

Efficient Plurality Consensus, or: The benefits of cleaning up from time to time

Petra Berenbrink¹,
Tom Friedetzky², George Giakkoupis³, and Peter Kling^{*1}

1 Simon Fraser University, Burnaby, Canada

2 Durham University, Durham, U.K.

3 INRIA, Rennes, France

Abstract

Plurality consensus considers a network of n nodes, each having one of k opinions. Nodes execute a (randomized) distributed protocol with the goal that all nodes adopt the *plurality* (the opinion initially supported by the most nodes). Communication is realized via the GOSSIP (or random phone call) model. A major open question has been whether there is a protocol for the complete graph that converges (w.h.p.) in polylogarithmic time and uses only polylogarithmic memory per node (local memory). We answer this question affirmatively.

We propose two protocols that need only mild assumptions on the bias in favor of the plurality. As an example of our results, consider the complete graph and an arbitrarily small constant multiplicative bias in favor of the plurality. Our first protocol achieves plurality consensus in $O(\log k \cdot \log \log n)$ rounds using $\log k + \Theta(\log \log k)$ bits of local memory. Our second protocol achieves plurality consensus in $O(\log n \cdot \log \log n)$ rounds using only $\log k + 4$ bits of local memory. This disproves a conjecture by Becchetti et al. (SODA'15) implying that any protocol with local memory $\log k + O(1)$ has worst-case runtime $\Omega(k)$. We provide similar bounds for much weaker bias assumptions. At the heart of our protocols lies an *undecided state*, an idea introduced by Angluin et al. (Distributed Computing'08).

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases plurality consensus, voting, majority, distributed, gossip

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Reaching *plurality consensus* is a fundamental problem in distributed computing. We consider this problem in a networked setting, where a graph is given in which each node initially holds one of $k \in \mathbb{N}$ *opinions*. The objective is the design of an efficient distributed protocol that ensures that eventually all nodes agree on the initial *plurality opinion*, which is the opinion that is initially supported by the most nodes. This problem is also referred to as *majority consensus* or *proportionate agreement* [1, 3, 19]. In accordance with [5] and others, we prefer to refer to it as *plurality consensus*, so as to make clear that the opinion eventually to be attained by all nodes need not initially have been absolute majority. One need not stray too far from the core of distributed computing to come across direct applications of plurality consensus: the handling of fault tolerance in parallel computing or the implementation of majority-based conflict resolution for CRCW PRAMs (and derivative models) are immediate examples.

Natural metrics for plurality consensus protocols are running time and memory overhead. The latter is the additional amount of (local) memory needed by each node above and beyond the bits

* pkling@sfu.ca. Supported in part by the Pacific Institute for the Mathematical Sciences.



required to store its current opinion. Results are typically expressed in terms of the number of nodes, the number of initial opinions, and the initial bias (between plurality and remaining opinions). Of particular interest, to us and in general, are fast protocols with small memory overhead. As in [5] we assume a synchronous PULL-based GOSSIP communication model on the underlying graph. Here, in each discrete round, every node may contact one neighbor and query that neighbor's opinion. We mostly assume the complete graph, but using a construction of [5] we can easily extend our results to regular expanders. We define the (relative) plurality gap $\gamma > 1$ as the ratio between the plurality opinion and the second most common opinion. We analyze two protocols in this paper:

- A protocol with running time $O(\log k \cdot \log \log_\gamma n + \log \log n)$ and memory overhead $\Theta(\log \log k)$.
- A protocol with running time $O(\log n \cdot \log \log_\gamma n)$ and memory overhead 4.

Plurality consensus is a member of the class of population dynamics, which are of great interest in fields as varied as epidemiology, physics, statistics, biology, chemistry, or sociology. All these have in common an initial population of agents with some initial properties and a protocol (dynamics) that in some manner changes the properties of given agents usually based on those of other agents. Specific models are as varied as the problems themselves; for instance, we may or may not have an underlying graph structure, a prescribed timing model, or restrictions on the amount and nature of communication. Other related dynamics are the lately en vogue voting protocols and Moran-type processes.

1.1 Related Work

In [3], Aspnes et al. consider $k=2$ initial opinions $\{x, y\}$ and the complete graph as neighborhood structure. They introduce a third state, referred to as *blank*, b , which is a crucial ingredient in their protocol and analysis. Their protocol works such that an activated node u picks another node v at random. Given the opinions of those two nodes, the transition now proceeds as follows: If u has a non-blank opinion and sees in v the other non-blank opinion then it changes to b , if it has a non-blank opinion and sees in v the same non-blank opinion or b then it maintains its opinion, and if it has the blank opinion then it just copies whatever it sees in v . The authors show that with high probability all n nodes reach consensus within $O(n \log n)$ many interactions (corresponding to parallel convergence time $O(\log n)$), and the consensus value is the plurality value provided its (absolute) initial bias is at least $\omega(\sqrt{n \log n})$. Each node needs to be able to store one of three values, x , y , or b .

In the case of two opinions the plurality problem can be solved by calculating the median of the opinions. In [11] the authors present such a protocol that converges in $O(\log n)$ rounds and has constant memory overhead, if the initial difference bias $c_1 - c_2$ is at least $\Omega(\sqrt{n \log n})$. In [2] the authors consider the plurality problem in a sequential setting where only one node can change its opinion at a time. They present a new protocol called Average and Conquer (AVC) that solves plurality exactly, in sequential time $O(n \log n / (s\epsilon) + \log n \log s)$, where ϵn ($\epsilon > 1/2$) is the size of the plurality opinion and s the number of states. In [7] the authors generalize the former result to general networks and k opinions. They introduce protocols that solve the plurality consensus problem that are based on an interesting relationship between plurality consensus and distributed load balancing.

In [5] Becchetti et al. take the model of [3] and generalise it to $k \geq 2$ initial opinions. They still use the blank state, which they however refer to as undecided. The authors express their results in terms of $\text{md}(\bar{c})$, the *monochromatic distance* of configuration c . Formally, $\text{md}(\bar{c}) = \sum_{i=1}^k c_i / c_1$, where c_i is the number of nodes with opinion i and $c_1 \geq c_2 \geq \dots \geq c_k$. Note that $\text{md}(\bar{c})$ is always between $O(1)$ and $O(k)$. Then authors show almost-tight bounds on convergence time. Formally, let $k = k(n)$ be any function such that $k = O((n/\log n)^{1/3})$, and consider any initial configuration with $c_1 \geq (1 + \alpha) \cdot c_2$, where $\alpha \geq 0$ is any arbitrarily-small constant. Their protocol converges in $O(\text{md}(\bar{c}) \cdot \log n)$, rounds (w.h.p.) and it has only constant memory overhead. They also show

that for $k = O((n/\log n)^{1/6})$ and any initial configuration the convergence time of their protocol is (w.h.p.) linear in the monochromatic distance. Finally, they show how to adapt their results to regular expanders using random walks to sample the opinion of nodes. As with [3], they require one state more than is necessary to store the actual opinion values. They conjecture that any protocol using $\log k + O(1)$ bits of memory has runtime at least linear in k in the worst case — as discussed later, we partially refute this conjecture. In [6] the authors consider the 3-majority dynamics. They show that for $k \leq n^\alpha$ (with constant α) the 3-majority dynamics converges to an *almost-consensus* state in time $O((k^2 \sqrt{\log n} + k \log n)(k + \log n))$. An almost-consensus state is defined as a state where all but a subset of size $O(n^\gamma)$ (for constant $\gamma < 1$) of the nodes support the same opinion. In [4] the authors consider the undecided dynamics in complete graphs in an asynchronous setting. They derive the time of convergence and an upper bound for the probability of error.

