in any step defective nodes generate fewer messages than good nodes, during I3 defective nodes generate fewer than $(\mathcal{T}-1)$ messages.

Adding the messages generated in the three intervals, we find that D_{max}, the maximum number of messages that p_d could have collected up to step T for rounds $(r-k\mathcal{T}+2)$ or larger, is smaller than $(\mathcal{T}-1)\cdot k\mathcal{T}$; at the same time, since by assumption p_d is in round r_d, D_{max} must equal at least $(r_d-(r-k\mathcal{T}+2))\cdot\mathcal{T}$. Therefore, we have that $(r_d-(r-k\mathcal{T}+2))\cdot\mathcal{T} < (\mathcal{T}-1)\cdot k\mathcal{T}$, which implies $r_d \leq r-(k-1)$, proving the corollary. □

Corollary 3. *Suppose a good node p_g is in round r at step T, and a node p_d is in round $r_d=(r-1)$ at step $T'\leq T$. p_d must have collected some message from good nodes in some round r_g, where $r-3\mathcal{T} < r_g \leq (r-1)$.*

Proof. By contradiction. Assume that node p_d is in round $(r-1)$ and has not collected any message from good nodes in any round r_g, where $r-3\mathcal{T} < r_g < r$.

Note that by T' p_d has collected no message for round r as well, for, if it had, its round number would be at least r. To see why, suppose p_d collected m_r for round r. By line 5 of Sandglass, p_d would then add to its set Rec_{p_d} both m_r and all the messages in the message coffer of m_r, including at least $\mathcal{T}$ messages for round $(r-1)$. Then, $\max_{|Rec_i(r)|\geq\mathcal{T}}(r)$ would be at least r at line 6, and p_d would update its round number to be at least r (line 7).

Therefore, p_d does not collect any message from good nodes in any round r_g, where $(r-3\mathcal{T}) < r_g \leq r$.

Then, by applying $k=3$ in Lemma 11, $r_d \leq (r-2)$. Contradiction. □

The following lemma formalizes the semantics of the unanimity counter uC included in every message; it states that the value of uC in a message that proposes v is equal to the number of consecutive rounds in which the sender of m has collected only messages that propose v.

Lemma 12. *Consider a message $m=(\cdot,\cdot,r,v,\cdot,uC>0,M)$. For $\forall m'=(\cdot,\cdot,r',v',\cdot,uC',\cdot)\in M$, where $r-uC\leq r'<r$, we have $v'=v$ and $uC'\geq uC-(r-r')$.*

Proof. By induction on uC.

Base case: $uC=1$ We are going to prove that if a node broadcasts a message $m=(\cdot,\cdot,r,v,\cdot,uC=1,M)$, then $\forall m'=(\cdot,\cdot,r',v',\cdot,uC',\cdot)\in M$, where $r'=(r-1)$, we have $v'=v$ and $uC'\geq uC-1=0$.

Establishing that $uC'\geq 0$ follows trivially from the protocol. Since $r'=(r-1)<r$, m' was added to M in line 10 of Sandglass (not in line 24). Note that Sandglass sets the value of $uCounter_i$ (at lines 16-19) only once per round, in the round's first step. Since by assumption the unanimity counter's value is 1, it must have been set at line 17; therefore, the condition at line 16 must be

satisfied. Thus, for all m' in $M_i(r-1)$, v' equals the value v broadcast in m at line 25.

Induction hypothesis Assume the lemma holds for $uC = k > 0$.

Induction step We are going to prove that the lemma holds for $uC = (k+1)$.

First, we prove that for any $m_{r-1} = (\cdot, \cdot, r-1, v_{r-1}, \cdot, uC_{r-1}, \cdot) \in M$, it holds that $uC_{r-1} \geq uC - 1 = (k+1) - 1 = k$ and $v = v_{r-1}$. Since $uC = (k+1) > 0$, the value of the minimum unanimity counter carried by all messages (including m_{r-1}) from round $(r-1)$ must be k; therefore, $uC_{r-1} \geq uC - 1 = k$. Finally, as in the Base Case, since *uCounter* is set at line 17 of Sandglass, the condition at line 16 is satisfied; therefore $v = v_{r-1}$.

Now, by line 10 of Sandglass, $\forall m' = (\cdot, \cdot, r', v', \cdot, uC', \cdot) \in M$, one of the following must be true:

Case 1 $r' = (r-1)$. It directly follows that $uC' \geq uC - 1$ and $v = v'$.

Case 2 There exists a message $m'' = (\cdot, \cdot, r'', v'', \cdot, uC'', M'') \in M$, where $r'' = (r-1)$ and $m' \in M''$. Since $r'' = (r-1)$, it follows again that $uC'' \geq uC - 1 = k$ and $v = v''$. Therefore, by the induction hypothesis, we have $\forall m_x = (\cdot, \cdot, r_x, v_x, \cdot, uC_x, \cdot) \in M''$, where $r'' - k \leq r_x < r''$, we have $v_x = v'' = v$ and $uC_x \geq uC'' - (r'' - r_x)$. Since $m' \in M''$, it follows that $v' = v$ and $uC' \geq uC'' - (r'' - r') \geq (uC - 1) - ((r-1) - r') \geq uC - (r - r')$.

□

Corollary 4. *If a good node p proposes v with $uCounter = uC$ for round r and $uC \geq 1$, then for any round r', where $r - uC \leq r' \leq r$, there exists a good node proposing v with uCounter at least $uC - (r - r')$.*

Proof. If a good node p proposes v with $uCounter = uC$ for round r and $uC \geq 1$, by Lemma 12, all the messages p collected for round r', where $r - uC \leq r' < r$ propose v with $uCounter \geq uC - (r - r')$. By Corollary 1, at least one of these messages is from a good node. Therefore, there exists a good node proposing v with *uCounter* at least $uC - (r - r')$ for round r'.

For $r' = r$, the corollary trivially holds, since p proposes v with $uCounter = uC$ for round r. □

Lemma 13. *If a good node p sends a message proposing v with $uCounter > 0$ for round r, no good node sends a message proposing $v' \neq v$ with $uCounter > 0$ for round $(r-1)$.*

Proof. By contradiction. Assume a good node p' sends a message, m', proposing $v' \neq v$ with $uCounter > 0$ for round $(r-1)$.

Let T' be the first step when p' is in round $(r-1)$, and let T be the first step when p is in round r.

Sandglass does not change the proposal value (v_i) or the priority counter ($uCounter_i$) during a round; therefore, p sends a message proposing v

with *uCounter* > 0 for round r at T; and p' sends a message proposing v' with *uCounter* > 0 for round $(r-1)$ at T'.

First, we are going to show that $T' = T$ by showing that neither $T' < T$ or $T' > T$ is possible.

Not $T' < T$ Assume $T' < T$. By model assumption, p will receive all the messages sent by good nodes on or before $T-1$, which include m'. Since T is the first step that p is in round r, the condition in line 6 of Sandglass is true, and all the messages p received for round $(r-1)$, including m', will be collected by p at line 10 at T. By lines 16-19, since m' is proposing v', it is impossible for p to propose v with non-zero *uCounter*.

Not $T' > T$ Assume $T' > T$. Since p is in round r at T, then by Lemma 3, p' is at least in round r at $(T+1)$. Then, by Lemma 9, it is impossible for p' to be in round $(r-1)$ at $T' \geq (T+1)$.

Therefore, $T' = T$, i.e. T is both the first step when p is in round r, and the first step when p' is in round $(r-1)$.

Now, we show that p' must have collected some message proposing v for round $(r-2)$.

By Corollary 1, there exists a message, $m_{all,r-1}$, from a good node for round $(r-1)$ that is collected by all the good nodes in round r, including p. We make two observations about $m_{all,r-1}$: (i) to be collected by T, $m_{all,r-1}$ must be sent before T; and (ii) since p proposes v with *uCounter* > 0, by Lemma 12, $m_{all,r-1}$ must propose v.

Let $m_{all,r-1} = (p_{all,r-1}, \cdot, r-1, v, \cdot, \cdot, M_{all,r-1})$. By lines 11-15 of Sandglass, v must be proposed by one of the round $(r-2)$ messages in $M_{all,r-1}$. Let one of the messages that propose v for round $(r-2)$ in $M_{all,r-1}$ be $m_{v,r-2}$.

Since p' is a good node, it must also have received $m_{all,r-1}$ by T. Therefore, by line 5 of Sandglass, all the messages in $M_{all,r-1}$, including $m_{v,r-2}$, are added to Rec' by p' at T.

Now, since we established that T is the first step when p' is in round $(r-1)$, p' updates its round number to $(r-1)$ at line 7 of Sandglass. Then, at line 10, p' collects all the round $(r-2)$ messages from the messages that it has received, including $m_{v,r-2}$. Since $m_{v,r-2}$ proposes v, by lines 16-19, it is impossible for p' to propose v' with *uCounter* > 0 for round $(r-1)$. Contradiction. □

A.3 Agreement

Our strategy for proving Agreement (see Definition 1) proceeds in two phases and with the help of two claims, detailed below. We begin by *assuming* that Claim 1 holds, and rely on it to prove Agreement in Lemma 14. We do not prove Claim 1 directly, however: instead, we find it easier to prove Claim 2, which implies Claim 1, thus establishing Agreement.

Claim 1. *Let p_d be the earliest good node to decide, in round r_d at step T_d. Suppose p_d decides v_d. Then, any good node p_g that in any step (whether before, at, or after T_d) finds itself in a round r_g, where $r_g \geq r_d$, proposes v_d for r_g.*

We now prove that, assuming Claim 1 holds, so does Agreement.

Lemma 14 (Agreement)**.** *If a good node decides a value v, then no good node decides a value other than v.*

Proof. Denote by U^D the value of the unanimity counter at which a node decides. By lines 20 and 21 of Sandglass, $U^D = (6\mathcal{T}+9)\mathcal{T}$. Let p_d be the first good node to decide; and suppose p_d decides v_d in round r_d at step T_d. By line 25, node p_d broadcasts at T_d a message $m_d = (p_d, \cdot, r_d, v_d, \cdot, uC_d)$, where $uC_d \geq U^D$.

By Lemma 12, all the messages p_d collected for any round $r_d - i$, $1 \leq i \leq U^D$, must be of the form $(\cdot, \cdot, r_d - i, v_d, \cdot, uCounter, \cdot)$, where $uCounter \geq U^D - i$.

By Corollary 2, at T_d no good node can be in a round earlier than $(r_d - 1)$; this implies, since p_d is the first good node to decide, that no good can decide prior to round $(r_d - 1)$.

We now show that it is impossible for any such node to decide on a value other than v_d – neither in $(r_d - 1)$, nor in r_d or in later rounds – thus proving the lemma.

