

# Communication-Efficient Randomized Consensus

Dan Alistarh\*      James Aspnes†      Valerie King‡      Jared Saia§

## Abstract

We consider the problem of consensus in the challenging *classic* model. In this model, the adversary is adaptive; it can choose which processors crash at any point during the course of the algorithm. Further, communication is via asynchronous message passing: there is no known upper bound on the time to send a message from one processor to another, and all messages and coin flips are seen by the adversary.

We describe a new randomized consensus protocol with expected message complexity  $O(n^2 \log^2 n)$  when fewer than  $n/2$  processes may fail by crashing. This is an almost-linear improvement over the best previously known protocol, and within logarithmic factors of a known  $\Omega(n^2)$  message lower bound. The protocol further ensures that no process sends more than  $O(n \log^3 n)$  messages in expectation, which is again within logarithmic factors of optimal. We also present a generalization of the algorithm to an arbitrary number of failures  $t$ , which uses expected  $O(nt + t^2 \log^2 t)$  total messages.

Our approach is to build a message-efficient, resilient mechanism for aggregating individual processor votes, implementing the message-passing equivalent of a weak shared coin. Roughly, in our protocol, a processor first announces its votes to small groups, then propagates them to increasingly larger groups as it generates more and more votes. To bound the number of messages that an individual process might have to send or receive, the protocol progressively increases the weight of generated votes. The main technical challenge is bounding the impact of votes that are still “in flight” (generated, but not fully propagated) on the final outcome of the shared coin, especially since such votes might have different weights. We achieve this by leveraging the structure of the algorithm, and a technical argument based on martingale concentration bounds. Overall, we show that it is possible to build an efficient message-passing implementation of a shared coin, and in the process (almost-optimally) solve the classic consensus problem in the asynchronous message-passing model.

## 1 Introduction

Consensus [30, 31] is arguably the most well-studied problem in distributed computing. The FLP impossibility result [23], showing that consensus could not be achieved deterministically in an asynchronous message-passing system with even one crash failure, sparked a flurry of research on overcoming this fundamental limitation, either by adding timing assumptions, e.g. [21], employing failure detectors, e.g. [19], or by relaxing progress conditions to allow for *randomization*, e.g. [14]. A significant amount of research went into isolating time and space complexity bounds for randomized consensus in the shared-memory model, e.g. [3, 5, 9, 10, 12, 17, 22], developing elegant and technically complex tools in the process. As a result, the time complexity of consensus in asynchronous shared memory is now well characterized: the tight bound on total number of steps is  $\Theta(n^2)$  [12], while the individual step bound is  $\tilde{\Theta}(n)$  [5].<sup>1</sup>

---

\*Microsoft Research Cambridge. Email: dan.alistarh@microsoft.com.

†Yale University, Department of Computer Science. Supported in part by NSF grant CCF-0916389. Email:aspnes@cs.yale.edu.

‡University of Victoria; Simons Institute for the Theory of Computing, Berkeley; Institute for Advanced Study, Princeton.

§University of New Mexico, Department of Computer Science. Supported in part by NSF CAREER Award 0644058 and NSF CCR-0313160.

<sup>1</sup>We consider a standard model with  $n$  processes,  $t < n/2$  of which may fail by crashing. The  $\tilde{\Theta}$  notation hides logarithmic factors.

Somewhat surprisingly, the complexity of randomized consensus in the other core model of distributed computing, the *asynchronous message-passing model*, is much less well understood. In this model, communication is via full-information, asynchronous message passing: there is no known upper bound on the time to send a message from one processor to another, and all messages are seen by the adversary. Further, as in the shared memory model, the adversary is adaptive; it can choose which processors crash at any point during the course of the algorithm. We refer to this as the *classic* model.

While simulations exist [11] allowing shared-memory algorithms to be translated to message-passing, their overhead is at least linear in the number of nodes. It is therefore natural to ask if message-efficient solutions for randomized consensus can be achieved, and in particular if quadratic shared-memory communication cost for consensus can be also achieved in message-passing systems against a strong, adaptive adversary.

In this paper, we propose a new randomized consensus protocol with expected message complexity  $O(n^2 \log^2 n)$  against a strong (adaptive) adversary, in an asynchronous message-passing model in which less than  $n/2$  processes may fail by crashing. This is an almost-linear improvement over the best previously known protocol. Our protocol is also *locally-efficient*, ensuring that no process sends or receives more than expected  $O(n \log^3 n)$  messages, which is within logarithmic factors of the linear lower bound [12]. We also provide a generalization to an arbitrary number of failures  $t < n/2$ , which uses  $O(nt + t^2 \log^2 t)$  messages.

Our general strategy is to construct a **message-efficient weak shared coin**. A weak shared coin with parameter  $\delta > 0$  is a protocol in which for each possible return value  $\pm 1$ , there is a probability of at least  $\delta$  that all processes return that value. We then show that that this shared coin can be used in a message-efficient consensus protocol modeled on a classic shared-memory protocol of Chandra [18].

Since early work by Bracha and Rachman [17], implementations of weak shared coins for shared-memory systems with an adaptive adversary have generally been based on voting. If the processes between them generate  $n^2$  votes of  $\pm 1$ , then the absolute value of the sum of these votes will be at least  $n$  with constant probability. If this event occurs, then even if the adversary hides  $\Theta(n)$  votes by crashing processes, the total vote seen by the survivors will still have the same sign as the actual total vote.

Such algorithms can be translated to a message-passing setting directly using the classic Attiya-Bar-Noy-Dolev (ABD) simulation [11]. The main idea of the simulation is that a write operation to a register is simulated by distributing a value to a majority of the processes (this is possible because of the assumption that a majority of the processes do not fail). Any subsequent read operation contacts a majority of the processes, and because the majorities overlap, this guarantees that any read sees the value of previous writes.

The obvious problem with this approach is that its message complexity is high: because ABD uses  $\Theta(n)$  messages to implement a write operation, and because each vote must be written before the next vote is generated if we are to guarantee that only  $O(n)$  votes are lost, the cost of this direct translation is  $\Theta(n^3)$  messages. Therefore, the question is whether this overhead can be eliminated.

To reduce both total and local message complexity, we employ two new ingredients.

The first is an algorithmic technique to reduce the message complexity of distributed vote counting by using a binary tree of process groups called **cohorts**, where each leaf corresponds to a process, and each internal node represents a cohort consisting of all processes in the subtree. Instead of announcing each new vote to all participants, new  $\pm 1$  votes are initially only announced to small cohorts at the bottom of the tree, but are propagated to increasingly large cohorts as more votes are generated. As the number of votes grows larger, the adversary must crash more and more processes to hide them. This generalizes the one-crash-one-vote guarantee used in shared-memory algorithms to a many-crashes-many-votes approach.

At the same time, this technique renders the algorithm message-efficient. Given a set of generated votes, the delayed propagation scheme ensures that each vote accounts for exactly one update at the leaf,  $1/2$  updates at the 2-neighborhood, and in general,  $1/2^i$  updates at the  $i$ th level of the tree. Practically, since the  $i$ th level cohort has  $2^i$  members, the propagation cost of a vote is exactly one message per tree level. In total, that is  $\log n$  messages per vote, amortized.

An limitation of the above scheme is that a fast process might have to generate all the  $\Theta(n^2)$  votes

itself in order to decide, which would lead to high individual message complexity. The second technical ingredient of our paper is a procedure for assigning increasing weight to a processes’ votes, which reduces individual complexity. This general idea has previously been used to reduce individual work for shared-memory randomized consensus [5, 7, 10]; however, we design and analyze a new weighting scheme that is customized for our vote-propagation mechanism.