A line of research which is related to the plurality consensus problem is the voting problem. The setting is the same, a network with n nodes is given and initially every node has one of k opinions. Here the goal is that all nodes agree on one opinion, which is not necessarily the plurality opinion. A sequential version of the voter model was introduced in [16]. The parallel voter model was first analyzed in [15]. The authors of [15] bound the expected consensus time in terms of the expected meeting time T_m of two random walks and show a bound of $O(T_m \cdot \log n) = O(n^3 \log n)$. The authors of [8] provide an improved upper bound of $O(1/(1 - \lambda_2) \cdot \log^4 n + \rho)$ on the expected consensus time for any graph G , where λ_2 is the second eigenvalue of the transition matrix of a random walk on G , and $\rho = (\sum_{u \in V(G)} d(u))^2 / \sum_{u \in V(G)} d^2(u)$ is the ratio of the square of the sum of node degrees over the sum of the squared degrees. The authors of [9, 10] consider a modification of the standard voter model with two opinions, which they call *two-sample voting*. In every round, each node chooses two of its neighbors randomly and adopts their opinion only if they both agree. For regular graphs and random regular graphs, it is shown that two-sample voting has a consensus time of $O(\log n)$ if the initial imbalance between the nodes having the two opinions is large enough. In [13] the authors consider a 2-choice voting protocol for k opinions in the complete graph. Their protocol converges to the majority opinion in time $O(k \cdot \log n)$, with high probability, if $k = O(n^\epsilon)$ for some small $\epsilon > 0$, and the initial absolute gap between largest and second-largest opinion is $\Omega(\sqrt{n \log n})$. They also show that there exist initial configurations where the $\Theta(k)$ bound on the run time is matched. Independently, they also give a protocol which is similar to our simple, first protocol (cf. Section 3) and has roughly the same voting time. Other related papers from literature about sensor networks include [12] (which considers *binary interval consensus*, which can be used to solve majorization) and [17] (which considers the plurality problem in a different distributed model and for constant k).

1.2 Our Contribution

In Section 3 we present and analyze protocol REPEATEDCLEANUP, which in a complete graph of n nodes and k opinions works in $O(\log k \cdot \log \log_\gamma n + \log \log n)$ rounds and has a memory overhead of $O(\log \log k)$ bits. The protocol is a synchronized version of the protocol from [6], in which the nodes “lose” their opinion and adopt the state “undecided” whenever they sample a node with a different opinion. Then they will adopt the “real” opinion of the next (i.e., not undecided) node they sample. The nodes keep changing their opinion until all nodes agree on the plurality opinion. In contrast, our protocol works in $O(\log \log_\gamma n)$ many phases, each of length $\Theta(\log k)$. In the first round of every phase the nodes sample a node and lose the opinion if the sampled node has a different opinion. They then use the rest of the phase to find a new “real” opinion. In [5] the authors show that for $k = O((n/\log n)^{1/6})$ and any initial configuration the convergence time of their protocol is linear in the monochromatic distance, which can be as large as $O(k)$. Hence, our protocol outperforms the lower bound for the protocol of [5]. Interestingly, the speed-up is reached by synchronizing the protocol, which can also be regarded as slowing it down. In our

protocol a node that just found a real opinion again waits until the beginning of the next round to sample another node, instead of doing that immediately.

The drawback of our first protocol is that it uses a counter to determine the end of a phase. In Section 4 we present a protocol that works in $O(\log n \cdot \log \log_\gamma n)$ rounds and has a memory overhead of only 4 bits. The main idea of the protocol is to slow down the progress in the individual phases by having nodes toss a biased coin with success probability $1/n$, basically replacing the deterministic counter by a probabilistic counterpart. Our result shows that a conjecture by Becchetti et al. [5], implying that any protocol with constant memory overhead has worst-case runtime $\Omega(k)$, does not hold if nodes have access to such a coin. Note that the coin toss is not necessary if nodes can decide whether they sampled themselves (or a marked node/leader).

A very recent, independent result by Ghaffari and Parter [14] suggests a protocol for plurality consensus with similar time and memory bounds as ours. They employ the same basic idea of cleanup and decision-accumulation rounds (cf. Section 3), which they name selection and recovery steps. Their final protocol differs in that they use some of the undecided nodes as *clock nodes* (which use the $\log k$ bits normally used to store the opinion to count time) to help synchronize other nodes.

2 Model & Notation

We consider protocols in the complete graph with $n \in \mathbb{N}$ nodes. Each node u has one of $k \in \mathbb{N}$ opinions $\text{op}_u \in \{1, 2, \dots, k\}$. We write $\text{op}_u = \perp$ to indicate that node u is undecided. Time is modelled in synchronous, discrete and parallel rounds and we assume a PULL-based GOSSIP model for communication (nodes can request information from one other node chosen uniformly at random). Note that each node needs at least $\lceil \log k \rceil$ bits of local memory (to store its current opinion). Any additional number of bits per node needed by a given protocol is called the protocol's *memory overhead*.

Notation. In the following, $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the L^1 and L^2 norms, respectively, that is, $\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$ and $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^n |x_i|^2}$ for an n -dimensional vector \mathbf{x} . For a real value $x > 0$, its binary logarithm is denoted by $\log x$ and its natural logarithm by $\ln x$. For an integer i , the shorthand $[i] := \{1, 2, \dots, i\}$ denotes the set of the first i integers. The phrase *with high probability* (w.h.p.) refers to probabilities of the form $1 - n^{-\Omega(1)}$.

At any point in time, the system can be described by a k -dimensional vector $\mathbf{x} = (x_1, x_2, \dots, x_k) \in \{0/n, 1/n, \dots, 1\}^k \subseteq [0, 1]^k$, where the i -th entry $x_i \in [0, 1]$ denotes the current fraction of nodes with opinion i . We call such a vector \mathbf{x} a *configuration*. Note that $1 - \|\mathbf{x}\|_1 \in [0, 1]$ is the fraction of undecided nodes. $X_i(t)$ is defined as the random variable that takes as its values the configuration produced by the protocol at the end of round $t \in \mathbb{N}_0$. We use $\bar{\mathbf{x}} := \mathbf{X}(0)$ to denote the initial (fixed) configuration. The random variable $Y(t) := 1 - \|\mathbf{X}(t)\|_1$ with values in $[0, 1]$ denotes the fraction of undecided nodes at the end of round t .

To measure how far we are from plurality consensus, we define two *plurality gap* notions: Assuming (w.l.o.g.) that 1 is a most common opinion, let $\psi(\mathbf{x}) := x_1 - \max_{i \neq 1} x_i \in [0, 1]$ be the *absolute plurality gap*. Similarly, $\gamma(\mathbf{x}) := x_1 / \max_{i \neq 1} x_i \geq 1$ is the *relative plurality gap*.