Not in $(r_d - 1)$ By Corollary 1, there exists a message for round $(r_d - 2)$ broadcast by a good node that is collected by every good node that is in round $(r_d - 1)$ or larger. Let the message be m_{r_d-2}. Since m_{r_d-2} is also collected by p_d when it decides with $uCounter = U^D$, m_{r_d-2} is of the form $(\cdot, \cdot, r_d - 2, v_d, \cdot, uCounter_{r_d-2}, \cdot)$, where $uCounter_{r_d-2} \geq U^D - 2$. Consider any good node p_g in round $(r_d - 1)$. Since p_g collects m_{r_d-2} that proposes v_d in round $(r_d - 2)$, by lines 16-19 of Sandglass, the $uCounter$ of any value other than v_d proposed by p_g must be 0. Therefore, by lines 20-21, it is impossible for p_g to decide any value other than v_d in round $(r_d - 1)$.

Not in $r \geq r_d$ Trivially follows from Claim 1, any good node p_g in a round r, where $r \geq r_d$, will propose v_d, and cannot decide any value other than v_d. □

Now we have shown if Claim 1 is true, Agreement is satisfied. To complete the proof, we proceed to show Claim 1 is true, and we are going to do it by proving the following claim that implies Claim 1. Claim 2 is at least as strong as Claim 1, since it adds to Claim 1 the additional requirement shown in bold.

Claim 2. *Let p_d be the earliest good node to decide, in round r_d at step T_d. Suppose p_d decides v_d. Then, any good node p_g that in any step (whether before, at, or after T_d) finds itself in a round r_g, where $r_g \geq r_d$, proposes v_d for r_g* ***with priority at least 1****.*

Now, before proving Claim 2, we prove an observation that is useful to prove the claim.

Let U^D be the value of the unanimity counter at which a node decides. Since p_d decides at T_d, Sandglass requires p_d to broadcast at T_d a message $m_d = (p_d, \cdot, r_d, v_d, \cdot, uC_d)$, where $uC_d \geq U^D$; therefore, by Lemma 12, all the

messages that p_d has collected for round r, where $r \geq r_d - U^D$, propose v_d, and their *uCounter* is at least $U^D - (r_d - r)$.

Definition 4 ((p_d, T_d)-D-form)**.** *Given a node p_d that decides v_d in round r_d at T_d, we say that a message for round r, where $r \geq r_d - U^D$, is* in (p_d, T_d)-D-form *if it proposes v_d and the uCounter is at least $U^D - (r_d - r)$. When p_d and T_d are clear from the context, we e say simply that the message is in* D-form.

It directly follows from Lemma 12 that all the messages that p_d collects from round $(r_d - U^D)$ to round $(r_d - 1)$ are in D-form.

Observation 1. *For any round r, where $r_d - U^D \leq r \leq r_d - 1$, let $T_r^\forall$ and $T_{r+1}^\forall$ be the earliest steps where all the good nodes are at least in round r and in round $(r+1)$, respectively. Consider the set that includes messages sent by defective nodes starting from $T_r^\forall$ and before $T_{r+1}^\forall$ and messages sent by good nodes for round r. The total number of messages not in* D-form *in this set is smaller than $\mathcal{T}$.*

Proof. Let $T_r^\exists$ be the earliest step when some good node is in round r. By Corollary 1, we know that $T_r^\exists$ exists for all r.

We are going to show that:

For any round r, where $r_d - U^D \leq r \leq r_d - 1$, all the messages sent by good nodes for round r before $T_{r+1}^\exists$ must be in D-form. (F1)

We prove two cases separately.

Case 1 $r_d - U^D \leq r \leq r_d - 2$. Consider the message, $m_{all,r+1}$, that all good nodes collect for round $(r+1)$. Since p_d also collects it, $m_{all,r+1}$ must be in D-form. Consider the node $p_{all,r+1}$ that sends $m_{all,r+1}$. By definition of $T_{r+1}^\exists$, $p_{all,r+1}$ enters round $(r+1)$ at or after $T_{r+1}^\exists$; therefore, $p_{all,r+1}$ must have collected all the messages sent by good nodes for round r before $T_{r+1}^\exists$. Since $m_{all,r+1}$ is in D-form, by Lemma 12, all the messages sent by good nodes for round r before $T_{r+1}^\exists$ must also be in D-form.

Case 2 $r = r_d - 1$, note that p_d decides in round r_d, and thus also sends a message in D-form for round r_d. Since p_d enters round r_d at or after $T_{r_d}^\exists$, p_d must have collected all the messages sent by good nodes for round $(r_d - 1)$ before $T_{r_d}^\exists$. Therefore, by Lemma 12, all messages sent by good nodes for round $r = r_d - 1$ before $T_{r+1}^\exists$ must also be in D-form.

Having established F1, we proceed to prove the observation.

By Corollary 2, all good nodes are at least in round r at $T_{r+1}^\exists$, i.e., good nodes are either in round r or in round $(r+1)$ at $T_{r+1}^\exists$. If all good nodes are in round $(r+1)$, then $T_{r+1}^\forall = T_{r+1}^\exists$; otherwise, by Lemma 3, $T_{r+1}^\forall = (T_{r+1}^\exists + 1)$. We consider these two cases separately.

Let S_r be the sequence of k steps starting from $T_r^\forall$ and up to $(T_{r+1}^\forall - 1)$ (perhaps $k = 1$). Let X equal the sum of (i) the number of messages sent

by good nodes for round r that are not in D-form, and (ii) the number of messages sent by defective nodes during S_r that are not in D-form. To prove Observation 1, it is sufficient to prove $X < \mathcal{T}$.

Let d_i denote the number of messages sent by defective nodes in the i-th step of S_r. Let g_i denote the number of messages sent by good nodes for round r in the i-th step of S_r.

Case 1 $T^{\forall}_{r+1} = T^{\exists}_{r+1}$. In this case, since all messages sent by good nodes for round r are sent before $T^{\exists}_{r+1}$, by F1, all messages sent by good nodes for round r are in D-form. Therefore, X is no more than the number of messages sent by defective nodes during S_r, i.e. $X \leq \Sigma^{k}_{i=1} d_i$.

Since, in this case, no good node is in round $(r+1)$ at $(T^{\forall}_{r+1} - 1)$, the number of messages sent by good nodes for round r before $(T^{\forall}_{r+1} - 1)$ is smaller than $\mathcal{T}$; otherwise, by line 6 of Sandglass, all good nodes would have proceeded to round $(r+1)$ at $(T^{\forall}_{r+1} - 1)$. Therefore, $\Sigma^{k-1}_{i=1} g_i < \mathcal{T}$. Then, by Lemma 1, $\Sigma^{k}_{i=1} d_i < \mathcal{T}$, therefore $X < \mathcal{T}$, an we are done.

Case 2 $T^{\forall}_{r+1} = T^{\exists}_{r+1} + 1$. Again, we proved in F1 that all messages sent by good nodes for round r before $T^{\exists}_{r+1}$ are in D-form. Therefore, X is no more than the sum of the messages sent by good node for round r at $T^{\exists}_{r+1}$ and the messages sent by defective nodes during S_r, i.e., $X \leq \Sigma^{k}_{i=1} d_i + g_k$.

Now we are going to show, using a set of inequalities, that $\Sigma^{k}_{i=1} d_k + g_k < \mathcal{T}$; $X < \mathcal{T}$ directly follows.

Let $\bar{d} = \frac{\Sigma^{k-1}_{i=1} d_i}{k-1}$, and $\bar{g} = \frac{\Sigma^{k-1}_{i=1} g_i}{k-1}$. Recall that, in all steps, good nodes outnumber defective nodes. Therefore, for all $1 \leq i \leq (k-1)$, we have $d_i \leq g_i - 1$. Then, for all $1 \leq i \leq (k-1)$, since $d_i \leq g_i - 1$ and $d_i + g_i \leq \mathcal{N}$, we have that $\bar{d} \leq \bar{g} - 1$ and $\bar{d} + \bar{g} \leq \mathcal{N}$. Dividing both inequalities by $\bar{g}$ yields $\frac{\bar{d}}{\bar{g}} \leq \min(1 - \frac{1}{\bar{g}}, \frac{\mathcal{N}}{\bar{g}} - 1)$. Note that the largest value of $\min(1 - \frac{1}{\bar{g}}, \frac{\mathcal{N}}{\bar{g}} - 1)$ occurs when $1 - \frac{1}{\bar{g}} = \frac{\mathcal{N}}{\bar{g}} - 1$; solving for $\bar{g}$ and plugging the solution back in gives us: $\min(1 - \frac{1}{\bar{g}}, \frac{\mathcal{N}}{\bar{g}} - 1) \leq (1 - \frac{2}{\mathcal{N}+1})$. Therefore, we have $\frac{\bar{d}}{\bar{g}} \leq (1 - \frac{2}{\mathcal{N}+1})$, and thus

$$\bar{d} \leq \bar{g} \cdot (1 - \frac{2}{\mathcal{N}+1}). \tag{1}$$

Since $T^{\forall}_{r+1} = T^{\exists}_{r+1} + 1$, some good node is still in round r at $T^{\exists}_{r+1}$; therefore, the number of messages sent by good nodes for round r before $T^{\exists}_{r+1}$ is smaller than $\mathcal{T}$; otherwise, by line 6 of Sandglass, all good nodes would have proceeded to round $(r+1)$ at $T^{\exists}_{r+1}$. Therefore, $\Sigma^{k-1}_{i=1} g_i < \mathcal{T}$, i.e.,

$$\bar{g} \cdot (k-1) < \mathcal{T}. \tag{2}$$

Since at least one good node is already in round $(r+1)$ at $T^{\exists}_{r+1}$, the number of good nodes in round r plus the number of defective nodes at $T^{\exists}_{r+1}$ is no more than $\mathcal{N} - 1$, i.e.,

$$g_k + d_k \leq \mathcal{N} - 1. \tag{3}$$

Since good nodes outnumber defective nodes in all steps, we have for all $1 \leq i \leq (k-1)$

$$d_i \leq \frac{\mathcal{N}-1}{2}. \tag{4}$$

Now we will show $(\Sigma_{i=1}^{k} d_i + g_k) < \mathcal{T}$.

$$\begin{aligned}
(\Sigma_{i=1}^{k} d_i + g_k) &= d_k + g_k + (k-1)\bar{d} \\
&\leq (\mathcal{N}-1) + (k-1)\bar{d} && \text{(By Inequality 3)} \\
&\leq (\mathcal{N}-1) + (k-1)\cdot\bar{g}\cdot(1-\frac{2}{N+1}) && \text{(By Inequality 1)} \\
&< (\mathcal{N}-1) + \mathcal{T}\cdot(1-\frac{2}{\mathcal{N}+1}) && \text{(By Inequality 2)} \\
&= (\mathcal{N}-1) + \mathcal{T} - \frac{2}{\mathcal{N}+1}\cdot\mathcal{T} \\
&\leq \mathcal{T} - \frac{2}{\mathcal{N}+1}\cdot\frac{\mathcal{N}^2}{2} + (\mathcal{N}-1) && \text{(Since } \mathcal{T} = \lceil\tfrac{\mathcal{N}^2}{2}\rceil \geq \tfrac{\mathcal{N}^2}{2}) \\
&= \mathcal{T} - \frac{1}{\mathcal{N}+1} < \mathcal{T}
\end{aligned}$$

This concludes the second case and thus the proof. □

We can now proceed to prove Claim 2.