In our scheme, each process doubles the weight of its votes every  $4n \log n$  votes, and we run the protocol until the total reported variance—the sum of the squares of the weights of all reported votes—exceeds  $n^2 \log n$ . Intuitively, this allows a process running alone to reach the threshold quickly, reducing per-process message complexity. This significantly complicates the termination argument, since a large number of generated votes, of various weights, could be still making their way to the root at the time when a process first notices the termination condition. We show that, with constant probability, this drift is not enough to influence the sign of the sum, by carefully bounding the weight of the extra votes via the structure of the algorithm and martingale concentration bounds. We thus obtain a constant-bias weak shared coin. The bounds on message complexity follow from bounds on the number of votes generated by any single process or by all the processes together before the variance threshold is reached.

We convert the shared coin construction into a consensus algorithm via a simple framework inspired by Chandra’s shared-memory consensus protocol [18], which in turn uses ideas from earlier consensus protocols of Chor, Israeli, and Li [20] and Aspnes and Herlihy [9]. Roughly, we associate each of the two possible decision values with a message-passing implementation of a max register [6, 8], whose value is incremented by the “team” of processes obtaining that value from the shared coin. If a process sees that its own team has fallen behind, it switches to the other team, and once one of the max register’s values surpasses the other by two, the corresponding team wins. Ties are broken (eventually) by having processes that do not observe a clear leader execute a weak shared coin. This simple protocol maintains the asymptotic complexity of the shared coin in expectation.

Finally, we present a more efficient variant of the protocol for the case where  $t = o(n)$ , based on the observation that we can “deputize” a subset of  $2t + 1$  of the processes to run the consensus protocol, and broadcast their result to all  $n$  processes. The resulting protocol has total message complexity  $O(nt + t^2 \log^2 t)$ , and  $O(n + t \log^3 t)$  individual message complexity.

Overall, we show that it is possible to build message-efficient weak shared coins and consensus in asynchronous message-passing systems. Our vote counting construction implements a message-efficient, asynchronous approximate trigger counter [27], which may be of independent interest. An interesting aspect of our constructions is that message sizes are small: since processes only communicate vote counts, messages only require  $O(\log n)$  bits of communication.

## 2 System Model and Problem Statement

We consider the standard asynchronous message-passing model, in which  $n$  processes communicate with each other by sending messages through channels. We assume that there are two uni-directional channels between any pair of processes. Communication is *asynchronous*, in that messages can be arbitrarily delayed by a channel. However, we assume that messages do not become corrupted.

Computation proceeds in a sequence of *steps*. At each step, a process checks incoming channels for new messages, then performs local computation, and sends new messages. A process may become *faulty*, in which case it ceases to perform local computation and to send new messages. A process is *correct* if it takes steps infinitely often during the execution. We assume that at most  $t < n/2$  processes may be faulty during the execution.

Message delivery and process faults are assumed to be controlled by a *strong (adaptive) adversary*. At any time during the computation, the adversary can examine the entire state of the system (in particular, the results of process coinflips), and decide on process faults and messages to be delivered.

The (*worst-case*) *message complexity* of an algorithm is simply the maximum, over all adversarial strate-

gies, of the total number of messages sent by processes running the algorithm. Without loss of generality, we assume that the adversary’s goal is to maximize the message complexity of our algorithm.

In the (*binary*) *randomized consensus* problem, each process starts with an input in  $\{0, 1\}$ , and returns a decision in  $\{0, 1\}$ . A correct protocol satisfies *agreement*: all processes that return from the protocol choose the same decision, *validity*: the decision must equal some process’s input, and *probabilistic termination*: every non-faulty process returns after a finite number of steps, with probability 1.

### 3 Related Work

The first shared-memory protocol for consensus was given by Chor, Israeli, and Li [20] for a weak adversary model, and is based on a race between processors to impose their proposals. Abrahamson [1] gave the first wait-free consensus protocol for a strong adversary, taking exponential time. Aspnes and Herlihy [9] gave the first polynomial-time protocol, which terminates in  $O(n^4)$  expected total steps. Subsequent work, e.g. [3, 10, 17, 32], continued to improve upper and lower bounds for this problem, until Attiya and Censor [12] showed a tight  $\Theta(n^2)$  bound on the total number of steps for asynchronous randomized consensus. In particular, their lower bound technique implies an  $\Omega(t(n - t))$  total message complexity lower bound and a  $\Omega(t)$  individual message complexity lower bound for consensus in the asynchronous message-passing model. Our  $(n/2 - 1)$ -resilient algorithms match both lower bounds within logarithmic factors, while the  $t$ -resilient variant matches the first lower bound within logarithmic factors.

To our knowledge, the best previously known upper bound for consensus in asynchronous message-passing requires  $\Theta(n^3)$  messages. This is obtained by simulating the elegant shared-memory protocol of Attiya and Censor-Hillel [12], using the simulation from [11]. A similar bound can be obtained by applying the same simulation to an  $O(n)$ -individual-work algorithm of Aspnes and Censor [7].

In the message passing model, significant work has focused on the problem of Byzantine agreement, which is identical to consensus except that the adversary *controls* up to  $t$  processes, and can cause them to deviate arbitrarily from the protocol. In 1983, Fischer, Lynch and Patterson [23] showed that no deterministic algorithm could solve consensus, and hence Byzantine agreement, in the classic model. In the same year, Ben-Or gave a randomized algorithm for Byzantine agreement which required an expected exponential communication rounds and number of messages [14]. His algorithm tolerated  $t < n/5$ . Subsequent work extended this idea in two directions: to solve message-passing Byzantine agreement faster and with higher resilience, and to solve agreement wait-free, in asynchronous shared-memory tolerating crash failures.

Resilience against Byzantine faults was improved to  $t < n/3$  in 1984 by Bracha [16]. However, the communication rounds and number of messages remained exponential in expectation. This resilience is the best possible for randomized Byzantine agreement [26]. In 2013, King and Saia gave the first algorithm for Byzantine agreement in the classic model with expected polynomial communication rounds and number of messages [28]. Their algorithm required in expectation  $O(n^{2.5})$  communication rounds,  $O(n^{6.5})$  messages, and  $O(n^{7.5})$  bits sent. It tolerated  $t < n/500$ . Unfortunately, local computation time was exponential. In 2014, the same authors achieved polynomial computation time. However, the new algorithm required expected  $O(n^3)$  communication rounds,  $O(n^7)$  messages, and  $O(n^8)$  bits sent. Further, the resilience decreased to  $t < 0.000028n$  [29].

### 4 A Message-Passing Max Register

To coordinate the recording of votes within a group, we use a message-passing max register [6]. The algorithm is adapted from [8], and is in turn based on the classic ABD implementation of a message-passing register [11]. The main change from [8] is that we allow for groups consisting of  $g < n$  processes. Recall that a max register maintains a value  $v$ , which is read using the *MaxRead* operation, and updated using the *MaxUpdate* operation. A *MaxUpdate*( $u$ ) operation changes the value only if  $u$  is higher than the current value  $v$ .