Assumptions. In order to guarantee convergence, we need some (mild) bias assumptions. Given the initial configuration $\bar{\mathbf{x}}$, without loss of generality (w.l.o.g.), we assume $\bar{x}_1 \geq \bar{x}_i$ for all $i \in [k]$ and refer to (the initially most common) opinion 1 as the *plurality opinion*. For most of the analysis we assume

```

1 sample a random node  $v$ 
2 if  $t \equiv 0 \pmod{T}$ : {cleanup}
3   if  $\text{op}_v \neq \text{op}_u$ :  $\text{op}_u \leftarrow \perp$ 
4 else: {decision-accumulation}
5   if  $\text{op}_u = \perp$ :  $\text{op}_u \leftarrow \text{op}_v$ 

```

■ **Listing 1** Protocol REPEATEDCLEANUP as executed by node u in round $t \in \mathbb{N}$. It works in phases of length $T = \Theta(\log k)$ and has memory overhead $\Theta(\log \log k)$.

$$\psi(\bar{\mathbf{x}}) = \omega\left(\frac{(\log n)^2}{\sqrt{n}}\right) \quad (1) \quad \text{and} \quad k = o\left(\frac{\sqrt{n}}{(\log n)^2}\right). \quad (2)$$

While Condition (1) is essential¹, Condition (2) is without loss of generality (it can be achieved by merging small opinions). We define $\rho = \rho(n) := (\log n)^2 / \sqrt{n}$ and call a configuration \mathbf{x} *biased* if $\psi(\mathbf{x}) = \omega(\rho)$. Our analysis assumes n to be at least some sufficiently large constant.

3 Plurality Consensus with $\log k + \Theta(\log \log k)$ Bits

We divide time into *phases*, each consisting of $T := 5 + 2 \log k$ rounds. In each round, every node u uses a PULL operation to sample a random node v and checks its opinion op_v . We distinguish two types of rounds:

- *Cleanup* rounds represent the first round of each phase. Here, u becomes undecided if op_v differs from u 's own opinion op_u (and keeps its opinion otherwise).
- *Decision-accumulation* rounds make up the remaining $T - 1$ rounds of a phase. Here, only undecided nodes act and simply adopt the pulled opinion op_v .

The synchronization of the steps in which nodes lose their opinion and adopt opinions of sampled nodes is key to the fast convergence. As we will see in the analysis they ensure that each phase increases the relative plurality gap exponentially. We call this protocol REPEATEDCLEANUP. See Listing 1 for a formal description. In the remainder of this section we prove the following theorem.

► **Theorem 1.** REPEATEDCLEANUP has a local memory overhead of $\Theta(\log \log k)$ bits. If started on a biased configuration $\bar{\mathbf{x}}$ with relative plurality gap $\gamma := \gamma(\bar{\mathbf{x}})$, then (w.h.p.) plurality consensus is achieved in $O(\log(k) \cdot \log \log_\gamma n + \log \log n)$ rounds.

We use the shorthand $\mathbf{X}(\tau, t) := \mathbf{X}((\tau - 1) \cdot T + t)$ to denote the configuration at the end of round t in phase τ . Similarly, $X_i(\tau, t)$ and $Y(\tau, t)$ denote the corresponding fractions of nodes. The following definition identifies opinions that are supported by only a small fraction of nodes (and, thus, are likely to vanish in a cleanup round), and configurations that do not have too many undecided nodes.

- An opinion i in configuration \mathbf{x} is *negligible* if $x_i \leq \rho$ (recall that $\rho = (\log n)^2 / \sqrt{n}$).
- A configuration \mathbf{x} is *alive* if $\|\mathbf{x}\|_1 \geq 1 - e^{-1}$.

The following simple observation already hints at the basic change of configurations during a phase: in expectation, the cleanup round squares the relative plurality gap (but reduces the

¹ Our results hold also under the slightly weaker condition $\psi(\bar{\mathbf{x}}) = \omega(\sqrt{\log(n)/n})$. However, Condition (1) significantly simplifies some parts of the analysis.

absolute support of each opinion), while the decision-accumulation rounds increase the absolute support of each opinion (not changing the relative plurality gap).

► **Observation 2.** Consider a configuration \mathbf{x} with $y := 1 - \|\mathbf{x}\|_1$. If the configuration \mathbf{X} results from a cleanup round on \mathbf{x} and $Y = 1 - \|\mathbf{X}\|_1$, then

$$\mathbb{E}[X_i] = x_i^2, \quad \mathbb{E}[\|\mathbf{X}\|_1] = \|\mathbf{x}\|_2^2, \quad \text{and} \quad \mathbb{E}[Y] = 2y - y^2. \quad (3)$$

Similarly, if \mathbf{X} results from a decision-accumulation round, then

$$\mathbb{E}[X_i] = (1+y) \cdot x_i, \quad \mathbb{E}[\|\mathbf{X}\|_1] = 2\|\mathbf{x}\|_1 - \|\mathbf{x}\|_2^2, \quad \text{and} \quad \mathbb{E}[Y] = y^2. \quad (4)$$

Analysis Overview. The bound on the memory overhead in Theorem 1 immediately follows from the protocol description. The runtime bound is proven in three steps: Lemma 7 shows that all non-plurality opinions become negligible during the first $O(\log \log_{\gamma(\bar{x})} n)$ phases. Lemma 8 proves that all these negligible opinions vanish within an additional constant number of phases. Lemma 9 shows that any remaining undecided nodes vanish in another $O(\log \log n)$ phases. Before we prove these key lemmas (Section 3.3), we show how configurations evolve during single rounds (Section 3.1) and single phases (Section 3.2).

3.1 Change During a Single Round

Our first two claims show concentration for the expected configuration change in cleanup and decision-accumulation rounds we saw in Observation 2. They follow by standard Chernoff bounds.

► **Claim 3** (Cleanup Round). Consider a configuration \mathbf{x} at the beginning of phase τ . Let $a > 0$ be a constant.

(a) Let $i \in [k]$ and $\delta := \sqrt{3a \cdot \log(n)/n} \cdot x_i^{-1}$. Then

$$\Pr[X_i(\tau, 1) \geq \max(\rho^2, (1+\delta) \cdot x_i^2) \mid \mathbf{X}(\tau, 0) = \mathbf{x}] \leq n^{-a} \quad \text{and} \quad (5)$$

$$\Pr[X_i(\tau, 1) \leq (1-\delta) \cdot x_i^2 \mid \mathbf{X}(\tau, 0) = \mathbf{x}] \leq n^{-a}. \quad (6)$$

(b) Let $\delta := \sqrt{3a \cdot \log(n)/n} \cdot \|\mathbf{x}\|_2^{-1}$. Then $\Pr[\|\mathbf{X}(\tau, 1)\|_1 \leq (1-\delta) \cdot \|\mathbf{x}\|_2^2 \mid \mathbf{X}(\tau, 0) = \mathbf{x}] \leq n^{-a}$.

► **Claim 4** (Decision-accumulation Round). Consider a configuration \mathbf{x} at the beginning of round $t+1 > 1$ in phase τ . Let $a > 0$ be a constant and $y := 1 - \|\mathbf{x}\|_1$ the fraction of undecided nodes.