Proof. We are going to prove that, for any step T, if at T a good node p_g is in round $r_g \geq r_d$, then p_g proposes v_d with *priority* at least 1.

Let $U^1 = 6\mathcal{T}$ be the *uCounter* value, such that if *uCounter* is greater or equals to U^1, then *priority* is at least 1. Thus $U^D = (U^1+3)\mathcal{T} + U^1$ is the *uCounter* value that, once reached, allows a node to decide (lines 21-22 of Sandglass).

As the first step of our proof, we establish the following fact:

If a good node is in r_g at T, a node that, before T, proposes $v' \neq v_d$, can be at most in round $(r_g - U^1 - 1)$. (F2)

Assuming F2 holds, Claim 2 follows easily. Since by F2 all nodes that propose in round $(r_g - U^1)$ before T must propose v_d, then, by line 17 of Sandglass, all nodes that propose in round $(r_g - U^1 + 1)$ before T must propose v_d with *uCounter* at least 1. A simple inductive argument then shows that all nodes that ever propose in round $(r_g - U^1 + i)$ before T, where $1 \leq i < U^1$, propose v_d with *uCounter* at least i. With $i = U^1 - 1$, messages sent for $r_{v'}$ before T must propose v_d with *uCounter* at least $(U^1 - 1)$; therefore, p_g must propose v_d with *uCounter* at least U^1 at T, i.e., with *priority* at least 1.

Before, proving F2, we introduce a useful notion: For each message m sent in round r, we consider the set of messages collected by the sender of m in round $(r-1)$; we call this set m's *bag* for round $(r-1)$.

Consider some node p' that sends $m_{r'}^{v'}$ proposing v' for round r'. By line 12 of Sandglass, p' must have collected for round $(r'-1)$ a message $m_{r'-1}^{v'}$ proposing v', whose *priority* was the largest among all messages in $m_{r'}^{v'}$'s *bag*. Inductively, consider message $m_{r'-i}^{v'}$: it must in turn contain in its bag a message $m_{r'-i-1}^{v'}$ proposing v', whose *priority* is the largest among all the messages in the bag. Therefore, there exists a chain of messages extending from round 1 to round r', where each of these messages proposes v'.

Consider these messages' bags. By construction of the chain, there exists exactly one bag per round, and at least one of the messages with the highest *priority* in each bag must be proposing v'.

Let $uCounter_i^{v'}$ be the value of *uCounter* of $m_i^{v'}$. By line 17 of Sandglass, for all $1 \leq i < r'$: $uCounter_i^{v'} \geq uCounter_{i+1}^{v'} - 1$. Therefore,

$$\text{for all } 1 \leq i < j \leq r' \text{: } uCounter_i^{v'} \geq uCounter_j^{v'} - (j-i). \tag{F3}$$

We now prove F2 by induction on the round number r_g that a good node, p_g, is in.

Induction Basis: $r_g = r_d$ Suppose a good node is in round r_g at T. We proceed by contradiction: assume that before T there exists a node, p', proposing v' in some round $r' > (r_g - U^1 - 1) = (r_d - U^1 - 1)$.

Proceeding as above, we construct the chain of messages for p' and consider the bags for every round from $(r_d - (U^D - U^1) - 1)$ to $(r_d - U^1 - 1)$. We will show that (i) at most one of these bags can contain messages in D-form; and (ii) the total number of messages not in D-form sent before T is not sufficient to fill these bags. Thus, it is impossible for a node that before T proposes $v' \neq v_d$ to advance up to round $(r_d - U^1)$, contradicting our assumption and proving the basis of the induction.

Proof of (i) We show that at most one of the bags for the rounds from $(r_d - (U^D - U^1) - 1)$ to $(r_d - U^1 - 1)$ contains messages in D-form, i.e., messages that propose v_d in round r, where $r \geq r_d - U^D$, with *uCounter* at least $U^D - (r_d - r)$.

By contradiction: assume more than one of the bags for the rounds from $(r_d - (U^D - U^1) - 1)$ to $(r_d - U^1 - 1)$ contains messages in D-form. Consider two such bags, for round r_1 and r_2 respectively, where $(r_d - (U^D - U^1) - 1) \leq r_1 < r_2 \leq (r_d - U^1 - 1)$. Consider now any message $m_{r_1}^d$ in D-form contained in the bag for round r_1; $m_{r_1}^d$ proposes v_d with $uCounter_{r_1}^d \geq U^D - (r_d - r_1)$. Similarly, any message $m_{r_2}^d$ in D-form contained in the bag for round r_2 proposes v_d with $uCounter_{r_2}^d \geq U^D - (r_d - r_2)$.

Now, let us consider the messages $m^{v'}_{r_1+1}$ and $m^{v'}_{r_2}$ on the chain. We showed above (F3) that $uCounter^{v'}_{r_1+1} \geq uCounter^{v'}_{r_2} - (r_2 - (r_1+1))$. Since there is a message in D-form that proposes $v_d \neq v'$ in the bag of round r_1, by line 17 of Sandglass, $uCounter^{v'}_{r_1+1} = 0$. Therefore, $uCounter^{v'}_{r_2} \leq uCounter^{v'}_{r_1+1} + (r_2 - (r_1+1)) = r_2 - (r_1+1) \leq r_2 - (r_d - (U^D - U^1)) = U^D - U^1 - (r_d - r_2)$. Then, by line 20 of Sandglass, $priority^{v'}_{r_2} \leq \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5)$. Recall that $m^d_{r_2}$ proposes v_d with $uCounter^d_{r_2} \geq U^D - (r_d - r_2)$. Then,

$$\begin{aligned} priority^d_{r_2} &\geq \max(0, \lfloor \frac{U^D - (r_d - r_2)}{\mathcal{T}} \rfloor - 5) \\ &= \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2) + 6\mathcal{T}}{\mathcal{T}} \rfloor - 5) \quad (\text{Since } U^1 = 6\mathcal{T}) \\ &= \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1). \end{aligned}$$

Since $r_2 > r_1 \geq (r_d - (U^D - U^1) - 1)$, i.e., $r_2 \geq (r_d - (U^D - U^1))$, we have $U^D - U^1 - (r_d - r_2) \geq 0$. Then, $\lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1 \geq 1$. Therefore, $priority^d_{r_2} \geq \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1 > \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5$ and $priority^d_{r_2} \geq 1 > 0$. Therefore, $priority^d_{r_2} > \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5) \geq priority^{v'}_{r_2}$.

Now, consider $m^{v'}_{r_2+1}$. It collects both $m^{v'}_{r_2}$ and $m^d_{r_2}$. Since $m^{v'}_{r_2+1}$ proposes v', $m^{v'}_{r_2}$ must be the message with the highest priority among the messages collected by $m^{v'}_{r_2+1}$ for round r_2. However, $priority^d_{r_2} > priority^{v'}_{r_2}$. Contradiction.

Proof of (ii) Now we established that at most one of the bags for the rounds from $(r_d - (U^D - U^1) - 1)$ to $(r_d - U^1 - 1)$ contains messages in D-form. That is, among these $(U^D - 2U^1 + 1)$ bags, $(U^D - 2U^1)$ of them contain only messages that are not in D-form. We will call these *ND-bags*. Since the size of each bag is at least $\mathcal{T}$,

$$\text{ND-bags contain at least } \mathcal{T} \cdot (U^D - 2U^1) \text{ messages.} \tag{F4}$$

Let us now compute the largest number of messages they can contain. ND-bags can only contain messages not in D-form in any round from $(r_d - (U^D - U^1) - 1)$ to $(r_d - U^1 - 1)$ sent by (A) good nodes; or (B) defective nodes.

Recall that $T^{\forall}_r$ is the earliest step where all good nodes are in round r.

By Lemma 5, $T^{\forall}_{r_d - (U^D - U^1) - 2}$ is the earliest step where some defective node can be in round $(r_d - (U^D - U^1) - 1)$. Then, the messages covered by case

(B) must have been sent between steps $T^{\forall}_{r_d-(U^D-U^1)-2}$ and $(T-1)$. We can then partition this range of steps into four consecutive subranges:

B1 : from $T^{\forall}_{r_d-(U^D-U^1)-2}$ to $(T^{\forall}_{r_d-(U^D-U^1)-1}-1)$

B2 : from $T^{\forall}_{r_d-(U^D-U^1)-1}$ to $(T^{\forall}_{r_d-U^1}-1)$

B3 : from $T^{\forall}_{r_d-U^1}$ to $(T^{\forall}_{r_d}-1)$

B4 : from $T^{\forall}_{r_d}$ to $(T-1)$

We now count the total number of messages covered by cases A and $B1$ to $B4$.

B1 By Lemma 10, the number of messages in $B1$ is at most $(\mathcal{T}-1)$.

A and B2 Consider, for any round r_b, where $(r_d-(U^D-U^1)-1) \leq r_b \leq (r_d-U^1-1)$, the set of messages S_{r_b} obtained by adding (i) messages sent by defective nodes starting from $T^{\forall}_{r_b}$ and before $T^{\forall}_{r_b+1}$; and (ii) messages not in D-form sent by good nodes for round r_b. By Observation 1, S_{r_b} contains fewer than $\mathcal{T}$ messages. Thus, the set

$$\bigcup_{r_b=(r_d-(U^D-U^1)-1)}^{(r_d-U^1-1)} S_{r_b},$$

which contains all messages covered by cases A and $B2$, consists of no more than $(\mathcal{T}-1)\cdot(U^D-2U^1+1)$ messages.

B3 By Lemma 10, the number of messages sent by defective nodes in the time interval from $T^{\forall}_r$ to $(T^{\forall}_{r+1}-1)$ is at most $(\mathcal{T}-1)$. Since $B3$ contains U^1such intervals, the number of messages sent in $B3$ is at most $(\mathcal{T}-1)\cdot U^1$.

B4 Note that p_g is still in round r_d at T, and that, by Lemma 3 and the definition of $T^{\forall}_{r_d}$, all good nodes are in round r_d from $T^{\forall}_{r_d}$ to $(T-1)$. Therefore, the number of messages good nodes generate during B4 is smaller than $\mathcal{T}$; otherwise, all good nodes would be at least in round (r_d+1) at T. Since good nodes outnumber defective nodes in any step, it follows that the number of messages sent by defective nodes between $T^{\forall}_{r_d}$ and $(T-1)$ is at most $(\mathcal{T}-1)$.