**Description.** We consider a group  $G$  of  $g$  processes, which implement the max register  $R$  collectively. Each process  $p_i$  in the group maintains a current value estimate  $v_i$  locally. The communicate procedure [11] broadcasts a request to all processes in the group  $G$ , and waits for at least  $\lceil g/2 \rceil$  replies.<sup>2</sup>

To perform a MaxRead, the process communicates a  $MaxRead(R)$  request to all other processes, setting its value  $v_i$  to be the maximum value received. Before returning this value, the process communicates a  $MaxReadACK(R, v_i)$  message. All processes receiving such a message will update their current estimate of  $R$ , if this value was less than  $v_i$ . If it receives at least  $\lceil g/2 \rceil$  replies, the caller returns  $v_i$  as the value read. This ensures that, if a process  $p_i$  returns  $v_i$ , no other process may later return a smaller value for  $R$ .

A MaxUpdate with input  $u$  is similar to a MaxRead: the process first communicates a  $MaxUpdate(R, u)$  message to the group, and waits for at least  $\lceil g/2 \rceil$  replies. Process  $p_i$  sets its estimate  $v_i$  to the maximum between  $u$  and the maximum value received in the first round, before communicating this value once more in a second broadcast round. Again, all processes receiving this message will update their current estimate of  $R$ , if necessary. The algorithm ensures the following properties.

**Lemma 1.** *The max register algorithm above implements a linearizable max register. If the communicate procedure broadcasts to a group  $G$  of processes of size  $g$ , then the message complexity of each operation is  $O(g)$ , and the operation succeeds if at most  $\lfloor g/2 \rfloor$  processes in the group are faulty.*

## 5 The Weak Shared Coin Algorithm

We now build a message-efficient asynchronous weak shared coin. Processes generate random votes, whose weight increases over time, and progressively communicate them to groups of nodes of increasing size. This implements a shared coin with constant bias, which in turn can be used to implement consensus.

**Vote Propagation.** The key ingredient is a message-efficient construction of an approximate asynchronous vote counter, which allows processes to maintain an estimate of the total number of votes generated, and of their sum and variance. The distributed vote counter is structured as a binary tree, where each process is associated with a leaf. Each subtree of height  $h$  is associated with a *cohort* of  $2^h$  processes, corresponding to its leaves. To each such subtree  $s$ , we associate a max register  $R_s$ , implemented as described above, whose value is maintained by all the processes in the corresponding cohort. For example, the value at each leaf is only maintained by the associated process, while the root value is tracked by all processes. (Please see Figure 1 for an illustration.)

The max register  $R_s$  corresponding to the subtree rooted at  $s$  maintains three values: the *count*, an estimate of the number of votes generated in the subtree, *total*, an estimate of the sum of generated votes, and *var*, an estimate of the variance of the generated votes. Values are ordered only by the first component. Practically, the implementation is identical to the max register described in the previous section, except that whenever sending the *count*, processes also send the associated *total* and *var*. Processes always adopt the tuple of maximum *count*. If a process receives two tuples with the same *count* but different *total/var* components, they adopt the one with the maximum total.

A process maintains max register estimates for each subtree it is part of. Please see Algorithm 1 for the pseudocode. In the  $k$ th iteration of the shared coin, the process generates a new vote with weight  $\pm w_k$  chosen as described in the next paragraph. After generating the vote, the process will propagate its current set of votes up to level  $r$ , the highest power of two which divides  $k$  (line 15). At each level from 1 (the leaf's parent) up to  $r$ , the process reads the max registers left and right children, and updates the  $\langle \text{count}, \text{total}, \text{var} \rangle$  of the parent to be the sum of the corresponding values at the child max registers (lines 17–20).

If  $n$  divides *total*, then the process also checks the count at the root. If the root variance count is greater than the threshold of  $K$  votes, the process returns the sign of the root total as its output from the shared coin.

<sup>2</sup>Since  $t < n/2$  processes may crash, and  $g$  may be small, a process may block while waiting for replies. This only affects the progress of the protocol, but not its safety. Our shared coin implementation will partition the  $n$  processes into max register groups, with the guarantee that some groups always make progress.

```

1 Let  $K = n^2 \log_2 n$ 
2 Let  $T = 4n \log_2 n$ 
3 count  $\leftarrow 0$ 
4 var  $\leftarrow 0$ 
5 total  $\leftarrow 0$ 
6 for  $k \leftarrow 1, 2, \dots, \infty$  do
7   Let  $w_k = 2^{\lfloor (k-1)/T \rfloor}$ 
8   Let vote =  $\pm w_k$  with equal probability
9   count  $\leftarrow$  count + 1
10  var  $\leftarrow$  var +  $w_k^2$ 
11  total  $\leftarrow$  total + vote
12  Write  $\langle$ count, var, total $\rangle$  to max register for my leaf
13  for  $j \leftarrow 1 \dots \log_2 n$  do
14    if  $2^j$  does not divide  $k$  then
15      break
16    Let  $s$  be my level- $j$  ancestor, with children  $s_\ell$  and  $s_r$ 
17    in parallel do
18      /* read left and right counts */
19       $\langle$ count $_\ell$ , var $_\ell$ , total $_\ell$  $\rangle \leftarrow$  ReadMax( $s_\ell$ )
20       $\langle$ count $_r$ , var $_r$ , total $_r$  $\rangle \leftarrow$  ReadMax( $s_r$ )
21      /* update the parent */
22      WriteMax( $s$ ,  $\langle$ count $_\ell$  + count $_r$ , var $_\ell$  + var $_r$ , total $_\ell$  + total $_r$  $\rangle$ )
23  if  $n$  divides  $k$  then
24     $\langle$ count $_{\text{root}}$ , var $_{\text{root}}$ , total $_{\text{root}}$  $\rangle \leftarrow$  ReadMax(root)
25    /* if the root variance exceeds the threshold */
26    if var $_{\text{root}} \geq K$  then
27      return sgn(total $_{\text{root}}$ ) /* return sign of root total */

```

**Algorithm 1:** Shared coin using increasing votes.

Otherwise, the process continues to generate votes.

**Vote Generation.** Each process generates votes with values  $\pm w_k$  in a series of epochs, each epoch containing of  $T = 4n \log_2 n$  loop iterations. Within each epoch, all votes have the same weight, and votes are propagated up the tree of max registers using the schedule described above. At the start of a new epoch, the weight of the votes doubles. This ensures that only  $O(\log n)$  epochs are needed until a single process can generate enough variance by itself to overcome the offsets between the observed vote and the generated vote due to delays in propagation up the tree.

Because votes have differing weights, we track the total variance of all votes included in a max register in addition to their number, and continue generating votes until this total variance exceeds a threshold  $K = n^2 \log_2 n$ , at which point the process returns the sign of the root total (line 24).

## 6 Algorithm Analysis

We prove that the algorithm in Section 5 implements a correct weak shared coin. We first analyze some of the properties of the tree-based vote counting structure. For simplicity, we assume that the number of processes  $n$  is a power of two. Due to space constraints, the complete argument is given in the Appendix.

**Vote Propagation.** The algorithm is based on the idea that, as processes take steps, counter values for the cohorts get increased, until, eventually, the root counter value surpasses the threshold, and processes start to return. We first provide a way of associating a set of generated votes to each counter value.

We say that a process  $p_i$  counts a number  $x_i$  of (consecutive) locally-generated votes to node  $s$  if, after generating the last such vote, process  $p_i$  updates the max register at  $s$  in its loop iteration. We prove that this

procedure has the following property:

**Lemma 2.** *Consider a subtree rooted at node  $s$  with  $\ell$  leaves, corresponding to member processes  $q_1, q_2, \dots, q_\ell$ . Let  $x_1, x_2, \dots, x_\ell$  be the number of votes most recently counted by processes  $q_1, q_2, \dots, q_\ell$  at node  $s$ , respectively. Then the value of the count component of the max register at  $s$  is at least  $\sum_{m=1}^{\ell} x_m$ .*

**Proof Strategy.** The proof can be divided into three steps. The first shows that the coin construction offers a good approximation of the generated votes, i.e. the total vote  $U^t$  observed in the root max register at any time  $t$  is close to the actual total generated vote  $V^t$  at the same time. The second step shows that when the threshold  $K$  is crossed at some time  $t$ , the **common votes total**  $|V^t|$  is likely to be far away from 0. The last step shows that, for any subsequent time  $t'$ , the combination of the approximation slack  $U^{t'} - V^{t'}$  and any **extra votes**  $V^{t'} - V^t$  observed by a particular process at time  $t'$  will not change the sign of the total vote.