(a) Let $i \in [k]$, $\delta' := \sqrt{13a \cdot \log n / (x_i y \cdot n)}$, and $\delta := \max(\delta', \delta'^2)$. Then

$$\Pr[X_i(\tau, t+1) \geq (1+\delta y) \cdot x_i \cdot (1+y) \mid \mathbf{X}(\tau, t) = \mathbf{x}] \leq n^{-a} \quad \text{and} \quad (7)$$

$$\Pr[X_i(\tau, t+1) \leq (1-\delta' y) \cdot x_i \cdot (1+y) \mid \mathbf{X}(\tau, t) = \mathbf{x}] \leq n^{-a}. \quad (8)$$

(b) Assume $y \leq 1 - \rho$. Then $\Pr[Y(\tau, t+1) \geq \max(\rho^2, y^{3/2}) \mid \mathbf{X}(\tau, t) = \mathbf{x}] \leq n^{-a}$.

3.2 Change During a Single Phase

Next, we use the effects of single rounds to show that (a) the property of being alive is (w.h.p.) invariant from phase to phase (Claim 5) and (b) the relative plurality gap of a biased configuration increases (w.h.p.) exponentially during a phase (Claim 6).

► **Claim 5.** Consider a configuration \mathbf{x} at the beginning of phase τ . Let $a > 0$ be a constant. Assume $\|\mathbf{x}\|_1 \geq 1 - e^{-1}$ and $\|\mathbf{x}\|_2^2 = \omega(\rho)$. Then

$$\Pr[\|\mathbf{X}(\tau+1,0)\|_1 < 1 - e^{-1} \mid \mathbf{X}(\tau,0) = \mathbf{x}] = n^{-a}. \quad (9)$$

Proof. Let us first consider the effect of the cleanup round. Claim 3(b) gives

$$\Pr[\|\mathbf{X}(\tau,1)\|_1 \leq \|\mathbf{x}\|_2^2/2 \mid \mathbf{X}(\tau,0) = \mathbf{x}] \leq n^{-a-1}, \quad (10)$$

where we used $\|\mathbf{x}\|_2^2 = \omega(\rho) = \omega(\sqrt{\log(n)/n})$, such that the involved δ -term becomes $o(1)$. Next, fix the configuration $\tilde{\mathbf{x}} = \mathbf{X}(\tau,1)$ after the cleanup round. Let $\tilde{y} := 1 - \|\tilde{\mathbf{x}}\|_1$ and assume $\tilde{y} \leq 1 - \|\mathbf{x}\|_2^2/2$ (this holds with probability at least $1 - n^{-a-1}$ due to Equation (10)). With the claim's assumption $\|\mathbf{x}\|_2^2 = \omega(\rho)$, this implies $\tilde{y} \leq 1 - \rho$ and we can apply Claim 4(b) to get

$$\Pr[Y(\tau,t+1) \geq \max(\rho^2, \tilde{y}^{3/2}) \mid \mathbf{X}(\tau,1) = \tilde{\mathbf{x}}] \leq n^{-a-1} \quad (11)$$

for a decision-accumulation round $t+1 > 1$. A union bound over the $T-1 = O(\log n)$ decision-accumulation rounds, combined with Equation (10) via the law of total probability, gives

$$\Pr\left[Y(\tau,T) \geq \max\left(\rho^2, \left(1 - \|\mathbf{x}\|_2^2/2\right)^{(3/2)^{T-1}}\right) \mid \mathbf{X}(\tau,0) = \mathbf{x}\right] \leq n^{-a}. \quad (12)$$

The second term of the maximum is at most $\exp(-(3/2)^{T-1} \cdot \|\mathbf{x}\|_2^2/2)$. For this to be at most some value $z > 0$, we need $T-1 \geq \log_{3/2}(2 \ln(1/z) / \|\mathbf{x}\|_2^2)$. Choosing $z := \exp(-k \cdot \|\mathbf{x}\|_2^2 / \|\mathbf{x}\|_1^2)$ and remembering our choice of T (see beginning of Section 3), we calculate

$$\log_{3/2}\left(\frac{2 \ln(1/z)}{\|\mathbf{x}\|_2^2}\right) = \log_{3/2}\left(2 / \|\mathbf{x}\|_1^2\right) + \log_{3/2}(k) \leq 4 + 2 \log(k) = T-1, \quad (13)$$

where we used the claim's assumption $\|\mathbf{x}\|_1 \geq 1 - e^{-1}$ to bound $\log_{3/2}(2 / \|\mathbf{x}\|_1^2) \leq 4$. Thus, Equation (12) implies that with probability at least $1 - n^{-a}$ we have $Y(\tau,T) \leq \max(\rho^2, z) \leq e^{-1}$, where we used $\rho^2 = o(1)$ and the Cauchy-Schwarz inequality to get $z = \exp(-k \cdot \|\mathbf{x}\|_2^2 / \|\mathbf{x}\|_1^2) \leq e^{-1}$. The claim follows with this from $\|\mathbf{X}(\tau+1,0)\|_1 = \|\mathbf{X}(\tau,T)\|_1 = 1 - Y(\tau,T)$. ◀

► **Claim 6.** Consider a configuration \mathbf{x} at the beginning of phase τ . Let $a > 0$ be a constant. For an opinion $i \neq 1$ assume $x_1 - x_i = \omega(\rho)$ and let $\gamma_i := \min(x_1/\rho, x_1/x_i)$. Then

$$\Pr\left[\frac{X_1(\tau+1,0)}{X_i(\tau+1,0)} \leq \gamma_i^{3/2} \mid \mathbf{X}(\tau,0) = \mathbf{x}\right] \leq n^{-a}. \quad (14)$$

Proof. For a configuration \mathbf{x}' define the shorthands $\gamma_i(\mathbf{x}') := x'_1/x'_i$ and $\psi_i(\mathbf{x}') := x'_1 - x'_i$. We make an inductive argument over γ_i during the phase (using ψ_i as an auxiliary tool for the induction). Applying Claim 3(a) to both opinions (once the upper and once the lower bound) yields

$$\Pr[X_1(\tau,1) \leq (1 - \delta_1) \cdot x_1^2 \mid \mathbf{X}(\tau,0) = \mathbf{x}] \leq n^{-a-1} \quad \text{and} \quad (15)$$

$$\Pr[X_i(\tau,1) \geq \max(\rho^2, (1 + \delta_i) \cdot x_i^2) \mid \mathbf{X}(\tau,0) = \mathbf{x}] \leq n^{-a-1}, \quad (16)$$

where $\delta_j := \sqrt{3a \cdot \log(n)/n} \cdot x_j^{-1}$ for $j \in \{1, i\}$. Let $\delta_\rho := \sqrt{3a \cdot \ln(n)/n} \cdot \rho^{-1}$ and $\delta := \min(\delta_i, \delta_\rho)$. Note that $\delta_1 \leq \delta$, since $x_1 \geq x_i + \omega(\rho)$ by the claim's assumption. With this, we bound the right-hand side in the probability of Equation (15) by $(1 - \delta) \cdot x_1^2$ and, similarly, the right-hand side in the probability of Equation (16) by $(1 + \delta) \cdot \max(\rho^2, x_i^2)$. Using a union bound and the inequality