Therefore, adding the number of messages in $B1$, A and $B2$, $B3$, and $B4$, ND-bags can contain no more than $(\mathcal{T}-1)+(\mathcal{T}-1)\cdot(U^D-2U^1+1)+(\mathcal{T}-1)\cdot U^1+(\mathcal{T}-1)$ messages, i.e.,

$$(\mathcal{T}-1)\cdot(U^D-U^1+3). \tag{5}$$

Recall F4: ND-bags contain at least

$$\mathcal{T}\cdot(U^D-2U^1) \text{ messages.} \tag{6}$$

Therefore, we have

$$(\mathcal{T}-1)\cdot(U^D-U^1+3)\geq\mathcal{T}\cdot(U^D-2U^1),$$

which we rewrite as $\mathcal{T}\cdot(U^D-U^1+3)-\mathcal{T}\cdot(U^D-2U^1)\geq U^D-U^1+3$, and finally as $U^D\leq(U^1+3)\mathcal{T}+U^1-3$.

However, since $U^D=(U^1+3)\mathcal{T}+U^1$, we have a contradiction. Q.E.D.

Induction hypothesis: $r_d\leq r_g\leq r_d+k$ We assume that if a good node is in r_g, where $r_d\leq r_g\leq r_d+k$, at T, then a node that, before T, proposes $v'\neq v_d$ can be at most in round (r_g-U^1-1). As we argued above, this is enough to easily show a version of Claim 2 limited to the case when $r_d\leq r_g\leq r_d+k$.
Induction step: $r_g=r_d+k+1$ Suppose a good node is in round $r_g=r_d+k+1$ at T, and that Claim 2 holds for any round r, where $r_d\leq r\leq r_d+k$. We will prove that if a good node is in r_g at T, then a node that, before T, proposes $v'\neq v_d$ can be at most in round (r_g-U^1-1).

We will assume, by contradiction, that there exists a node p' that before T proposes v' in some round $r'>(r_g-U^1-1)$. We will consider the following two cases: (1) $(r_g-U^1-1)\leq r_d-1$, i.e., $r_g\leq r_d+U^1$; and (2) $(r_g-U^1-1)>r_d-1$, i.e., $r_g>r_d+U^1$.

Case 1: $r_g\leq r_d+U^1$ Proceeding as above, we construct the chain of messages for p' and consider the bags for every round from $(r_d-(U^D-U^1)-1)$ to (r_g-U^1-1). With $r_g\leq r_d+U^1$, we have $(r_g-U^1-1)\leq r_d-1$.

We will show that (i) at most one of these bags can contain messages in D-form; and (ii) the total number of messages not in D-form sent before T is not sufficient to fill these bags. Thus, it is impossible for a node that before T proposes $v'\neq v_d$ to advance up to round (r_g-U^1), contradicting our assumption. The proof for Case 1 is very similar to how we proved the induction basis; we present it in full for completeness.

Proof of (i) We will show that at most one of the bags for the rounds from $(r_d-(U^D-U^1)-1)$ to (r_g-U^1-1) contains messages in D-form.

By contradiction: assume more than one of the bags for the rounds from $(r_d-(U^D-U^1)-1)$ to (r_g-U^1-1) contains messages in D-form. Consider two such bags, for round r_1 and r_2 respectively, where $(r_d-(U^D-U^1)-1)\leq r_1<r_2\leq(r_g-U^1-1)\leq(r_d-1)$. Consider now any message $m^d_{r_1}$ in D-form contained in the bag for round r_1; $m^d_{r_1}$ proposes v_d with $uCounter^d_{r_1}\geq U^D-(r_d-r_1)$. Similarly, any message $m^d_{r_2}$ in D-form contained in the bag for round r_2 proposes v_d with $uCounter^d_{r_2}\geq U^D-(r_d-r_2)$.

Now, let us consider the messages $m^{v'}_{r_1+1}$ and $m^{v'}_{r_2}$ on the chain. We showed above (F3) that $uCounter^{v'}_{r_1+1}\geq uCounter^{v'}_{r_2}-(r_2-(r_1+$

1)). Since there is a message in D-form that proposes $v_d \neq v'$ in the bag of round r_1, by line 17 of Sandglass, $uCounter^{v'}_{r_1+1} = 0$. Therefore, $uCounter^{v'}_{r_2} \leq uCounter^{v'}_{r_1+1} + (r_2 - (r_1+1)) = r_2 - (r_1 + 1) \leq r_2 - (r_d - (U^D - U^1)) = U^D - U^1 - (r_d - r_2)$. Then, by line 20 of Sandglass, $priority^{v'}_{r_2} \leq \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5)$.

Recall that $m^d_{r_2}$ proposes v_d with $uCounter^d_{r_2} \geq U^D - (r_d - r_2)$. Then,

$$\begin{aligned} priority^d_{r_2} &\geq \max(0, \lfloor \frac{U^D - (r_d - r_2)}{\mathcal{T}} \rfloor - 5) \\ &= \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2) + 6\mathcal{T}}{\mathcal{T}} \rfloor - 5) \\ &\qquad\qquad \text{(Since } U^1 = 6\mathcal{T}) \\ &= \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1). \end{aligned}$$

Since $r_2 > r_1 \geq (r_d - (U^D - U^1) - 1)$, i.e., $r_2 \geq (r_d - (U^D - U^1))$, we have $U^D - U^1 - (r_d - r_2) \geq 0$. Then, $\lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1 \geq 1$. Therefore, $priority^d_{r_2} \geq \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1 > \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5$ and $priority^d_{r_2} \geq 1 > 0$. Therefore, $priority^d_{r_2} > \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5) \geq priority^{v'}_{r_2}$.

Now, consider $m^{v'}_{r_2+1}$. It collects both $m^{v'}_{r_2}$ and $m^d_{r_2}$. Since $m^{v'}_{r_2+1}$ proposes v', $m^{v'}_{r_2}$ must be the message with the highest priority among the messages collected by $m^{v'}_{r_2+1}$ for round r_2. However, $priority^d_{r_2} > priority^{v'}_{r_2}$. Contradiction.

Proof of (ii) Now we established that at most one of the bags for the rounds from $(r_d - (U^D - U^1) - 1)$ to $(r_g - U^1 - 1)$ contains messages in D-form. That is, among these $(U^D - 2U^1 + (r_g - r_d) + 1)$ bags, $(U^D - 2U^1 + (r_g - r_d))$ of them contain only messages that are not in D-form. We will call these *ND-bags*. Since the size of each bag is at least $\mathcal{T}$,

$$\text{ND-bags contain at least } \mathcal{T} \cdot (U^D - 2U^1 + (r_g - r_d)) \text{ messages.} \tag{F5}$$

Let us now compute the largest number of messages they can contain. ND-bags can only contain messages not in D-form in any round from $(r_d - (U^D - U^1) - 1)$ to $(r_g - U^1 - 1)$ sent by (A) good nodes; or (B) defective nodes.

By Lemma 5, $T^{\forall}_{r_d - (U^D - U^1) - 2}$ is the earliest step where some defective node can be in round $(r_d - (U^D - U^1) - 1)$. Then, the messages covered by case (B) must have been sent from step $T^{\forall}_{r_d - (U^D - U^1) - 2}$

to step $(T-1)$. We can then partition this range of steps into four consecutive subranges:

B1 : from $T^{\forall}_{r_d-(U^D-U^1)-2}$ to $(T^{\forall}_{r_d-(U^D-U^1)-1}-1)$

B2 : from $T^{\forall}_{r_d-(U^D-U^1)-1}$ to $(T^{\forall}_{r_g-U^1}-1)$

B3 : from $T^{\forall}_{r_g-U^1}$ to $(T^{\forall}_{r_g}-1)$

B4 : from $T^{\forall}_{r_g}$ to $(T-1)$

We now count the total number of messages covered by cases A and $B1$ to $B4$.

B1 By Lemma 10, the number of messages sent in $B1$ is at most $(\mathcal{T}-1)$.

A and B2 Consider, for any round r_b, where $(r_d-(U^D-U^1)-1) \leq r_b \leq (r_g-U^1-1)$, the set of messages S_{r_b} obtained by adding (i) messages sent by defective nodes starting from $T^{\forall}_{r_b}$ and before $T^{\forall}_{r_b+1}$; and (ii) messages not in D-form sent by good nodes for round r_b. By Observation 1, S_{r_b} contains fewer than $\mathcal{T}$ messages. Thus, the set

$$\bigcup_{r_b=(r_g-(U^D-U^1)-1)}^{(r_d-U^1-1)} S_{r_b},$$

which contains all messages covered by cases A and $B2$, consists of no more than $(\mathcal{T}-1)\cdot(U^D-2U^1+(r_g-r_d)+1)$.

B3 By Lemma 10, the number of messages sent by defective nodes in the time interval from $T^{\forall}_r$ to $(T^{\forall}_{r+1}-1)$ is at most $(\mathcal{T}-1)$. Since $B3$ contains U^1such intervals, the number of messages sent in $B3$ is at most $(\mathcal{T}-1)\cdot U^1$.

B4 Note that p_g is still in round r_g at T, and that, by Lemma 3 and the definition of $T^{\forall}_{r_g}$, all good nodes are in round r_g from $T^{\forall}_{r_g}$ to $(T-1)$. Therefore, the number of messages good nodes generate during $B4$ is smaller than $\mathcal{T}$; otherwise, all good nodes would be at least in round (r_g+1) at T. Since good nodes outnumber defective nodes in any step, it follows that the number of messages sent by defective nodes between $T^{\forall}_{r_g}$ and $(T-1)$ is at most $(\mathcal{T}-1)$.

Therefore, adding the number of messages in $B1$, A and $B2$, $B3$, and $B4$, ND-bags can contain no more than $(\mathcal{T}-1)+(\mathcal{T}-1)\cdot(U^D-2U^1+(r_g-r_d)+1)+(\mathcal{T}-1)\cdot U^1+(\mathcal{T}-1)$ messages, i.e.,

$$(\mathcal{T}-1)\cdot(U^D-U^1+(r_g-r_d)+3). \tag{7}$$

Recall F5: ND-bags contain at least

$$\mathcal{T}\cdot(U^D-2U^1+(r_g-r_d)) \text{ messages.} \tag{8}$$

Therefore, we have

$$\begin{aligned}
&(\mathcal{T}-1)\cdot(U^D-U^1+(r_g-r_d)+3) \geq \mathcal{T}\cdot(U^D-2U^1+(r_g-r_d))\\
\Rightarrow\ &(\mathcal{T}-1)\cdot(U^D-U^1+(r_g-r_d)+3) \geq (\mathcal{T}-1)\cdot(U^D-2U^1+(r_g-r_d))\\
&\qquad\qquad + (U^D-2U^1+(r_g-r_d))\\
&\Rightarrow (\mathcal{T}-1)\cdot(U^1+3) \geq U^D-2U^1+(r_g-r_d)\\
&\Rightarrow (\mathcal{T}-1)\cdot(U^1+3) \geq (U^1+3)\mathcal{T}-U^1+(r_g-r_d)\\
&\qquad\qquad \text{(since } U^D=(U^1+3)\mathcal{T}+U^1\text{)}\\
&\qquad\qquad \Rightarrow 0 \geq 3+(r_g-r_d)
\end{aligned}$$

However, since $r_g \geq r_d$, we have a contradiction. Q.E.D.