The first step involves a detailed analysis of what votes may be omitted from the visible total combined with an extension of the Azuma-Hoeffding inequality [10]; the second step requires use of a martingale Central Limit Theorem [25, Theorem 3.2]; the last follows from an application of Kolmogorov's inequality. We begin by stating a few technical claims, whose proof can be found in the Appendix.

**Lemma 3.** *Fix some execution of Algorithm 1, let  $n_i$  be the number of votes generated by process  $p_i$  during this execution, and let  $w_{n_i}$  be the weight of the  $n_i$ -th vote. Then*

1.  $\sum_{i=1}^n \sum_{j=1}^{n_i} w_j^2 \leq \frac{K+2n^2}{1-8n/T} = O(n^2 \log n)$ .
2.  $w_{n_i} \leq \sqrt{1 + \frac{4K+8n^2}{T-8n}} = O(\sqrt{n})$ .
3. For all  $j$ ,  $n_j = O(n \log^2 n)$ .
4.  $\sum_i w_{n_i}^2 \leq n + \frac{4K+8n^2}{T-8n} = O(n)$ .
5.  $\sum_i n_i = O(n^2 \log n)$ .

For any adversary strategy  $\mathcal{A}$ , let  $\tau_{\mathcal{A}}$  be a stopping time corresponding to the first time  $t$  at which  $U^{\text{root}}[t].\text{var} \geq K$ . We will use a martingale Central Limit Theorem to show that  $V^{\text{root}}[\tau_{\mathcal{A}}].\text{total}$  converges to a normal distribution as  $n$  grows, when suitably scaled. This will then be used to show that all processes observe a population of common votes whose total is likely to be far from zero. The notation  $X \xrightarrow{p} Y$  means that  $X$  converges in probability to  $Y$ , and  $Y \xrightarrow{d} Y$  means that  $X$  converges in distribution to  $Y$ . We show that the following convergence holds (please see the Appendix for the proof).

**Lemma 4.** *Let  $\{\mathcal{A}_n\}$  be a family of adversary strategies, one for each number of processes  $n \in \mathbb{N}$ . Let  $\tau_n = \tau_{\mathcal{A}_n}$  be as above. Then*

$$\frac{V^{\text{root}}[\tau_n].\text{total}}{\sqrt{K}} \xrightarrow{d} N(0, 1). \quad (1)$$

Once this is established, for each subtree  $s$  and time  $t_s$ , let  $D^s[t_s] = V^s[t_s].\text{total} - U^s[t_s].\text{total}$  be the difference between the generated votes in  $s$  at time  $t_s$  and the votes reported to the max register corresponding to  $s$  at time  $t_s$ . Let  $s_\ell$  and  $s_r$  be the left and right subtrees of  $s$ , and let  $t_{s_\ell}$  and  $t_{s_r}$  be the times at which the values added to produce  $U^s[t_s].\text{total}$  were read from these subtrees. Recursing over all proper subtrees of  $s$ , we obtain that

$$D^s[t_s] = V^s[t_s].\text{total} - U^s[t_s].\text{total} = \sum_{s'} \left( V^{s'}[t_{\text{parent}(s')}].\text{total} - V^{s'}[t_{s'}].\text{total} \right),$$

where  $s'$  ranges over all proper subtrees of  $s$ .

To bound this sum, we consider each (horizontal) layer of the tree separately, and observe that the missing interval of votes  $\left(V^{s'}[t_{\text{parent}(s')}] \cdot \text{total} - V^{s'}[t_{s'}] \cdot \text{total}\right)$  for each subtree  $s$  in layer  $h$  consists of at most  $2^h$  votes by each of at most  $2^h$  processes. For each process  $p_i$  individually, the variance of its  $2^h$  heaviest votes, using Lemma 3, is at most  $2^h \left(1 + (4/T) \sum_{j=1}^{n_i} w_j^2\right)$ . If we sum the total variance of at most  $2^h$  votes from all processes, we get at most

$$2^h \left( n^2 + (4/T) \sum_{i=1}^n \sum_{j=1}^{n_i} w_j^2 \right) \leq 2^h \left( n^2 + \frac{K + 2n^2}{1 - 8n/T} \right),$$

again using Lemma 3.

We would like to use this bound on the total variance across all missing intervals to show that the sum of the total votes across all missing intervals is not too large. Intuitively, if we can apply a bound to the total variance on a particular interval, we expect the Azuma-Hoeffding inequality to do this for us. But there is a complication in that the total variance for an interval may depend in a complicated way on the actions taken by the adversary during the interval. So instead, we attack the problem indirectly, by adopting a different characterization of the relevant intervals of votes and letting the adversary choose between them to obtain the actual intervals that contributed to  $D^{\text{root}}[t]$ . We will use the following extended version of the classic Azuma-Hoeffding inequality [10]:

**Lemma 5** ([10, Theorem 4.5]). *Let  $\{S_0, \mathcal{F}_0\}$ ,  $0 \leq i \leq n$  be a zero-mean martingale with difference sequence  $\{X_i\}$ . Let  $w_i$  be measurable  $\mathcal{F}_{i-1}$ , and suppose that for all  $i$ ,  $|X_i| \leq w_i$  with probability 1; and that there exists a bound  $W$  such that  $\sum_{i=1}^n w_i^2 \leq W$  with probability 1. Then for any  $\lambda > 0$ ,*

$$\Pr[S_n \geq \lambda] \leq e^{-\lambda^2/2W}. \quad (2)$$

Fix an adversary strategy. For each subtree  $s$ , let  $X_1^s, X_2^s, \dots$  be the sequence of votes generated in  $s$ . For each  $s, t$ , and  $W$ , let  $Y_i^{tsW} = X_i^s$  if (a) at least  $t$  votes have been generated by all processes before  $X_i^s$  is generated, and (b)  $\sum_{j < i} (Y_j^{tsW})^2 + (X_i^s)^2 \leq W$ . Otherwise, let  $Y_i^{tsW}$  be 0. If we let  $\mathcal{F}_i$  be generated by all votes preceding  $X_i^s$ , then the events (a) and (b) are measurable  $\mathcal{F}_i$ , so  $\{Y_i^{tsW}, \mathcal{F}_i\}$  forms a martingale. Furthermore, since only the sign of  $Y_i^{tsW}$  is unpredictable, we can define  $w_i = (Y_i^{tsW})^2$  to fit Lemma 5. From (b), we have that  $\sum w_i^2 \leq W$  always. It follows that, for any  $c > 0$ ,

$$\Pr \left[ \sum_i Y_i^{tsW} \geq \sqrt{2cW \ln n} \right] \leq e^{-c \ln n} = n^{-c}.$$

There are polynomially many choices for the parameters  $t, s$ , and  $W$ . Union bounding over all such choices shows that, for  $c$  sufficiently large, with high probability  $\sum_i Y_i^{tsW}$  is bounded by  $\sqrt{2cW \ln n}$  for all such intervals. We now use this to show the following.