$(1-x)/(1+x) \geq (1-2x)$ for $x \in [0,1]$, with probability at least $1-2n^{-a-1}$ we have $\gamma_i(\mathbf{X}(\tau,1)) \geq (1-2\delta) \cdot \gamma_i^2$. Similarly, we also have $\psi_i(\mathbf{X}(\tau,1)) = \omega(\rho^2)$. For $x_i \leq \rho$, this follows immediately from Equations (15) and (16) (since $x_1^2 = \omega(\rho^2)$ and $x_i^2 = O(\rho^2)$). For $x_i > \rho$, we calculate

$$\begin{aligned} \psi_i(\mathbf{X}(\tau,1)) &= X_1(\tau,1) - X_i(\tau,1) \geq (1-\delta_1) \cdot x_1^2 - (1+\delta_i) \cdot x_i^2 \\ &= \left(x_1^2 - x_i^2 - x_1 \cdot \sqrt{\frac{3a \cdot \log n}{n}} - x_i \cdot \sqrt{\frac{3a \cdot \log n}{n}} \right) \\ &= (x_1 + x_i) \cdot \left(x_1 - x_i - 2\sqrt{\frac{3a \cdot \log n}{n}} \right) = \omega(\rho) \cdot (\omega(\rho) - o(\rho)) = \omega(\rho^2). \end{aligned} \quad (17)$$

Next, fix the configuration $\tilde{\mathbf{x}} = X(\tau, t)$ and let $\tilde{y} := 1 - \|\tilde{\mathbf{x}}\|_1$. Assume $\psi_i(\tilde{\mathbf{x}}) = \omega(\rho^2)$ and $\tilde{x}_i \geq \rho^2$. Applying Claim 4(a) to both opinions (once the upper and once the lower bound) yields

$$\Pr[X_1(\tau, t+1) \leq (1 - \tilde{\delta}'_1 \tilde{y}) \cdot \tilde{x}_1 \cdot (1 + \tilde{y}) \mid \mathbf{X}(\tau, t) = \tilde{\mathbf{x}}] \leq n^{-a-1} \quad \text{and} \quad (18)$$

$$\Pr[X_i(\tau, t+1) \geq (1 + \tilde{\delta}'_i \tilde{y}) \cdot \tilde{x}_i \cdot (1 + \tilde{y}) \mid \mathbf{X}(\tau, t) = \tilde{\mathbf{x}}] \leq n^{-a-1}, \quad (19)$$

where $\tilde{\delta}'_j := \sqrt{13a \cdot \log n / (\tilde{x}_j \tilde{y} \cdot n)}$ and $\tilde{\delta}_j := \max(\tilde{\delta}'_j, \tilde{\delta}_j^2)$ for $j \in \{1, i\}$. Note that $\tilde{\delta}'_1 \leq \tilde{\delta}_i$, since $\tilde{x}_1 \geq \tilde{x}_i + \omega(\rho^2)$ by our assumption. Thus, as before we can combine Equations (18) and (19) via a union bound to get that with probability at least $1 - 2n^{-a-1}$ we have $\gamma_i(\mathbf{X}(\tau, t+1)) \geq (1 - 2\tilde{\delta}_i \tilde{y}) \cdot \gamma_i(\tilde{\mathbf{x}})$ and – analogous to the calculation in Equation (17) – $\psi_i(\mathbf{X}(\tau, t+1)) = \omega(\rho^2)$. Now, define the error $\tilde{\delta} := 2\tilde{\delta}_i \tilde{y}$ and note that

$$\tilde{\delta} = \max \left(\sqrt{\frac{13a \cdot \log n}{\tilde{x}_i n}} \cdot \sqrt{\tilde{y}}, \frac{13a \cdot \log n}{\tilde{x}_i n} \right) \leq \sqrt{\frac{13a \cdot \log n}{\tilde{x}_i n}} = O((\log n)^{-3/2}), \quad (20)$$

where we used the assumption $\tilde{x}_i \geq \rho^2 = \omega(\log(n)/n)$. In particular, this implies that (w.h.p.) we have $\gamma_i(\mathbf{X}(\tau, t+1)) \geq (1 - \tilde{\delta}) \cdot \gamma_i(\tilde{\mathbf{x}})$, $\psi_i(\mathbf{X}(\tau, t+1)) = \omega(\rho^2)$, and $X_i(\tau, t+1) \geq X_i(\tau, t) = \tilde{x}_i \geq \rho^2$. Moreover, the error $\tilde{\delta}$ is non-increasing in t (since \tilde{x}_i is non-decreasing in t), such that we get the largest error for $t=1$, such that we can apply the above recursively. Applying this via the chain rule to the $T-1 = O(\log n)$ decision accumulation rounds yields

$$\Pr \left[\gamma_i(\mathbf{X}(\tau, T)) \geq (1 - \tilde{\delta})^{T-1} \cdot \frac{\tilde{x}_1}{\max(\rho^2, \tilde{x}_i)} \mid \mathbf{X}(\tau, 1) = \tilde{\mathbf{x}} \right] \geq 1 - O(\log n) \cdot 2n^{-a-1}. \quad (21)$$

Since $\tilde{\delta} = O((\log n)^{-3/2})$, we can bound the error term $(1 - \tilde{\delta})^{T-1} \geq \exp(-(T-1) \cdot 2\tilde{\delta}) \geq 1 - 2T \cdot \tilde{\delta}$. That is, the error due to the decision-accumulation rounds is $1 - \delta'$ for $\delta' := 2T \cdot \tilde{\delta} = O((\log n)^{-1/2})$. We now combine our result for the cleanup round and the decision-accumulation rounds via the law of total probability to get that with probability at least $1 - n^{-a}$ we have $\gamma_i(\mathbf{X}(\tau, T)) \geq (1 - 2\delta)(1 - \delta') \cdot \gamma_i^2$.

Using $\check{\delta} := 2\max(2\delta, \delta') = O(\log(n) \cdot \sqrt{\log(n)/n}) \cdot \min(x_i^{-1}, \rho^{-1})$, we get $\gamma_i(\mathbf{X}(\tau, T)) \geq (1 - \check{\delta}) \cdot \gamma_i^2$. It merely remains to verify that $1 - \check{\delta} \geq \gamma_i^{-1/2}$. This is equivalent to $\check{\delta} \leq 1 - \gamma_i^{-1/2}$. If $\gamma_i = \Omega(1)$, this holds trivially since $\check{\delta} = o(1)$. So assume $\gamma_i = 1 + \varepsilon$ for a suitable $\varepsilon \in (0, 1]$. For this range, we have $1 - \gamma_i^{-1/2} \geq \varepsilon/4$, such that it is sufficient to show $\check{\delta} \leq \varepsilon/4$. By the claim's assumption $x_1 - x_i = \omega(\rho)$, we get $\varepsilon = \gamma_i - 1 = x_1/x_i - 1 = \omega(\rho) \cdot x_i^{-1}$. On the other hand, we have $\check{\delta} = O(\log(n) \cdot \sqrt{\log(n)/n}) \cdot \min(x_i^{-1}, \rho^{-1}) \leq O(\sqrt{(\log n)^3/n}) \cdot x_i^{-1} = O(\rho) \cdot x_i^{-1}$. This finishes the proof. \blacktriangleleft

3.3 Wrapping up the Analysis

We now have the tools to prove the three key lemmas mentioned before (which immediately imply Theorem 1). We first show that after $O(\log \log_{\gamma(\tilde{\mathbf{x}})} n)$ phases, (w.h.p.) all opinions $i \geq 2$ are negligible and the configuration is still alive.