Case 2: $r_g > r_d + U^1$ Again, we construct the chain of messages for p' and consider the bags for every round from $(r_d-(U^D-U^1)-1)$ to (r_g-U^1-1).

We will show that:

- (i) At most one of the bags for rounds from $(r_d-(U^D-U^1)-1)$ to (r_d-1) contains messages in D-form. That is, among these (U^D-U^1+1) bags, (U^D-U^1) of them contain only messages that are not in D-form. We will call these *ND-bags*.
- (ii) Among the bags for rounds from r_d to (r_g-U^1-1), at most one in every U^1 bags can contain messages from good nodes. That is, among these $(r_g-r_d-U^1)$ bags, $(r_g-r_d-U^1-\lceil\frac{r_g-r_d-U^1}{U^1}\rceil)$ of them contain only messages from defective nodes. We will call these *Def-bags*.
- (iii) The sum of (1) the messages not in D-form for round $(r_d-(U^D-U^1)-1)$ to (r_d-1), and (2) the messages sent by defective nodes for round r_d to (r_g-U^1-1) before T, is not sufficient to fill ND-bags and Def-bags.

Thus, it is impossible for a node that before T proposes $v' \neq v_d$ to advance up to round (r_g-U^1), contradicting our assumption.

Proof of (i) By contradiction: assume more than one of the bags for the rounds from $(r_d-(U^D-U^1)-1)$ to (r_d-1) contains messages in D-form. Consider two such bags, for round r_1 and r_2 respectively, where $(r_d-(U^D-U^1)-1) \leq r_1 < r_2 \leq (r_d-1)$. Consider now any message $m^d_{r_1}$ in D-form contained in the bag for round r_1; $m^d_{r_1}$ proposes v_d with $uCounter^d_{r_1} \geq U^D-(r_d-r_1)$. Similarly, any message $m^d_{r_2}$ in D-form contained in the bag for round r_2 proposes v_d with $uCounter^d_{r_2} \geq U^D-(r_d-r_2)$.

Now, let us consider the messages $m^{v'}_{r_1+1}$ and $m^{v'}_{r_2}$ on the chain. We showed above (F3) that $uCounter^{v'}_{r_1+1} \geq uCounter^{v'}_{r_2}-(r_2-(r_1+1))$. Since there is a message in D-form that proposes $v_d \neq v'$ in

the bag of round r_1, by line 17 of Sandglass, $uCounter_{r_1+1}^{v'} = 0$. Therefore, $uCounter_{r_2}^{v'} \leq uCounter_{r_1+1}^{v'} + (r_2 - (r_1+1)) = r_2 - (r_1 + 1) \leq r_2 - (r_d - (U^D - U^1)) = U^D - U^1 - (r_d - r_2)$. Then, by line 20 of Sandglass, $priority_{r_2}^{v'} \leq \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5)$.
Recall that $m_{r_2}^d$ proposes v_d with $uCounter_{r_2}^d \geq U^D - (r_d - r_2)$. Then,

$$\begin{aligned} priority_{r_2}^d &\geq \max(0, \lfloor \frac{U^D - (r_d - r_2)}{\mathcal{T}} \rfloor - 5) \\ &= \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2) + 6\mathcal{T}}{\mathcal{T}} \rfloor - 5) \\ &\qquad\qquad \text{(Since } U^1 = 6\mathcal{T}\text{)} \\ &= \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1). \end{aligned}$$

Since $r_2 > r_1 \geq (r_d - (U^D - U^1) - 1)$, i.e., $r_2 \geq (r_d - (U^D - U^1))$, we have $U^D - U^1 - (r_d - r_2) \geq 0$. Then, $\lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1 \geq 1$. Therefore, $priority_{r_2}^d \geq \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor + 1 > \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5$ and $priority_{r_2}^d \geq 1 > 0$. Therefore, $priority_{r_2}^d > \max(0, \lfloor \frac{U^D - U^1 - (r_d - r_2)}{\mathcal{T}} \rfloor - 5) \geq priority_{r_2}^{v'}$.

Now, consider $m_{r_2+1}^{v'}$. It collects both $m_{r_2}^{v'}$ and $m_{r_2}^d$, and $m_{r_2}^{v'}$. Since $m_{r_2+1}^{v'}$ proposes v', $m_{r_2}^{v'}$ must be the message with the highest priority among the messages collected by $m_{r_2+1}^{v'}$ for round r_2. However, $priority_{r_2}^d > priority_{r_2}^{v'}$. Contradiction.

Proof of (ii) We will show that among the bags for rounds r_d to $(r_g - U^1 - 1)$, at most one in every U^1 bags can contain messages from good nodes.

By contradiction: assume there exist two bags containing messages from good nodes, for round r_1 and r_2 respectively, where $r_d \leq r_1 < r_2 \leq (r_g - U^1 - 1)$, and $r_2 - r_1 < U^1$. Consider now any message $m_{r_1}^g$ from a good node contained in the bag for round r_1; by induction hypothesis, $m_{r_1}^g$ proposes v_d with $priority_{r_1}^g \geq 1$. Similarly, any message $m_{r_2}^g$ from a good node contained in the bag for round r_2 proposes v_d with $priority_{r_2}^g \geq 1$.

Now, let us consider messages $m_{r_1+1}^{v'}$ and $m_{r_2}^{v'}$ on the chain. We showed above (F3) that $uCounter_{r_1+1}^{v'} \geq uCounter_{r_2}^{v'} - (r_2 - (r_1+1))$. Since there is a message from a good node proposing $v_d \neq v'$ in the bag of round r_1, by line 17 of Sandglass, $uCounter_{r_1+1}^{v'} = 0$. Therefore, $uCounter_{r_2}^{v'} \leq uCounter_{r_1+1}^{v'} + (r_2 - (r_1+1)) = r_2 - r_1 - 1 < U^1$. Then, by line 20 of Sandglass, $priority_{r_2}^{v'} < 1$.

Recall that $m^g_{r_2}$ proposes v_d with $priority^g_{r_2} \geq 1$. Then, we have $priority^g_{r_2} > priority^{v'}_{r_2}$.

Now, consider $m^{v'}_{r_2+1}$. It collects both $m^{v'}_{r_2}$ and $m^g_{r_2}$. Since $m^{v'}_{r_2+1}$ proposes v', $m^{v'}_{r_2}$ must be the message with the highest priority among the messages collected by $m^{v'}_{r_2+1}$ for round r_2. However, $priority^g_{r_2} > priority^{v'}_{r_2}$. Contradiction.

Proof of (iii) Now we established that:

(i) Among the bags for rounds from $(r_d-(U^D-U^1)-1)$ to (r_d-1), (U^D-U^1) ND-bags contain only messages that are not in D-form.

(ii) Among the bags for rounds from r_d to (r_g-U^1-1), $(r_g-r_d-U^1-\lceil\frac{r_g-r_d-U^1}{U^1}\rceil)$ Def-bags contain only messages from defective nodes.

Since each bag contains at least $\mathcal{T}$ messages, ND-bags and Def-bags contain collectively at least

$$\mathcal{T}\cdot(U^D-U^1+(r_g-r_d-U^1-\lceil\frac{r_g-r_d-U^1}{U^1}\rceil)) \text{ messages.} \quad \text{(F6)}$$

Let us now compute the largest number of messages they can contain. ND-bags and Def-bags can contain messages not in D-form in any round from $(r_d-(U^D-U^1)-1)$ to (r_d-1) sent by either (A) good nodes or (B) defective nodes, and (C) messages sent for round r_d to r_g-U^1-1 by defective nodes.

By Lemma 5, $T^{\forall}_{r_d-(U^D-U^1)-2}$ is the earliest step where some defective node can be in round $(r_d-(U^D-U^1)-1)$. Then, the messages covered by case (B) and (C) must have been sent from step $T^{\forall}_{r_d-(U^D-U^1)-2}$ to step $(T-1)$. We can then partition this range of steps into four consecutive subranges:

BC1 : from $T^{\forall}_{r_d-(U^D-U^1)-2}$ to $(T^{\forall}_{r_d-(U^D-U^1)-1}-1)$

BC2 : from $T^{\forall}_{r_d-(U^D-U^1)-1}$ to $(T^{\forall}_{r_d}-1)$

BC3 : from $T^{\forall}_{r_d}$ to $(T^{\forall}_{r_g}-1)$

BC4 : from $T^{\forall}_{r_g}$ to $(T-1)$

We now count the total number of messages covered by cases A and $BC1$ to $BC4$.

BC1 By Lemma 10, the number of messages in $BC1$ is at most $(\mathcal{T}-1)$.

A and BC2 Consider, for any round r_b, where $(r_d-(U^D-U^1)-1) \leq r_b \leq (r_d-1)$, the set of messages S_{r_b} obtained by adding

(i) messages sent by defective nodes starting from $T^{\forall}_{r_b}$ and before $T^{\forall}_{r_b+1}$; and (ii) messages not in D-form sent by good nodes for round r_b. By Observation 1, S_{r_b} contains fewer than $\mathcal{T}$ messages. Thus, the set

$$\bigcup_{r_b=(r_d-(U^D-U^1)-1)}^{(r_d-1)} S_{r_b},$$

which contains all messages covered by cases A and $BC2$, consists of no more than $(\mathcal{T}-1)\cdot(U^D-U^1+1)$ messages.

BC3 By Lemma 10, the number of messages sent by defective nodes in the time interval from $T^{\forall}_r$ to $(T^{\forall}_{r+1}-1)$ is at most $(\mathcal{T}-1)$. Since $B3$ contains (r_g-r_d) such intervals, the number of messages sent in $BC3$ is at most $(\mathcal{T}-1)\cdot(r_g-r_d)$.

BC4 Note that p_g is still in round r_g at T, and that, by Lemma 3 and the definition of $T^{\forall}_{r_g}$, all good nodes are in round r_g from $T^{\forall}_{r_g}$ to $(T-1)$. Therefore, the number of messages good nodes generate during $BC4$ is smaller than $\mathcal{T}$; otherwise, all good nodes would be at least in round (r_g+1) at T. Since good nodes outnumber defective nodes in any step, it follows that the number of messages sent by defective nodes between $T^{\forall}_{r_g}$ and $(T-1)$ is at most $(\mathcal{T}-1)$.