**Lemma 6.** *For any adversary strategy and sufficiently large  $n$ , with probability  $1 - o(1)$ , it holds that at all times  $t$ ,*

$$|V^{\text{root}}[t] \cdot \text{total} - U^{\text{root}}[t] \cdot \text{total}| \leq 6n \sqrt{\log_2 n}.$$

*Proof.* We are trying to bound  $D^{\text{root}}[t] = \sum_s (V^s[t_{\text{parent}(s)}] - V^s[t_s])$ , where  $s$  ranges over all proper subtrees of the tree and for each  $s$  of size  $2^h$ , the interval  $(t_s, t_{\text{parent}(s)})$  includes at most  $2^h$  votes for each process.

Suppose that for each  $t, s, W$ , it holds that  $Y^{tsW} \leq \sqrt{9W \ln n}$ . By the preceding argument, each such event fails with probability at most  $n^{-9/2}$ . There are  $O(n^2 \log n)$  choices for  $t$ ,  $O(n)$  choices for  $s$ , and



$O(n^2 \log n)$  choices for  $W$ , so taking a union bound over all choices of  $Y^{tsW}$  not occurring shows that this event occurs with probability  $O(n^{-1/2} \log^2 n) = o(1)$ .

If all  $Y^{tsW}$  are bounded, then it holds deterministically that

$$\begin{aligned} \sum_s (V^s[t_{\text{parent}(s)}] - V^s[t_s]) &= \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} (V^s[t_{\text{parent}(s)}] - V^s[t_s]) \\ &= \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} Y^{tsW_s} \leq \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} \sqrt{2cW_s \ln n} \\ &= \sqrt{2c \ln n} \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} \sqrt{W_s}, \end{aligned}$$

where  $W_s$  is the total variance of the votes generated by  $s$  in the interval  $(t_s, t_{\text{parent}(s)})$ .

Note that this inequality does not depend on analyzing the interaction between voting and when processes read and write the max registers. For the purposes of computing the total offset we are effectively allowing the adversary to choose what intervals it includes retrospectively, after carrying out whatever strategy it likes for maximizing the probability that any particular values  $Y^{tsW}$  are too big.

Because each process  $i$  in  $s$  generates at most  $2^h$  votes, and each such vote has variance at most  $w_{n_i}^2$ , we have

$$W_s \leq 2^h \sum_{i \in s} w_{n_i}^2.$$

Furthermore, the subtrees at any fixed level  $h$  partition the set of processes, so applying Lemma 3 gives

$$\sum_{s, |s|=2^h} W_s \leq \sum_{s, |s|=2^h} 2^h \sum_{i \in s} w_{n_i}^2 = 2^h \sum_{s, |s|=2^h} \sum_{i \in s} w_{n_i}^2 = 2^h \sum_i w_{n_i}^2 \leq 2^h \left( n + \frac{4K + 8n^2}{T - 8n} \right).$$

By concavity of square root,  $\sum \sqrt{x_i}$  is maximized for non-negative  $x_i$  constrained by a fixed bound on  $\sum x_i$  by setting all  $x_i$  equal. Setting all  $n/2^h$  values  $W_s$  equal gives the following upper bound.

$$\begin{aligned} W_s &\leq \frac{2^h}{n} \cdot 2^h \left( n + \frac{4K + 8n^2}{T - 8n} \right) = 2^{2h} \left( 1 + \frac{4K + 8n^2}{Tn - 8n^2} \right), \text{ and thus} \\ \sqrt{W_s} &\leq 2^h \sqrt{1 + \frac{4K + 8n^2}{Tn - 8n^2}}, \end{aligned}$$

which gives the bound

$$\begin{aligned} &\sqrt{9 \ln n} \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} \sqrt{W_s} \leq \sqrt{9 \ln n} \sum_{h=0}^{\log_2 n - 1} 2^h \sqrt{1 + \frac{4K + 8n^2}{Tn - 8n^2}} \\ &= \sqrt{9 \ln n \left( 1 + \frac{4K + 8n^2}{Tn - 8n^2} \right)} \sum_{h=0}^{\log_2 n - 1} 2^h = 3(n-1) \sqrt{\ln n} \sqrt{1 + \frac{4K + 8n^2}{Tn - 8n^2}} \\ &= 3(n-1) \sqrt{\ln n} \sqrt{1 + \frac{\log_2 n + 2}{\log_2 n - 2}} \leq 6n \sqrt{\log_2 n}, \end{aligned}$$

when  $n$  is sufficiently large. The last step uses the fact that for  $K = n^2 \log n$  and  $T = 4n \log n$ , the value under the radical converges to 2 in the limit, and  $3\sqrt{2/\ln 2} < 6$ .  $\square$

For the last step of the proof, we need to show that the **extra votes** that arrive after  $K$  variance has been accumulated are not enough to push  $V^{\text{root}}[t]$  close to the origin. For this, we use Kolmogorov’s inequality, a martingale analogue to Chebyshev’s inequality, which says that if we are given a zero-mean martingale  $\{S_i, \mathcal{F}_i\}$  with bounded variance, then  $\Pr[\exists i \leq n : |S_i| \geq \lambda] \leq \frac{\lambda^2}{\text{Var}[S_n]}$ .

Consider the martingale  $S_1, S_2, \dots$  where  $S_i$  is the sum of the first  $i$  votes after  $V^{\text{root}}[i].\text{var}$  first passes  $K$ . Then from Lemma 3,

$$\text{Var}[S_i] \leq \frac{K + 2n^2}{1 - 8n/T} - K = \frac{8K(n/T) + 2n^2}{1 - 8n/T} = \frac{8n^2 + 2n^2}{1 - 2/(\log_2 n)} = O(n^2).$$

So for any fixed  $c$ , the probability that  $|S_i|$  exceeds  $cK$  for any  $i$  is  $O(1/\log n) = o(1)$ .

**Final argument.** From Lemma 4, we have that the total common vote  $V^{\text{root}}[\tau_n].\text{total}$  converges in distribution to  $N(0, 1)$  when scaled by  $\sqrt{K} = n\sqrt{\log_2 n}$ . In particular, for any fixed constant  $c$ , there is a constant probability  $\pi_c > 0$  that for sufficiently large  $n$ ,  $\Pr[V^{\text{root}}[\tau_n] \geq cn\sqrt{\log_2 n}] \geq \pi_c$ .

Let  $c$  be 7. Then with probability  $\pi_7 - o(1)$ , all of the following events occur:

1. The common vote  $V^{\text{root}}[\tau_n].\text{total}$  exceeds  $7n\sqrt{\log_2 n}$ ;
2. For any  $i$ , the next  $i$  votes have sum  $o(n\sqrt{\log_2 n})$ ;
3. The vote  $U^{\text{root}}[t].\text{total}$  observed by any process differs from  $V^{\text{root}}[t].\text{total}$  by at most  $6n\sqrt{\log_2 n}$ .

If this occurs, then every process observes, for some  $t$ ,  $U^{\text{root}}[t].\text{total} \geq 7n\sqrt{\log_2 n} - 6n\sqrt{\log_2 n} - o(n\sqrt{\log n}) > 0$ . In other words, all processes return the same value  $+1$  with constant probability for sufficiently large  $n$ . By symmetry, the same is true for  $-1$ . We have therefore constructed a weak shared coin with constant agreement probability.