► **Lemma 7.** Consider an initial configuration $\bar{\mathbf{x}}$ that is alive and for which $\psi(\bar{\mathbf{x}}) = \omega(\rho)$. Define $\tau_1 := \log_{3/2} \log_{\gamma(\bar{\mathbf{x}})} n$. Then $\Pr \left[\bigcap_{i=2}^k (X_i(\tau_1, 0) \leq \rho) \wedge \|\mathbf{X}(\tau_1, 0)\|_1 \geq 1 - e^{-1} \mid \mathbf{X}(1, 0) = \bar{\mathbf{x}} \right] \geq 1 - n^{-2}$.

Proof. Fix a phase $\tau \in \mathbb{N}$ and let \mathbf{x} denote the configuration at the beginning of phase τ . Assume \mathbf{x} to be alive, biased, and 1 to be the plurality opinion. Let $\gamma := \min(x_1/\rho, \gamma(\mathbf{x}))$. By Claim 5, with probability at least $1 - n^{-3}$ we have $\|\mathbf{X}(\tau+1, 0)\|_1 \geq 1 - e^{-1}$. Combined with Claim 6 via a union bound over all opinions we get

$$\Pr \left[\gamma(\mathbf{X}(\tau+1, 0)) \leq \gamma^{3/2} \vee \|\mathbf{X}(\tau+1, 0)\|_1 < 1 - e^{-1} \mid \mathbf{X}(\tau, 0) = \mathbf{x} \right] \leq 2n^{-3}. \quad (22)$$

Thus, as long as there is at least one non-negligible opinion $i \geq 2$, (w.h.p.) the relative plurality gap increases exponentially and the configuration stays alive. Moreover, note that $\mathbf{X}(\tau+1, 0)$ being alive and the increased relative gap between x_1 and $\max(\rho, x_i)$ for any other opinion i implies $\psi(\mathbf{X}(\tau+1, 0)) = \omega(\rho)$. Thus, we can iterate this argument. To this end, for any $\tau \in \mathbb{N}$ define the event

$$\mathcal{E}_{\tau+1} := \left(\gamma(\mathbf{X}(\tau+1, 0)) > \gamma_\tau^{3/2} \right) \wedge (\|\mathbf{X}(\tau+1, 0)\|_1 \geq 1 - e^{-1}) \wedge \psi(\mathbf{X}(\tau+1, 0)) = \omega(\rho), \quad (23)$$

where $\gamma_\tau := \min(X_1(\tau, 0)/\rho, \gamma(\mathbf{X}(\tau, 0)))$. Above we proved $\Pr[\mathcal{E}_{\tau+1} \mid \mathcal{E}_\tau, \mathbf{X}(\tau, 0) = \mathbf{x}] \geq 1 - 2n^{-3}$. Using the definition of conditional probability, we get $\Pr[\bigcap_{\tau \leq \tau_1} \mathcal{E}_\tau] \geq (1 - 2n^{-3})^{\tau_1} \geq 1 - n^{-2}$, where we used (by Assumption (1)) $\tau_1 \leq 2 \log \frac{\log n}{\log \gamma(\bar{\mathbf{x}})} \leq 2 \log \frac{\log n}{\log(1 + \psi(\bar{\mathbf{x}}))} \leq 2 \log \frac{2 \log n}{\psi(\bar{\mathbf{x}})} = o(\log n)$. Finally, our choice of τ_1 guarantees $\gamma(\bar{\mathbf{x}})^{(3/2)^{\tau_1}} = n \geq \rho^{-1}$, such that all opinions are negligible at the start of phase τ_1 . ◀

The next lemma shows that once we are in a configuration that is alive and where all opinions $i \geq 2$ are negligible, (w.h.p.) all these negligible opinions vanish within a constant number of additional phases.

► **Lemma 8.** Consider a configuration \mathbf{x} that is alive and for which $x_i \leq \rho$ for all $i \neq 1$. Define $\tau_2 := \tau_1 + 3$. Then $\Pr \left[\sum_{i=2}^k X_i(\tau_2, 0) = 0 \wedge \|\mathbf{X}(\tau_2, 0)\|_1 \geq 1 - e^{-1} \mid \mathbf{X}(\tau_1, 0) = \mathbf{x} \right] \geq 1 - n^{-2}$.

Proof. Applying Claim 6 yields $\Pr \left[\bigcup_{i \neq 1} \left(\frac{X_i(\tau_1+1, 0)}{X_i(\tau_1+1, 0)} \leq \gamma_\rho^{3/2} \right) \mid \mathbf{X}(\tau_1, 0) = \mathbf{x} \right] \leq n^{-3}$, where $\gamma_\rho = x_1/\rho$ as in Claim 6. Since $\|\mathbf{x}\|_1 \geq 1 - e^{-1}$ but $x_i \leq \rho$ for all $i \geq 2$, we must have $x_1 = \Omega(1)$ (or $\|\mathbf{x}\|_1 \leq k \cdot \rho + o(1) = o(1)$ would contradict \mathbf{x} being alive). In particular, we get $\gamma_\rho = \Omega(1/\rho)$. Combining this with Claim 5 via a union bound, with probability at least $1 - 2n^{-3}$ configuration $\mathbf{X}(\tau_1+1, 0)$ is alive and $X_i(\tau_1+1, 0) < \gamma_\rho^{-3/2} = O(\rho^{3/2})$ for all $i \geq 2$. Now consider the cleanup round of phase τ_1+1 for an opinion $i \geq 2$ with $X_i(\tau_1+1, 0) = O(\rho^{3/2})$. The probability that even one node of such an opinion remains decided after the cleanup round is at most $X_i(\tau_1+1, 0) \cdot n \cdot X_i(\tau_1+1, 0) = O(\rho^3 n)$. Repeating this for a constant number of phases (note that we can use our high probability bounds to guarantee that the configuration stays alive and the plurality gap high enough) and applying the geometric distribution, we get that the probability for an opinion $i \geq 2$ to survive c more phases is at most $O(\rho^{3c} n)$. The probability that even one of the $k-1 \leq \sqrt{n}$ negligible opinions survives these c phases is $O(\rho^{3c} \cdot n^{3/2}) = o(n^{-(c-1) \cdot 5/4})$. The claim's statement follows for $c=2$. ◀

Our last lemma shows that once we reached a configuration that is alive and only the plurality opinion is left, (w.h.p.) all nodes adopt the plurality in $O(\log \log n)$ phases.

► **Lemma 9.** For \mathbf{x} with $\sum_{i \geq 2} x_i = 0$ and $\|\mathbf{x}\|_1 = x_1 \geq 1 - e^{-1}$ let $\tau_3 := \tau_2 + \ln \log n + 1 = \tau_2 + O(\log \log n)$. Then $\Pr[X_1(\tau_3, 0) = 1 \mid \mathbf{X}(\tau_2, 0) = \mathbf{x}] \geq 1 - n^{-2}$.