Therefore, adding the number of messages in $BC1$, A and $BC2$, $BC3$, and $BC4$, ND-bags and Def-bags can contain no more than $(\mathcal{T}-1)+(\mathcal{T}-1)\cdot(U^D-U^1+1)+(\mathcal{T}-1)\cdot(r_g-r_d)+(\mathcal{T}-1)$ messages, i.e.,

$$(\mathcal{T}-1)\cdot(U^D-U^1+(r_g-r_d)+3). \tag{9}$$

Recall F6: ND-bags and Def-bags contain at least

$$\mathcal{T}\cdot(U^D-U^1+(r_g-r_d-U^1-\lceil\frac{r_g-r_d-U^1}{U^1}\rceil)) \text{ messages.} \tag{F6}$$

Therefore, we have

$$(\mathcal{T}-1)\cdot(U^D-U^1+(r_g-r_d)+3)\geq\mathcal{T}\cdot(U^D-U^1+(r_g-r_d-U^1-\lceil\frac{r_g-r_d-U^1}{U^1}\rceil))$$

$$\Rightarrow\mathcal{T}\cdot(U^1+3+\lceil\frac{r_g-r_d-U^1}{U^1}\rceil)\geq U^D-U^1+(r_g-r_d)+3$$

$$\Rightarrow\mathcal{T}\cdot(U^1+3+\lceil\frac{r_g-r_d-U^1}{U^1}\rceil)\geq(U^1+3)\mathcal{T}+(r_g-r_d)+3$$

(Since $U^D=((U^1+3)\mathcal{T}+U^1)$)

$$\Rightarrow\lceil\frac{r_g-r_d}{6\mathcal{T}}\rceil-1\geq\frac{(r_g-r_d)+3}{\mathcal{T}}$$

(Since $U^1=6\mathcal{T}$)

$$\Rightarrow\frac{r_g-r_d}{6\mathcal{T}}>\frac{(r_g-r_d)+3}{\mathcal{T}}$$

(Since $\frac{r_g-r_d}{6\mathcal{T}}>(\lceil\frac{r_g-r_d}{6\mathcal{T}}\rceil-1)$)

However, since $0 < (r_g - r_d) < ((r_g - r_d) + 3)$ and $(6\mathcal{T}) > \mathcal{T} > 0$, we have a contradiction. Q.E.D.

This concludes our proof of Agreement. □

A.4 Termination

The Termination property requires good nodes that stay active to eventually decide. Sandglass's Termination guarantee is probabilistic: For Termination to hold, Sandglass needs to be lucky. To help us prove that luck befalls Sandglass with probability 1, we introduce the interdependent notions of *lucky period*, *lucky value*, and *lucky round*.

Intuitively, a lucky period is a sequence of steps that leads to a decision: all nodes that are active in the step that immediately follows the end of the lucky period are guaranteed to decide in that step, if not earlier. The quality that makes a period lucky is straightforward. Recall that in Sandglass, if a node receives distinct highest priority proposals in the previous round, it can choose uniformly at random among them which one it is going to propose in the current round. During a lucky period, all the random choices that occur in a given round just happen to select the same value – the *lucky value* for that round. We give below a simple rule that defines what constitutes the lucky value for any given round spanned by the lucky period. To prove that Sandglass guarantees Termination with probability 1, we will proceed in two steps. First, we will show that the unanimity counter of all good nodes that are active during the last step of a lucky period reaches a value that allows them to decide. Second, we will prove that lucky periods occur with non-zero probability. Since in any infinite execution lucky periods appear infinitely often, it follows that any good node that stays active, no matter when it joins, is guaranteed to eventually decide.

Lucky value The rule that determines the lucky value for a given round r is defined in terms of two sets. The first, $C(r, p)$, is independently computed by every node p as the set of messages for round r defined by line 11 of Sandglass; it contains the highest-priority messages p collected for round $(r-1)$. The second set, $O(r)$, contains a (possibly empty) subset of good nodes, and is defined across all good nodes that enter round r at any time. It contains any good node p_g that meets the following two criteria: (1) p_g has collected exactly one highest priority value in round $(r-1)$ (which p_g is then required to propose in round r) and (2) one of the messages sent by p_g in round r is collected by all good nodes in round $(r+1)$. Note that if $O(r)$ contains multiple good nodes, they may differ in the single highest priority value they have collected.

We dub the first round of a lucky period a *lucky round*. The lucky value $v_\ell(r)$ for a given round r of a given lucky period is defined inductively, with the base case defined by that period's lucky round, r_{start}, as follows:

- When $r = r_{start}$:

  If $O(r_{start}) \neq \emptyset$ and $\forall p \in O(r_{start})$, $v \in C(r_{start}, p)$, then $v_\ell(r_{start}) = v$.

Otherwise, $v_\ell(r_{start})$ is arbitrarily set to one of the initial values. We will assume, without loss of generality, that $v_\ell(r_{start})$ is set to a.

- When $r > r_{start}$:

  If $O(r) \neq \emptyset$ and $\forall p \in O(r)$, $v \in C(r,p)$, then $v_\ell(r) = v$.

  Otherwise, $v_\ell(r) = v_\ell(r-1)$.

Lucky period. We already saw that, informally, a lucky round is the first round of a lucky period. To define these notions more precisely, we introduce the following definitions, which we will use extensively in our Termination proof:

- $T_1(r)$: The earliest step where some node, possibly defective, is in round r.
- r_{lock}: The round with index $(r + 6\mathcal{T})$. We will prove that, if r is a lucky round then, in *every round* from $(r_{lock}+1)$ to the end of the lucky period, v_ℓ is the same as the lucky value of round r_{lock}, and all good nodes propose the lucky value of round r_{lock}.
- $T(r_{lock})$: The earliest step where some node is in round r_{lock}.
- $\overrightarrow{P_\ell}$: A constant, equal to $(6\mathcal{T} + \lceil \frac{(6\mathcal{T}-1)\cdot U^D + 18\mathcal{T}}{5} \rceil)$, which denotes the number of rounds spanned by a lucky period, i.e., all rounds from the period's lucky round r_{start} to round $(r_{start} + \overrightarrow{P_\ell} - 1)$ (or, equivalently, round $(r_{lock} + \lceil \frac{(6\mathcal{T}-1)\cdot U^D + 18\mathcal{T}}{5} \rceil - 1)$.
- $T_D(r)$: The earliest step where all good nodes are in round $(r + \overrightarrow{P_\ell})$ or later. We will prove that, if r is a lucky round, then all good nodes decide by step $T_D(r)$.

We then say that r_{start} is a *lucky round* if, in every step during the lucky period from $T_1(r_{start})$ to $(T_D(r_{start}) - 1)$, whenever the set $C(r,p)$ of a node p in round r (where $r_{start} \leq r < r_{start} + \overrightarrow{P_\ell}$) holds multiple values, p randomly chooses to propose that round's lucky value, i.e., $v_\ell(r)$.

We now prove two observations that are useful for the proof for termination.

Observation 2. *Suppose r_{start} is lucky and consider round r, where $r_{start} \leq r < r_{start} + \overrightarrow{P_\ell}$. If $v_\ell(r) = v$, then all good nodes in round $(r+1)$ collect at least one message proposing v for round r.*

Proof. By contradiction. Assume $v_\ell(r) = v$ and that some good node in round $(r+1)$ does not collect v for round r.

Let $A(r)$ be the set of good nodes whose messages for round r are collected by all the good nodes in round $(r+1)$. By Corollary 1, $A(r) \neq \emptyset$. Since some good node does not collect v for round r, it follows that none of the good nodes

in $A(r)$ proposes v for round r, i.e. all the good nodes in $A(r)$ propose $v' \neq v$ for round r.

Since $v_\ell(r) = v$, for any node p, if $v \in C(r,p)$, then p must propose v for round r. Note that all the good nodes in $A(r)$ propose v' for round r, therefore, for any $p_g \in A(r)$, $C(r,p_g)$ only contains v'. That is, $O(r) = A(r) \neq \emptyset$. By definition of v_ℓ, $v_\ell(r)$ should be set to v'. Contradiction. □

Observation 3 (Necessary condition for v_ℓ flipping)**.** *If r_{start} is lucky, then for any r where $r_{start} < r < r_{start} + \overleftarrow{P_\ell}$, $v_\ell(r)$ is different from $v_\ell(r-1)$ only if some good node collects from round $(r-1)$ some message proposing $v_\ell(r)$ with priority at least 1.*

Proof. Assume $v_\ell(r-1) = v'$ and $v_\ell(r) = v$, where $v' \neq v$.

Since $r > r_{start}$ and $v' \neq v$, by definition of v_ℓ, $O(r) \neq \emptyset$ and for any $p \in O(r)$, $v \in C(r,p)$. Now we consider the values that one such good node, $p_g \in O(r)$, collects from round $(r-1)$. By Observation 2, p_g collects at least one message proposing v' from round $(r-1)$. However, only v is in $C(r,p_g)$. Therefore, p must have collected a message proposing v with a higher *priority* than v', that is, at least 1.

□

Observation 4. *If some good node in round r at T collects from round $(r-1)$ some message m proposing v with priority at least 1, then there exists a good node proposing v with uCounter larger than $3\mathcal{T}$ in round r_g, where $r - 3\mathcal{T} < r_g \leq r-1$.*

Proof. Consider the node p that at step $T' < T$ sends m, which proposes v with *priority* at least 1 for round $(r-1)$. By Corollary 3, p must have collected a message from a good node m_g by T' for round r_g, where $r - 3\mathcal{T} < r_g \leq r-1$. Since p sends m with *priority* 1, *i.e.* $uCounter_p \geq 6\mathcal{T}$; then, by Lemma 12, m_g must propose v' with $uCounter_g \geq uCounter_p - ((r-1) - r_g) \geq 6\mathcal{T} - ((r-1) - r_g) > 3\mathcal{T}$. □

Observation 5. *If round r_{start} is a lucky round, then all good nodes active at step $T_D(r_{start})$ have decided by $T_D(r_{start})$.*

Proof. Recall that $T(r_{lock})$ is the earliest step where some node is in round r_{lock}, and $r_{lock} = r_{start} + 6\mathcal{T}$. Let $v_\ell(r_{lock})$ be v.

The proof proceeds in two main steps. In Step 1, we will prove that:

For any r, where $r_{lock} < r \leq r_{start} + \overleftarrow{P_\ell} - 1$, $v_\ell(r) = v_\ell(r_{lock}) = v$, and all good nodes propose v for round r. (F7)

In Step 2, relying on F7, we are going to prove that, for any good node p_g, the *uCounter* of $v = v_\ell(r_{lock})$ at $T_D(r_{start})$ will be at least U^D, upon which p_g will decide v.