**Theorem 1.** *Algorithm 1 implements a weak shared coin with constant bias, message complexity  $O(n^2 \log^2 n)$ , and with a bound of  $O(n \log^3 n)$  on the number of messages sent and received by any one process.*

*Proof.* We have just shown that Algorithm 1 implements a weak shared coin with constant bias, and from Lemma 3 we know that the maximum number of votes generated by any single process is  $O(n \log^2 n)$ . Because each process communicates with a subtree of  $2^h$  other processes every  $2^{-h}$  votes, each level of the tree contributes  $\Theta(1)$  amortized outgoing messages and incoming responses per vote, for a total of  $\Theta(\log n)$  messages per vote, or  $O(n \log^3 n)$  messages altogether.

In addition, we must count messages received and sent by a process  $p$  as part of the max register implementation. Here for each process  $q$  in  $p$ ’s level- $h$  subtree,  $p$  may incur  $O(1)$  messages every  $2^h$  votes generated by  $q$ . Each such process  $q$  generates at most  $O(n \log^2 n)$  votes, and there are  $2^h$  such processes  $q$ . So  $p$  incurs a total of  $O(n \log^2 n)$  votes from its level- $h$  subtree. Summing over all  $\log n$  levels gives the same  $O(n \log^3 n)$  bound on messages as for max-register operations initiated by  $p$ .

This gives the final bound of  $O(n \log^3 n)$  messages per process. Applying the same reasoning to the total vote bound from Lemma 3 yields the bound of  $O(n^2 \log^2 n)$  on total message complexity.  $\square$

## 7 Consensus Protocol and Extension for General $t$

**Consensus.** We now describe how to convert a message-efficient weak shared coin into message-efficient consensus. We adapt a shared-memory consensus protocol, due to Chandra [18], which, like many shared-memory consensus protocols has the **early binding** property identified by Aguilera and Toueg [2] as necessary to ensure correctness of a consensus protocol using a weak shared coin.

Chandra’s protocol uses two arrays of bits to track the speed of processes with preference 0 or 1. The mechanism of the protocol is similar to previous protocols of Chor, Israeli, and Li [20] and Aspnes and

Herlihy [9]: if a process observes that the other team has advanced beyond it, it adopts that value, and if it observes that all processes with different preferences are two or more rounds behind, it decides on its current preference secure in the knowledge that they will switch sides before they catch up. The arrays of bits effectively function as a max register, so it is natural to replace them with two max registers  $m[0]$  and  $m[1]$ , initially set to 0, implemented as in Section 4. We obtain Algorithm 2, whose complete description, pseudocode, and proof are given in Section B of the Appendix.

**Theorem 2.** *Let  $\text{SharedCoin}_r$ , for each  $r$ , be a shared coin protocol with constant agreement parameter, individual message complexity  $T_1(n)$ , and total message complexity  $T(n)$ . Then Algorithm 2 implements a consensus protocol with expected individual message complexity  $O(T_1(n) + n)$  and total message complexity  $O(T(n) + n^2)$ .*

**Extension for General  $t$ .** We can decrease the message complexity of the protocol by taking advantage of values of  $t = o(n)$ . The basic idea is to reduce message complexity by “deputizing” a set of  $2t + 1$  processes to run the consensus protocol described in Section B and produce an output value, which they broadcast to all other participants. For this, we fix processes  $p_1, \dots, p_{2t+1}$  to be the group of processes running the consensus protocol, which we call the *deputies*. When executing an instance of the protocol, each process first sends a *Start* message to the deputies. If the process is a deputy, it waits to receive *Start* notifications from  $n - t$  processes. Upon receiving these notifications, the process runs the consensus algorithm described above, where the only participants are processes  $p_1, \dots, p_{2t+1}$ . Upon completing this protocol, each deputy broadcasts a  $\langle \text{Result}, \text{value} \rangle$  message to all processes, and returns the decided value. If the process is not a deputy, then it simply waits for a *Result* message from one of the deputies, and returns the corresponding value. Correctness follows from the previous arguments.

**Theorem 3.** *Let  $n, t > 0$  be parameters such that  $t < n$ . The algorithm described above implements randomized consensus using  $O(nt + t^2 \log^2 t)$  expected total messages, and  $O(n + t \log^3 t)$  expected messages per process.*

## 8 Conclusions and Future Work

We have addressed the problem of consensus in the classic asynchronous message-passing model. Besides theoretical importance, our work is also motivated by the recent surge in practical interest for asynchronous message-passing primitives for reducing contention both at the operating system level [13], and at the application level, e.g. [15, 24]. In particular, we note that practitioners have already proposed implementations with (weaker) agreement-like semantics in the context of distributed transactional memory [15].

We have described a randomized algorithm for consensus with expected message complexity  $O(n^2 \log^2 n)$  that tolerates  $t < n/2$  crash faults; this algorithm also has the desirable property that each process sends and receives expected  $O(n \log^3 n)$  messages on average, and message size is logarithmic. We also present a generalization that uses expected  $O(nt + t^2 \log^2 t)$  messages.

Two conspicuous open problems remain. The first is whether we can close the remaining poly-logarithmic gap for the message cost of consensus in the classic model. Second, can we use techniques from this paper to help close the gap for message-cost of Byzantine agreement in the classic model? To the best of our knowledge, the current lower bound for message cost of Byzantine agreement is  $\Omega(n^2)$ , while the best upper bound is  $O(n^{6.5})$  — a significant gap.

## References

- [1] Karl Abrahamson. On achieving consensus using a shared memory. In *Proceedings of the Seventh Annual ACM Symposium on Principles of Distributed Computing*, PODC '88, pages 291–302, New York, NY, USA, 1988. ACM.

- [2] Marcos K. Aguilera and Sam Toueg. The correctness proof of Ben-Or’s randomized consensus algorithm. *Distributed Computing*, 25(5):371–381, 2012.
- [3] James Aspnes. Lower bounds for distributed coin-flipping and randomized consensus. *Journal of the ACM*, 45(3):415–450, May 1998.
- [4] James Aspnes. Randomized protocols for asynchronous consensus. *Distributed Computing*, 16(2–3):165–175, September 2003.
- [5] James Aspnes, Hagit Attiya, and Keren Censor. Randomized consensus in expected  $O(n \log n)$  individual work. In *PODC ’08: Proceedings of the Twenty-Seventh ACM Symposium on Principles of Distributed Computing*, pages 325–334, August 2008.
- [6] James Aspnes, Hagit Attiya, and Keren Censor-Hillel. Polylogarithmic concurrent data structures from monotone circuits. *J. ACM*, 59(1):2, 2012.
- [7] James Aspnes and Keren Censor. Approximate shared-memory counting despite a strong adversary. *ACM Transactions on Algorithms*, 6(2):1–23, 2010.
- [8] James Aspnes and Keren Censor-Hillel. Atomic snapshots in  $O(\log^3 n)$  steps using randomized helping. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC 2013)*, pages 254–268. 2013. Full version available at: <http://cs-www.cs.yale.edu/homes/aspnes/papers/disc2013-journal.pdf>.
- [9] James Aspnes and Maurice Herlihy. Fast randomized consensus using shared memory. *Journal of Algorithms*, 11(3):441–461, September 1990.
- [10] James Aspnes and Orli Waarts. Randomized consensus in expected  $o(n \log^2 n)$  operations per processor. *SIAM Journal on Computing*, 25(5):1024–1044, October 1996.
- [11] Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. Sharing memory robustly in message-passing systems. *J. ACM*, 42(1):124–142, January 1995.
- [12] Hagit Attiya and Keren Censor. Tight bounds for asynchronous randomized consensus. *J. ACM*, 55(5):20:1–20:26, November 2008.
- [13] Andrew Baumann, Paul Barham, Pierre-Evariste Dagand, Tim Harris, Rebecca Isaacs, Simon Peter, Timothy Roscoe, Adrian Schüpbach, and Akhilesh Singhanian. The multikernel: A new os architecture for scalable multicore systems. In *Proceedings of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles, SOSP ’09*, pages 29–44, New York, NY, USA, 2009. ACM.
- [14] 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, PODC ’83*, pages 27–30, New York, NY, USA, 1983. ACM.
- [15] Annette Bieniusa and Thomas Fuhrmann. Consistency in hindsight: A fully decentralized stm algorithm. In *IPDPS*, pages 1–12. IEEE, 2010.
- [16] Gabriel Bracha. An asynchronous  $[(n - 1)/3]$ -resilient consensus protocol. In *PODC ’84: Proceedings of the third annual ACM symposium on Principles of distributed computing*, pages 154–162, New York, NY, USA, 1984. ACM.
- [17] Gabriel Bracha and Ophir Rachman. Randomized consensus in expected  $O(n^2 \log n)$  operations. In Sam Toueg, Paul G. Spirakis, and Lefteris M. Kirousis, editors, *WDAG*, volume 579 of *Lecture Notes in Computer Science*, pages 143–150. Springer, 1991.