Proof. We will show that the number of undecided nodes decreases exponentially from phase to phase, until their number is so low that (w.h.p.) they vanish within a constant number of decision-accumulation rounds. To this end, we use two basic high probability bounds: for cleanup rounds of a phase τ on any configuration \mathbf{x}' with $\sum_{i \geq 2} x'_i = 0$ and $\|\mathbf{x}'\|_1 = x'_1 \geq 2/3$, we use

$$\Pr[Y(\tau,1) \geq \max(\rho^2, 2.05 \cdot y') \mid \mathbf{X}(\tau,0) = \mathbf{x}'] \leq n^{-a}. \quad (24)$$

To see this, first note that² we have $\Pr[Y(\tau,1) \geq (1+\delta x'_1) \cdot y' \cdot (1+x'_1) \mid \mathbf{X}(\tau,0) = \mathbf{x}'] \leq n^{-a}$ for $\delta' := \sqrt{13a \cdot \log n / (x'_1 y' \cdot n)}$ and $\delta := \max(\delta', \delta'^2)$. Equation (24) then follows by distinguishing whether $y' \geq \rho^2/4$ or not.

To show the exponential decrease of the undecided nodes, assume we are given a configuration \mathbf{x}'' with $y'' < \max(\rho^2, 2.05 \cdot y')$ after the cleanup round of a phase τ . We apply Claim 4(b) to the first $t_* := 5 \leq T-1$ decision-accumulation rounds of phase τ and use a union bound to get $\Pr[Y(\tau,1+t_*) \geq \max(\rho^2, y''^{(3/2)^{t_*}}) \mid \mathbf{X}(\tau,1) = \mathbf{x}''] \leq t_* \cdot n^{-a}$. Since $y' \leq e^{-1}$, we get $(2.05 \cdot y')^{(3/2)^5} \leq y'^2$. Combining these observations with Equation (24) we get $\Pr[Y(\tau+1,0) \geq \max(\rho^2, y'^2) \mid \mathbf{X}(\tau,0) = \mathbf{x}'] \leq 6n^{-a}$. In particular, the resulting configuration is still alive at the start of phase $\tau+1$, so we can iterate to get

$$\Pr[Y(\tau_2+t,0) \geq \max(\rho^2, y'^{2^t}) \mid \mathbf{X}(\tau_2,0) = \mathbf{x}] \leq 6t \cdot n^{-a}. \quad (25)$$

Since $y \leq e^{-1}$, for $t = \ln \log n$ we get $\max(\rho^2, y'^{2^t}) = \rho^2$. Now, consider a configuration \mathbf{x}' with $y' < \rho^2$. Equation (24) yields $\Pr[Y(\tau_2+t,1) \geq 2.05\rho^2 \mid \mathbf{X}(\tau_2+t,0) = \mathbf{x}'] \leq n^{-a}$. If $Y(\tau_2+t,1) < 2.05\rho^2$, then the probability that even one undecided node remains undecided after the first decision-accumulation round is at most $Y(\tau_2+t,1) \cdot n \cdot Y(\tau_2+t,1) \leq 5\rho^4 n \leq n^{-3/4}$. Similar to the proof of Lemma 8, we boost this probability using a geometric random variable by considering the first $4 \leq T-1$ consecutive decision-accumulation rounds, such that $\Pr[Y(\tau_2+t,5) \geq 0 \mid \mathbf{X}(\tau_2+t,0) = \mathbf{x}'] \leq n^{-a} + n^{-3}$. Combined with Equation (25), we get the desired statement. \blacktriangleleft

4 Plurality Consensus with $\log k + 4$ Bits

The non-constant memory overhead of REPEATEDCLEANUP is due to the round counter used to synchronize phases. We now present a protocol that avoids this counter. Each node u stores its opinion $\text{op}_u \in [k]$, a *phase counter* $p_u \in \mathbb{N}$, and a *state variable* $s_u \in \{S, (U,I), (D,I), (U,A), (D,A)\}$. Our description and analysis assume p_u to be an arbitrary integer. While this would result in a non-constant memory overhead, we will prove that (w.h.p.) $|p_u - p_v| \leq 1$ for any two nodes u, v and any round. Thus, the actual implementation can restrict p_u to $\{0, 1, 2\}$, such that we get a memory overhead of $\log(3 \cdot 5) \leq 4$ bits. We call this protocol CONSTOVERHEAD. Our main result is the following theorem:

► Theorem 10. *CONSTOVERHEAD has a local memory overhead of 4 bits (15 states). If started on a biased configuration $\tilde{\mathbf{x}}$ with relative plurality gap $\gamma := \gamma(\tilde{\mathbf{x}})$, then (w.h.p.) plurality consensus is achieved in $O(\log n \cdot \log \log_\gamma n)$ rounds.*

The rest of this section describes the protocol as well as the underlying idea and gives a sketch of the analysis. More details can be found in the full version.


² This follows exactly like Claim 4(a) when switching x_i for y' and y for x'_1 . If only one opinion is left, undecided nodes in cleanup rounds increase exactly as decided nodes in decision-accumulation rounds.

```

1 sample a random node  $v$ 
2 if  $s_u = S$ : {cleanup}
3   if  $p_v \geq p_u$  or ( $p_v = p_u - 1$  and  $s_v = (D, A)$ ):
4     if  $op_v = op_u$ :  $s_u \leftarrow (D, I)$ 
5     else:  $s_u \leftarrow (U, I)$ 
6   elseif  $s_u \in \{(U, I), (D, I)\}$ : {slow transition}
7     if ( $p_v = p_u$  and  $v$  active) or  $p_v > p_u$  or  $\text{Ber}(1/n) = 1$ :
8       become active
9   elseif  $s_u = (U, A)$ : {decision-accumulation}
10    if ( $p_v = p_u$  and  $v$  decided) or  $p_v > p_u$ :
11       $op_u \leftarrow op_v$  and become decided
12  elseif  $s_u = (D, A)$ : {slow transition}
13  if  $p_v > p_u$  or  $\text{Ber}(1/n) = 1$ :  $(p_u, s_u) \leftarrow (p_u + 1, S)$ 

```

■ **Listing 2** Protocol CONSTOVERHEAD as executed by node u in round $t \in \mathbb{N}$. This description assumes p_u to be an arbitrary integer.



■ **Figure 1** Illustration how nodes transition through phases. Red (thick) state transitions highlight slow transitions that ensure that nodes do not get out of sync too much.

Protocol Overview. Initially, each node u starts with $p_u = 1$ and $s_u = S$. We call u *undecided* if $s_u = (U, \cdot)$, *decided* if $s_u = (D, \cdot)$, *inactive* if $s_u = (\cdot, I)$, and *active* if $s_u = (\cdot, A)$. At the start of a round, each node u samples a random node v and uses v 's data to transition through its phase. Listing 2 gives the formal protocol description and Figure 1 an illustration.

We start with a high level description of the protocol which we gradually refine. Basically, CONSTOVERHEAD mimics the synchronized behavior of REPEATEDCLEANUP. As in our first protocol, there are two ways to transition through a phase: if a node u samples another node v of the same opinion, it keeps its opinion and has to wait for the other nodes to catch up. While REPEATEDCLEANUP implemented this waiting via a counter, CONSTOVERHEAD uses *slow transitions* (highlighted red/thick in Figure 1). Otherwise, if u samples a node v of a different opinion, it becomes undecided and keeps sampling nodes until it finds a new opinion.