Step 1 To prove F7, we are going to prove:

> No good node in round r, where $r_{lock} \leq r \leq r_{start} + \overleftrightarrow{P_\ell} - 1$, collects a message proposing $v' \neq v$ for round $(r-1)$ with *priority* larger than 0. (F8)

Assuming F8 is true, then it is easy to show F7 is true as follows. By combining F8 and Observation 3, we can conclude that $v_\ell(r_{lock})$ is the lucky value for all rounds from r_{lock} to $(r_{start} + \overleftrightarrow{P_\ell} - 1)$. Now, consider any round r, where $r_{lock} + 1 \leq r \leq r_{start} + \overleftrightarrow{P_\ell} - 1$. By Observation 2, since $v_\ell(r-1) = v$, all good nodes in round r collect at least one message proposing v for round $(r-1)$. By F8, we know that no good node in round r collects a message proposing v' for round $(r-1)$ with *priority* larger than 0 . Therefore, any good node in round r either collects only v, or collects both v and v', where the *priority* of v' is 0. Since r_{start} is a lucky round, all good nodes propose v for round r, proving F7.

We are going to prove F8 by contradiction. Let r' be the earliest round in the range from r_{lock} to $(r_{start} + \overleftrightarrow{P_\ell} - 1)$, where some good node, currently in r', collects a message from round $(r'-1)$ proposing $v' \neq v$ with *priority* at least 1.

We are going to prove that (i) there exists a round r_g, where $(r_{lock} - 3\mathcal{T}) < r_g \leq (r'-1)$, such that (a) a good node p_g proposes v' in round r_g with $uCounter_g > 3\mathcal{T}$, and (b) $v_\ell(r_g - 1) = v'$; and (ii) $(r_g - 1)$ can be neither smaller nor larger than r_{lock}. Since $(r_g - 1) \neq r_{lock}$, as their v_ℓ values are different, this leads to a contradiction.

Proof of (i) Since some good node in round r' at T collects a message proposing v' for round $(r'-1)$ with *priority* at least 1, then, by Observation 4, there exists a good node p_g proposing v' in round r_g, where $r' - 3\mathcal{T} < r_g \leq r' - 1$ with $uCounter_g > 3\mathcal{T}$. Now with $r_g > (r' - 3\mathcal{T}) \geq (r_{lock} - 3\mathcal{T})$ (by definition of r'), we have $(r_{lock} - 3\mathcal{T}) < r_g \leq (r'-1)$, establishing (a).

Now, to establish (b), we show that $v_\ell(r_g - 1) = v'$. Since $uCounter_g > 0$, by line 17 of Sandglass, p_g collects only v' in round $(r_g - 1)$. Note that, by Observation 2, all good nodes in round r_g, including p_g, collect at least one message proposing $v_\ell(r_g - 1)$ for round $(r_g - 1)$. Therefore, $v_\ell(r_g - 1)$ must be equal to v', *i.e.*, $v_\ell(r_g - 1) = v'$.

Having proved (a) and (b), we proved (i).

Proof of (ii) Consider the round r_g that exists by (i). We know that $(r_g - 1) \neq r_{lock}$. We now show that $(r_g - 1)$ can be neither smaller nor larger than r_{lock}, which leads to a contradiction.

Case 1 $(r_g - 1) < r_{lock}$

We are going to show that, under the assumption of $(r_g - 1) < r_{lock}$, it is possible to prove two statements, S1 and S2, that are in contradiction with each other.

S1: *There exists a round r, where $r_g - 3\mathcal{T} < r \leq r_{\text{lock}} - 1$, in which a good node proposes v with $uCounter > 3\mathcal{T}$.*
Since $v_\ell(r_g - 1) = v'$ and $v_\ell(r_{lock}) = v$, there must exist a round r_c between r_g and r_{lock} where v_ℓ changes from v' to v, *i.e.* $v_\ell(r_c - 1) = v'$ and $v_\ell(r_c) = v$.
By Observation 3, some good node in round r_c collects from round $(r_c - 1)$ some message proposing v with *priority* at least 1. Then, by Observation 4, there exists a good node p_v in round r_v, where $r_c - 3\mathcal{T} < r_v \leq r_c - 1$, proposing v with $uCounter > 3\mathcal{T}$. Since $r_g \leq r_c \leq r_{lock}$, we have proved S1: there exists a round r_v, where $r_g - 3\mathcal{T} < r_v \leq r_{lock} - 1$, such that a good node p_v in r_v proposes v with $uCounter > 3\mathcal{T}$.

S2: *No good node in any round r, where $r_g - 3\mathcal{T} < r \leq r_{\text{lock}} - 1$, proposes v with $uCounter > 3\mathcal{T}$.*
Recall that, by (i), p_g proposes v' with $uCounter_g \geq 3\mathcal{T} + 1$ in round r_g. Consider any round r between $r_g - 3\mathcal{T}$ and r_g. By Corollary 4, there exists a good node in round r proposing v with $uCounter$ at least $(uCounter_g - (r_g - r))$, which is a value greater than 0. Then, by Lemma 13, we can draw a first conclusion: no good node can propose v with $uCounter > 0$ for any round r, where $r_g - 3\mathcal{T} - 1 \leq r \leq r_g - 1$.
When r is equal to $(r_g - 1)$, this means that no good node proposes v with $uCounter > 0$ in round $(r_g - 1)$. Then, by Corollary 4, we can further infer that no good node proposes v with $uCounter > 3\mathcal{T}$ in any round between r_g and $(r_g + 3\mathcal{T})$. Since $r_g \geq (r_{lock} - 3\mathcal{T})$, *i.e.*, $(r_{lock} - 1) \leq r_g + 3\mathcal{T} - 1$, we can draw a second conclusion: for any round r, where $r_g \leq r \leq r_{lock} - 1$, no good node proposes v with $uCounter > 3\mathcal{T}$.
Combining our two conclusions, we have that no good node can propose v with $uCounter > 3\mathcal{T}$ in any round r, where $r_g - 3\mathcal{T} < r \leq r_{lock} - 1$, proving S2.

Since S1 and S2 contradict each other, and we were able to prove them under the assumption that $r_g - 1 < r_{lock}$, we conclude that Case 1 is impossible.

Case 2 $(r_g - 1) > r_{lock}$
Since $v_\ell(r_{lock}) = v$ and we proved that $v_\ell(r_g - 1) = v'$, then in some round r_c, where $r_{lock} < r_c \leq (r_g - 1)$, $v_\ell(r_c - 1) = v$ and $v_\ell(r_c) = v'$. By Observation 3, some good node in round r_c must collect a message proposing v' with *priority* at least 1 from

round $(r_c - 1)$.
Recall that r' is the earliest round in the range from r_{lock} to $(r_{start} + \overleftarrow{P_\ell} - 1)$, where some good node, currently in r', collects a message from round $(r' - 1)$ proposing $v' \neq v$ with *priority* at least 1; and that $r_g \leq r' - 1$. Therefore, $r_{lock} < r_c < r_g < r'$.
However, by assumption, r' is the earliest round in which some node collects a message proposing v' with *priority* at least 1. Contradiction.

This concludes the proof that F8 holds. Recall that, as we showed above, F8 implies F7:

For any r, where $r_{lock} < r \leq r_{start} + \overleftarrow{P_\ell} - 1$, $v_\ell(r) = v_\ell(r_{lock}) = v$, and all good nodes propose v for round r. (F7)

which is now also proved.

Step 2 Now, we are going to show that, for any good node p_g that is active at $T_D(r_{start})$, the *uCounter* of $v = v_\ell(r_{lock})$ at $T_D(r_{start})$ will be at least U^D. This is the condition upon which p_g will decide v.

The key technical hurdle we need to clear is to prove following fact:

A node that proposes $v' \neq v$ before $T_D(r_{start})$ can be at most in round $(r_{start} + \overleftarrow{P_\ell} - U^D - 1)$. (F9)

Assuming F9 holds, it follows easily that all good nodes that are active at $T_D(r_{start})$ must have decided by $T_D(r_{start})$. Here is why.

Since by F9 all nodes that propose in round $(r_{start} + \overleftarrow{P_\ell} - U^D)$ before $T_D(r_{start})$ must propose v, then, by line 17 of Sandglass, all nodes that propose in round $(r_{start} + \overleftarrow{P_\ell} - U^D + 1)$ before $T_D(r_{start})$ must propose v with *uCounter* at least 1. A simple inductive argument then shows that all nodes that ever propose in round $(r_{start} + \overleftarrow{P_\ell} - U^D + i)$ before $T_D(r_{start})$, where $1 \leq i < U^D$, propose v with *uCounter* at least i. With $i = U^D - 1$, messages sent for round $(r_{start} + \overleftarrow{P_\ell} - 1)$ before $T_D(r_{start})$ must propose v with *uCounter* at least $(U^D - 1)$. Note that by Lemma 3 and because p_g is active at step $T_D(r_{start})$, p_g enters round $(r_{start} + \overleftarrow{P_\ell})$ either at step $(T_D(r_{start}) - 1)$ or at step $T_D(r_{start})$. In both cases, p_g proposes v with *uCounter* at least U^D, *i.e.*, with *priority* at least $(6\mathcal{T} + 4)$, and decides by lines 21-22 of Sandglass. Therefore, if p_g is active at step $T_D(r_{start})$, it must have decided by step $T_D(r_{start})$.

To prove F9, we use again the notion of *bags* that we introduced in the proof for Agreement. We quickly review it below.

For each message m sent in round r, m's *bag* for round $(r-1)$ is the set of messages collected by the sender of m in round $(r-1)$.

Recall that, if some node p' sends a message $m_{r'}^{v'}$ proposing v' for round r', then there exists a chain of messages extending from round 1 to round r', where (a) each message on the chain proposes v', and (b) the i-th message on the chain was one of the highest *priority* messages collected from round i by the sender of the $(i+1)$-th message.

To each message in the chain corresponds a bag: by definition, the bag of the chain's i-th message is the bag for round $(i-1)$. Thus, in the chain there exists exactly one bag per round, and at least one of the messages with the highest *priority* in each bag must be proposing v'.

Let $uCounter_i^{v'}$ be the value of *uCounter* of the i-th message on the chain. By line 17 of Sandglass, $\forall i : 2 \leq i < r'$: $uCounter_i^{v'} \geq uCounter_{i+1}^{v'} - 1$. Therefore, as we saw, the following holds:

$$\forall i : 2 \leq i < j \leq r' \text{: } uCounter_i^{v'} \geq uCounter_j^{v'} - (j-i). \qquad \text{(F3)}$$

We are now ready to prove F9. We proceed by contradiction.

Assume there exists a node p' that before $T_D(r_{start})$ uses a message m' to propose v' in some round $r' > (r_{start} + \overleftarrow{P_\ell} - U^D - 1)$. Consider the chain of messages associated with m' and, in particular, the bags for every round from $(r_{lock}+1)$ to $(r_{start} + \overleftarrow{P_\ell} - U^D - 1)$.