- [18] Tushar Deepak Chandra. Polylog randomized wait-free consensus. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, pages 166–175, Philadelphia, Pennsylvania, USA, 23–26 May 1996.
- [19] Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, March 1996.
- [20] Benny Chor, Amos Israeli, and Ming Li. On processor coordination using asynchronous hardware. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing*, PODC '87, pages 86–97, New York, NY, USA, 1987. ACM.
- [21] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, April 1988.
- [22] Faith Fich, Maurice Herlihy, and Nir Shavit. On the space complexity of randomized synchronization. *J. ACM*, 45(5):843–862, September 1998.
- [23] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985.
- [24] Vincent Gramoli, Rachid Guerraoui, and Vasileios Trigonakis. Tm2c: A software transactional memory for many-cores. In *Proceedings of the 7th ACM European Conference on Computer Systems*, EuroSys '12, pages 351–364, New York, NY, USA, 2012. ACM.
- [25] P. Hall and C.C. Heyde. *Martingale Limit Theory and Its Application*. Academic Press, 1980.
- [26] A. Karlin and A. Yao. Probabilistic lower bounds for byzantine agreement and clock synchronization. Unpublished manuscript.
- [27] Ram Keralapura, Graham Cormode, and Jeyashankher Ramamirtham. Communication-efficient distributed monitoring of thresholded counts. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, SIGMOD '06, pages 289–300, New York, NY, USA, 2006. ACM.
- [28] Valerie King and Jared Saia. Byzantine agreement in polynomial expected time. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2013.
- [29] Valerie King and Jared Saia. Faster agreement via a spectral method for detecting malicious behavior. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.
- [30] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982.
- [31] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, April 1980.
- [32] M. Saks, N. Shavit, , and H. Woll. Optimal time randomized consensus - making resilient algorithms fast in practice. In *Proc. of the 2nd ACM Symposium on Discrete Algorithms (SODA)*, pages 351–362, 1991.




Figure 1: Structure of the cohorts. A process maintains count and sum estimates for each cohort it is a member of.

## A Complete Algorithm Analysis

In this section, we prove correctness of the weak shared coin algorithm in Section 5. We first analyze some of the properties of the tree-based vote counting structure. For simplicity, we assume that the number of processes  $n$  is a power of two.

**Vote Propagation.** The algorithm is based on the idea that, as processes take steps, counter values for the cohorts get increased, until, eventually, the root counter value surpasses the threshold, and processes start to return. We first provide a way of associating a set of generated votes to each counter value.

We say that a process  $p_i$  counts a number  $x_i$  of (consecutive) locally-generated votes to node  $s$  if, after generating the last such vote, process  $p_i$  updates the max register at  $s$  in its loop iteration. See Figure 1 for an illustration. We prove that this procedure has the following property:

**Lemma 7.** *Consider a subtree rooted at node  $s$  with  $\ell$  leaves, and let  $x_1, x_2, \dots, x_\ell$  be the number of votes most recently counted by member processes  $q_1, q_2, \dots, q_\ell$  at node  $s$ , respectively. Then the value of the count component of the max register at  $s$  is at least  $\sum_{m=1}^{\ell} x_m$ .*

*Proof.* We proceed by induction on the height of the subtree rooted at  $s$ , proving the following claim: Given the subtree rooted at  $s$  with corresponding processes  $q_1, q_2, \dots, q_\ell$ , with  $x_1, x_2, \dots, x_\ell$  being the number of votes most recently counted by each process at node  $s$ , there exists a process  $q_i$  with  $1 \leq i \leq \ell$  that performs a MaxUpdate with value at least  $\sum_{m=1}^{\ell} x_m$  on the max register at node  $s$  when counting votes  $x_i$  at  $s$ .

If  $s$  is a leaf, the claim is straightforward, since each process writes the vote count, total, and variance at its leaf after generating the vote. (We can regard the single-writer register at the leaf as a max register.) For the induction step, let  $v$  be the left child of  $s$ , and  $r$  be the right child. By the step, there exists a process  $p_i$  which updates the left child max register to  $\sum_{m=1}^{\ell/2} x_m$ , and a process  $p_j$  which updates the right child max register to  $\sum_{m=\ell/2+1}^{\ell} x_m$ .

We now focus on the value of the max register at  $s$  after both processes  $p_i$  and  $p_j$  complete their MaxUpdate operations at  $s$ . Crucially, notice that, since both processes first read the max register at  $v$ , then the one at  $r$ , one of the operations must necessarily “see” the value the other wrote at the corresponding child. Therefore, one of the processes, say  $p_i$ , must write a value of at least  $\sum_{m=1}^{\ell} x_m$  to the max register at  $s$ , as claimed.  $\square$

**Proof Strategy.** To show that the algorithm works, we need to show (a) that the total vote  $U^t$  observed in the root max register at any time  $t$  is close to the actual total generated vote  $V^t$ ; (b) that when the threshold  $S$  is crossed, the total **common votes**  $|V^t|$  is likely to be far away from 0; and (c) the combination of  $U^{t'} - V^{t'}$  and any subsequent **extra votes**  $V^{t'} - V^t$  observed by a particular process at time  $t'$  will not change the sign of the total vote. The first involves a detailed analysis of what votes may be omitted from the visible total combined with a version of the Azuma-Hoeffding inequality taken from [10]; the second requires use of a martingale Central Limit Theorem [25, Theorem 3.2]; the last follows from an application of Kolmogorov’s inequality.

**Lemma 8.** Fix some execution of Algorithm 1, and let  $n_i$  be the number of votes generated by process  $p_i$  during this execution. Then

$$\sum_{i=1}^n \sum_{j=1}^{n_i} w_j^2 \leq \frac{K + 2n^2}{1 - 8n/T} = O(n^2 \log n). \quad (3)$$

*Proof.* Examining the algorithm, we see that every  $n$  votes, process  $p_i$  both propagates all of its votes to the root max register and checks the total variance against  $K$ . It follows that (a) the root max register includes all but at most  $n$  votes for each process, and (b) each process generates at most  $n$  additional votes after  $\text{var}_{\text{root}}$  passes  $K$ . Since each of these  $2n$  missing and extra votes from process  $p_i$  has weight bounded by  $w_{n_i}$ , we get a bound

$$\begin{aligned} \sum_{i=1}^n \sum_{j=1}^{n_i} w_j^2 &\leq K + \sum_{i=1}^n 2nw_{n_i}^2 \\ &\leq K + \sum_{i=1}^n 2n \left( 1 + \frac{4}{T} \sum_{j=1}^{n_i} w_j^2 \right) \\ &= K + 2n^2 + \frac{8n}{T} \sum_{i=1}^n \sum_{j=1}^{n_i} w_j^2, \end{aligned}$$

which, rearranged, gives (3). □

We can combine the preceding lemmas to get a bound on the maximum weight of any vote.