Becoming decided or undecided is modeled by entering the *inactive decided* state (D, I) or *inactive undecided* state (U, I) , respectively. This transition from S to one of these inactive states corresponds to cleanup rounds of REPEATEDCLEANUP. Now, there is a slow transition to the *active decided* state (D, A) and *active undecided* state (U, A) . This transition from inactive to active ensures that (a) decided nodes do not enter the next phase too early (which could require a large phase counter) and (b) undecided nodes do not sample decided nodes too early (which could result in a skewed distribution, since the number of decided nodes might be too small). Once an undecided node u becomes active, it keeps sampling nodes until it finds a new opinion. This corresponds to decision-accumulation rounds of REPEATEDCLEANUP. Note that the transition from (D, A) to S (at which the phase counter is increased) is slow. This ensures that not too many nodes enter the next phase before all undecided nodes found a new opinion (which could,

as before, require a large phase counter).

Some Subtleties. While the above reflects the basic behavior of our protocol, we omitted some details. Let us make a few important and useful observations:

- A node u always checks whether the sampled node v is not too far behind.
- Nodes do not explicitly forget their opinion when becoming undecided but simply overwrite their old opinion when they find a new opinion. In particular, if a node from an earlier state asks for u 's opinion while u is undecided, u answers with its most recent opinion (the provision of which does not cost us extra in terms of memory as we simply retain the information where it was; the undecided state is separately encoded in the already accounted-for $\Theta(1)$ additional bits).
- Slow transitions $s \rightarrow s'$ between two states s and s' basically simulate PULL-rumor spreading [18]: Nodes in state s' or later are “informed”, while all other nodes are “uninformed”. When u samples an informed node (a node in s' or beyond), it can cross the slow transition (independent of the sampled opinion). The Bernoulli trial $\text{Ber}(1/n)$ in slow transitions ensures that, eventually, at least one node is “informed”. Without it, no node could cross such a transition.

Analysis Overview. Given that `CONSTOVERHEAD` is based on the same principle as `REPEATEDCLEANUP`, namely employing synchronized cleanup rounds to increase the relative plurality gap exponentially from phase to phase, it is natural to use a similar analysis. The major difficulty stems from the fact that our synchronization primitive is now probabilistic (slow transitions) instead of deterministic (counter). In particular, in the case of `CONSTOVERHEAD` there is no guarantee that nodes wait at slow transitions for other nodes to catch up; in fact, there will typically be a few nodes that proceed early on over slow transitions. This might disturb our analysis in two ways: (a) if nodes could proceed arbitrarily far ahead, their phase counter could become arbitrarily high, resulting in a non-constant memory overhead, and (b) if a small group of nodes with a non-plurality opinion were “lucky” and proceeded fast, these nodes might cause more and more latecomers to adopt a non-plurality opinion.

The major tool to address both of these issues are two probabilistic synchronization results. The first shows that (w.h.p.) there is a period/stage of $O(\log n)$ consecutive rounds such that:

- If $n - \text{polylog}(n)$ nodes are in state (D, A) of phase $\tau - 1$ or state S of phase τ , and all remaining nodes are in one of the two inactive states (\cdot, I) of phase τ ,
- then at the end of this period, $n - \text{polylog}(n)$ nodes are in one of the two inactive states (\cdot, I) of phase τ , and all remaining nodes are in one of the two active states (\cdot, A) of phase τ .

The second result shows that (w.h.p.) there is a period/stage of $O(\log n)$ consecutive rounds such that:

- If $n - \text{polylog}(n)$ nodes are in one of the two inactive states (\cdot, I) of phase τ , and all remaining nodes are in one of the two active states (\cdot, A) of phase τ ,
- then at the end of this period, $n - \text{polylog}(n)$ nodes are in state (D, A) of phase τ or state S of phase $\tau + 1$, and all remaining nodes are in one of the two inactive states (\cdot, I) of phase τ .

Note that the final condition of the first stage fits perfectly into the assumption of the second stage and vice versa. We call the first stage the *cleanup stage* and the second stage the *decision-accumulation stage* (in the style of the corresponding round in `REPEATEDCLEANUP`). Since these stages are well-separated, we can prove an analogue of Claim 3 and an analogue of Claim 4. Equipped with these, the remainder of the proof is basically identical to the proof of `REPEATEDCLEANUP` (which was completely based on these concentration bounds).

References

- 1 Mohammed Amin Abdullah and Moez Draief. Majority consensus on random graphs of a given degree sequence. *CoRR*, abs/1209.5025, 2012. URL: <http://arxiv.org/abs/1209.5025>.
- 2 Dan Alistarh, Rati Gelashvili, and Milan Vojnovic. Fast and exact majority in population protocols. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 47–56, 2015. doi:10.1145/2767386.2767429.
- 3 Dana Angluin, James Aspnes, and David Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, 2008. doi:10.1007/s00446-008-0059-z.
- 4 Arta Babaee and Moez Draief. Distributed multivalued consensus. *The Computer Journal*, 57(8):1132–1140, 2014. doi:10.1093/comjnl/bxt026.
- 5 Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. Plurality consensus in the gossip model. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 371–390, 2015. doi:10.1137/1.9781611973730.27.
- 6 Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, and Luca Trevisan. Stabilizing consensus with many opinions. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 620–635, 2016. doi:10.1137/1.9781611974331.ch46.
- 7 P. Berenbrink, T. Friedetzky, P. Kling, F. Mallmann-Trenn, and C. Wastell. Plurality Consensus via Shuffling: Lessons Learned from Load Balancing. *ArXiv e-prints*, February 2016. arXiv:1602.01342.
- 8 Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. Coalescing random walks and voting on graphs. In *Proceedings of the 31st ACM Symposium on Principles of Distributed Computing (PODC)*, pages 47–56, 2012. doi:10.1145/2332432.2332440.
- 9 Colin Cooper, Robert Elsässer, and Tomasz Radzik. The power of two choices in distributed voting. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 435–446, 2014. doi:10.1007/978-3-662-43951-7_37.
- 10 Colin Cooper, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. Coalescing walks on rotor-router systems. In *Proceedings of the 22nd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 444–458, 2015. doi:10.1007/978-3-319-25258-2_31.
- 11 Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing consensus with the power of two choices. In *Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 149–158, 2011. doi:10.1145/1989493.1989516.
- 12 Moez Draief and Milan Vojnović. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimisation*, 50(3):1087–1109, 2012.
- 13 Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Frederik Mallmann-Trenn, and Horst Trinker. Efficient k -party voting with two choices. *CoRR*, abs/1602.04667, 2016. URL: <http://arxiv.org/abs/1602.04667>.
- 14 Mohsen Ghaffari and Merav Parter. A polylogarithmic gossip algorithm for plurality consensus. In *Proceedings of the 35th ACM Symposium on Principles of Distributed Computing (PODC)*, 2016. to appear.
- 15 Y. Hassin and D. Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Information and Computation*, 171(2):248–268, 2001. doi:10.1006/inco.2001.3088.
- 16 R. Holley and T. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, 3(4):643–663, 1975. URL: <http://www.jstor.org/stable/2959329>.

XX:14 Efficient Plurality Consensus

- 17 K. Jung, B. Y. Kim, and M. Vojnovic. Distributed ranking in networks with limited memory and communication. In *Proceedings of the 2012 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 980–984, 2012. doi:10.1109/ISIT.2012.6284710.
- 18 Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. Randomized rumor spreading. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000. doi:10.1109/SFCS.2000.892324.
- 19 Etienne Perron, Dinkar Vasudevan, and Milan Vojnovic. Using three states for binary consensus on complete graphs. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 2527–2535, 2009. doi:10.1109/INFOCOM.2009.5062181.