We will show that:

- (i) Among the bags for rounds from $(r_{lock}+1)$ to $(r_{start}+\overleftarrow{P_\ell}-U^D-1)$, at most one in every U^1 bags can contain messages from good nodes. That is, among these $(r_{start} + \overleftarrow{P_\ell} - U^D - r_{lock} - 2)$ bags, $(r_{start} + \overleftarrow{P_\ell} - U^D - r_{lock} - 2 - \lceil \frac{r_{start}+\overleftarrow{P_\ell}-U^D-r_{lock}-2}{U^1} \rceil) = (\overleftarrow{P_\ell} - 6\mathcal{T} - U^D - 2 - \lceil \frac{\overleftarrow{P_\ell}-6\mathcal{T}-U^D-2}{U^1} \rceil)$ of them contain only messages from defective nodes. We will call these *Def*-bags.
- (ii) The number of messages sent by defective nodes for round $(r_{lock}+1)$ to $(r_{start} + \overleftarrow{P_\ell} - U^D - 1)$ before $T_D(r_{start})$ is not sufficient to fill all *Def*-bags.

Thus, it is impossible for a node that before T proposes $v' \neq v_d$ to advance up to round $(r_{start} + \overleftarrow{P_\ell} - U^D - 1)$, contradicting our assumption.

Proof of (i) We will show that among the bags for rounds $(r_{lock}+1)$ to $(r_{start} + \overleftarrow{P_\ell} - U^D - 1)$, at most one in every U^1 bags contains messages from good nodes.

By contradiction: assume there exist two bags containing messages from good nodes, for round r_1 and r_2 respectively, where $(r_{lock}+1) \leq r_1 < r_2 \leq (r_{start}+\overleftarrow{P_\ell}-U^D-1)$, and $r_2 - r_1 < U^1$. Consider now any message $m_{r_1}^g$ from a good node contained in the bag for round r_1.

Since r_1 is within the lucky period that begins in r_{start}, by F7, $m^g_{r_1}$ proposes v. Similarly, any message $m^g_{r_2}$ from a good node contained in the bag for round r_2 proposes v.

Now, let us consider messages $m^{v'}_{r_1+1}$ and $m^{v'}_{r_2}$ on the chain. We showed above (F3) that $uCounter^{v'}_{r_1+1} \geq uCounter^{v'}_{r_2} - (r_2 - (r_1 + 1))$. Since $m^g_{r_1}$ proposing $v \neq v'$ is in the bag of round r_1, by line 17 of Sandglass, $uCounter^{v'}_{r_1+1} = 0$. Therefore, $uCounter^{v'}_{r_2} \leq uCounter^{v'}_{r_1+1} + (r_2 - (r_1 + 1)) = r_2 - r_1 - 1 < U^1$. Then, by line 20 of Sandglass, $priority^{v'}_{r_2} = 0$.

However, recall that $m^{v'}_{r_2}$ is one of the messages with the largest *priority* among all messages in $m^{v'}_{r_2+1}$'s bag. Therefore, no message collected by $m^{v'}_{r_2+1}$ from round r_2 proposes v' with *priority* greater than 0. Note that $m^{v'}_{r_2+1}$ also collects $m^g_{r_2}$, which proposes v. Therefore, consider the set of values with the highest *priority* that $m^{v'}_{r_2+1}$ collected from round r_2. Either that set contains only v, when some v is proposed with *priority* greater than 0; or it contains both v and v', when both values are proposed with *priority* equal to 0. In either case, since $(r_2 + 1) \leq (r_{start} + \overrightarrow{P_\ell} - U^D) < (r_{start} + \overrightarrow{P_\ell})$ is within the lucky period, and $v_\ell(r_2 + 1) = v_\ell(r_{lock}) = v$, $m^{v'}_{r_2+1}$ must propose v. However, by construction $m^{v'}_{r_2+1}$ should propose v'. Contradiction.

Proof of (ii) We have just established that, among the bags for rounds from $(r_{lock}+1)$ to $(r_{start}+\overrightarrow{P_\ell}-U^D-1)$, $(r_{start}+\overrightarrow{P_\ell}-U^D-r_{lock}-2-\lceil\frac{r_{start}+\overrightarrow{P_\ell}-U^D-r_{lock}-2}{U^1}\rceil)$ *Def*-bags contain only messages from defective nodes.

Since each bag contains at least $\mathcal{T}$ messages, *Def*-bags contain at least $\mathcal{T} \cdot (\overrightarrow{P_\ell} - 6\mathcal{T} - U^D - 2 - \lceil\frac{\overrightarrow{P_\ell}-6\mathcal{T}-U^D-2}{U^1}\rceil)$ messages.

Let us now compute the largest number of messages these *Def*-bags can contain.

Def-bags can only contain messages sent for round $(r_{lock} + 1)$ to $(r_{start}+\overrightarrow{P_\ell}-U^D-1)$ by defective nodes. By Lemma 5, $T^\forall_{r_{lock}}$ is the earliest step where some defective node can be in round $(r_{lock} + 1)$. Then, the messages in *Def*-bags must have been sent from step $T^\forall_{r_{lock}}$ to step $(T_D(r_{start}) - 1)$. By Lemma 10, the number of messages sent by defective nodes in the time interval from $T^\forall_r$ to $(T^\forall_{r+1} - 1)$ is at most $(\mathcal{T} - 1)$. Since $T_D(r_{start}) = T^\forall_{r_{start}+\overrightarrow{P_\ell}}$, the period from step $T^\forall_{r_{lock}}$ to step $(T_D(r_{start}) - 1)$ covers $(r_{start} + \overrightarrow{P_\ell} - r_{lock}) = (\overrightarrow{P_\ell} - 6\mathcal{T})$ such intervals; thus, the number of messages sent by defective nodes in this period is at most $(\mathcal{T} - 1) \cdot (\overrightarrow{P_\ell} - 6\mathcal{T})$.

Recall that *Def*-bags contain at least $\mathcal{T} \cdot (\overrightarrow{P_\ell} - 6\mathcal{T} - U^D - 2 -$

$\lceil \frac{\overrightarrow{P_\ell}-6\mathcal{T}-U^D-2}{U^1} \rceil$) messages. Therefore, we have

$$
\begin{aligned}
&(\mathcal{T}-1)\cdot(\overrightarrow{P_\ell}-6\mathcal{T}) \geq \mathcal{T}\cdot(\overrightarrow{P_\ell}-6\mathcal{T}-U^D-2-\lceil\frac{\overrightarrow{P_\ell}-6\mathcal{T}-U^D-2}{U^1}\rceil)\\
\Rightarrow &(\mathcal{T}-1)\cdot C \geq \mathcal{T}\cdot(C-U^D-2-\lceil\frac{C-U^D-2}{U^1}\rceil)\\
&\qquad\qquad \text{(where C} = \lceil\tfrac{(6\mathcal{T}-1)\cdot U^D+18\mathcal{T}}{5}\rceil\text{, and } \overrightarrow{P_\ell}=6\mathcal{T}+C)\\
\Rightarrow &U^D+2+\lceil\frac{C-U^D-2}{U^1}\rceil \geq \frac{C}{\mathcal{T}}\\
\Rightarrow &U^D+3+\frac{C-U^D-2}{U^1} > \frac{C}{\mathcal{T}}\\
&\qquad\qquad \text{(Since } (\tfrac{C-U^D-2}{U^1}+1) > \lceil\tfrac{C-U^D-2}{U^1}\rceil)\\
\Rightarrow &U^D+3 > \frac{5C+U^D+2}{6\mathcal{T}} \qquad\qquad \text{(Since } U^1=6\mathcal{T})\\
\Rightarrow &C < \frac{6\mathcal{T}(U^D+3)-U^D-2}{5} = \frac{(6\mathcal{T}-1)U^D+18\mathcal{T}-2}{5}
\end{aligned}
$$

However, since $C = \lceil\frac{(6\mathcal{T}-1)\cdot U^D+18\mathcal{T}}{5}\rceil > \frac{(6\mathcal{T}-1)U^D+18\mathcal{T}-2}{5}$, we have a contradiction. □

Lemma 15 (Termination with probability 1)**.** *Every good node that remains active decides with probability 1.*

Proof. By Observation 5, for any round r_{start}, if r_{start} is a lucky round, then all good nodes that are active in step $T_D(r_{start})$ decide a value in round $r_{start} + 6\mathcal{T} + \lceil\frac{(6\mathcal{T}-1)\cdot U^D+18\mathcal{T}}{5}\rceil$. Let $\overrightarrow{P_\ell} = 6\mathcal{T} + \lceil\frac{(6\mathcal{T}-1)\cdot U^D+18\mathcal{T}}{5}\rceil$.

Let $S = \{1 + k\cdot(\overrightarrow{P_\ell}+1) | k \in N\}$. S is a set containing infinitely many numbers of rounds, that the events of each of them being lucky are mutually independent.

Let $T^\forall(r)$ be the earliest step where all good nodes are in round r. For any round r, $T_1(r)$, which is earliest step that any good node can be in round r, is no earlier than $T^\forall(r-1)$ by Lemma 5; and $T_D(r)$ is defined as $T^\forall(r+\overrightarrow{P_\ell})$. Consider the period consisting of the steps from $T_1(r)$ to $T_D(r)-1$, i.e., from $T^\forall(r-1)$ to $T^\forall(r+\overrightarrow{P_\ell})-1$; this period covers all rounds from r to $(r+\overrightarrow{P_\ell}-1)$. For this period to be lucky and for r to be a lucky round, we require any node that must randomly select the value it will propose in any round between r and $r+\overrightarrow{P_\ell}-1$ to select the lucky value v_ℓ for its current round.

Consider the events that correspond to rounds in S being lucky. Since the lucky periods for rounds in S are not overlapping, these events are mutually independent.

Now we are going to show that all rounds in S are lucky with non-zero probability. Consider round r in S. In each step of r's lucky period, there are at most N nodes in the system. Each node that makes a random choice in one of the rounds covered by the lucky period chooses the round's lucky value

with probability $\frac{1}{2}$. Therefore, in every step of the lucky period, the probability that all nodes that make random choices select the lucky value for their current round is at least $\frac{1}{2^N}$. By Lemma 4, it takes at most $\mathcal{T} \cdot (\overrightarrow{P_\ell} + 1)$ steps from $T_1(r)$ to $(T_D(r) - 1)$. Therefore, the probability that any round in S is lucky is at least $\frac{1}{2^{N \cdot \mathcal{T} \cdot (\overrightarrow{P_\ell}+1)}} > 0$.

Now, consider any good node p_g that joins in round r_g at any step T and stays active. Recall that, by Lemma 4, good nodes are guaranteed to eventually reach any arbitrary round. Since there are infinitely many rounds r in S where $T_D(r) > T$, with probability 1 there exists a round $r \in S$ such that (1) r is a lucky round; and (2) $T_D(r) \geq T$. Then, by Lemma 4, p_g will eventually reach $T_D(r)$ and, by Observation 5, decide. □