**Lemma 9.** For any execution of Algorithm 1 and any process  $p_i$ ,

$$w_{n_i} \leq \sqrt{1 + \frac{4K + 8n^2}{T - 8n}} = O(\sqrt{n}). \quad (4)$$

*Proof.* We have

$$\begin{aligned} w_{n_i}^2 &\leq 1 + \frac{4}{T} \sum_{j=1}^{n_i} w_j^2 \\ &\leq 1 + \frac{4}{T} \cdot \frac{K + 2n^2}{1 - 8n/T} \\ &= 1 + \frac{4K + 8n^2}{T - 8n}. \end{aligned}$$

Now take the square root of both sides. □

**Corollary 1.** For any execution of Algorithm 1, for all  $i$ ,  $n_i = O(n \log^2 n)$ .

*Proof.* The maximum number of votes generated by any single process is constrained by

$$w_{n_i} = 2^{\lfloor (n_i - 1)/T \rfloor} = O(\sqrt{n})$$

giving

$$n_i = O(T \log(\sqrt{n})) = O(n \log^2 n). \quad \square$$

A similar argument to that used in Lemma 9 gives a bound on the sum, over all processes, of the square of the maximum vote. This will turn out to be useful later, and also gives better bound on the total number of generated votes than simply multiplying the bound in Corollary 1 by  $n$ .

**Lemma 10.** *For any execution of Algorithm 1,*

$$\sum_i w_{n_i}^2 \leq n + \frac{4K + 8n^2}{T - 8n} = O(n). \quad (5)$$

*Proof.* We have

$$\begin{aligned} \sum_i w_{n_i}^2 &\leq \sum_i \left( 1 + \frac{4}{T} \sum_{j=1}^{n_i} w_j^2 \right) \\ &\leq n + \frac{4}{T} \cdot \frac{K + 2n^2}{1 - 8n/T} \\ &= n + \frac{4K + 8n^2}{T - 8n}. \end{aligned}$$

□

**Corollary 2.** *In any execution of Algorithm 1,*

$$\sum_i n_i = O(n^2 \log n).$$

*Proof.* From Lemma 10, we have

$$\sum_i w_{n_i}^2 = \sum_i 4^{\lfloor (n_i-1)/T \rfloor} \leq an,$$

for sufficiently large  $n$  and some constant  $a$ .

The floor is inconvenient, so we remove it by taking

$$\sum_i 4^{(n_i-1)/T-1} \leq \sum_i 4^{\lfloor (n_i-1)/T \rfloor} \leq an.$$

Convexity of  $4^{(x-1)/T-1}$  implies

$$\begin{aligned} n \cdot 4^{(\frac{1}{n} \sum_i n_i - 1)/T-1} &\leq \sum_i 4^{(n_i-1)/T-1} \\ &\leq an, \end{aligned}$$

which gives

$$\begin{aligned} \sum_i n_i &\leq n(T \log_4(4a) + 1) \\ &= O(n^2 \log n). \end{aligned}$$

□



For any adversary strategy  $\mathcal{A}$ , let  $\tau_{\mathcal{A}}$  be a stopping time corresponding to the first time  $t$  at which  $U^{\text{root}}[t].\text{var} \geq K$ . We will use a martingale Central Limit Theorem to show that  $V^{\text{root}}[\tau_{\mathcal{A}}].\text{total}$  converges to a normal distribution as  $n$  grows, when suitably scaled. This will then be used to show that all processes observe a population of common votes whose total is likely to be far from zero.

In particular, we use the following simplified version of [25, Theorem 3.2], taken from [7]. The notation  $X \xrightarrow{p} Y$  means that  $X$  converges in probability to  $Y$ , and  $Y \xrightarrow{d} Y$  means that  $X$  converges in distribution to  $Y$ .

**Lemma 11.** *Let  $\{S_{mt}, \mathcal{F}_{mt}\}$ ,  $1 \leq t \leq k_m$ ,  $m \geq 1$  be a martingale array where for each fixed  $m$ ,  $\{S_{mt}, \mathcal{F}_{mt}\}$  is a zero-mean martingale with difference sequence  $X_{mt} = S_{mt} - S_{m,t-1}$ . If*

1.  $\max_t |X_{mt}| \xrightarrow{p} 0$ ,
2.  $\sum_t X_{mt}^2 \xrightarrow{p} 1$ ,
3.  $E[\max_t X_{mt}^2]$  is bounded in  $m$ , and
4.  $\mathcal{F}_{m,t} \subseteq \mathcal{F}_{m+1,t}$  for  $1 \leq t \leq k_m$ ,  $m \geq 1$ ,

then  $S_{mt} \xrightarrow{d} N(0, 1)$ , where  $N(0, 1)$  has a normal distribution with zero mean and unit variance.

**Lemma 12.** *Let  $\{\mathcal{A}_n\}$  be a family of adversary strategies, one for each number of processes  $n \in \mathbb{N}$ . Let  $\tau_n = \tau_{\mathcal{A}_n}$  be as above. Then*

$$\frac{V^{\text{root}}[\tau_n].\text{total}}{\sqrt{K}} \xrightarrow{d} N(0, 1). \quad (6)$$

*Proof.* Let  $Y_{nt}$  be the  $t$ -th vote generated in the execution governed by  $\mathcal{A}_n$ , and let  $X_{nt} = Y_{nt}/\sqrt{K}$  if  $t \leq \tau_n$  and 0 otherwise. Then

$$\frac{V^{\text{root}}[\tau_n].\text{total}}{\sqrt{K}} = \sum_{i=1}^{\tau_n} Y_{ni}/\sqrt{K} = \sum_{i=1}^{\infty} X_{ni} = \lim_{t \rightarrow \infty} S_{nt}.$$

So to prove (6), we need only verify that the random variables  $X_{nt}$  satisfy the conditions of Lemma 11 for an appropriate choice of  $\sigma$ -algebras  $\mathcal{F}_{n,t}$ . That each  $X_{nt}$  has zero mean is immediate from their definition. Again following the approach of [7], we will let  $\mathcal{F}_{n,t}$  be generated by the signs of the first  $t$  votes, and will assume that these signs are consistent across all hypothetical executions for each  $n$ , giving  $\mathcal{F}_{n,t} = \mathcal{F}_{n',t}$  for all  $n, n'$ . This trivially satisfies the last condition.

For the first condition, from Lemma 9 we have

$$\max_t |X_{nt}| \leq \max_i \frac{w_{n_i}}{\sqrt{K}} \leq \sqrt{\frac{1}{K} + \frac{4 + 8n^2/K}{T - 8n}} = \sqrt{\frac{1}{n^2 \log_2 n} + \frac{4 + 8/\log_2 n}{4n \log_2 n - n}}.$$

Since the denominator dominates the numerator in both fractions under the square root, this converges to zero always.

For the second condition, observe that  $K \sum_t X_{nt}^2$  is the sum of the squares of the weights of all votes generated by  $\tau_n$ . This quantity is at least  $K$  (as otherwise  $U^{\text{root}}[\tau_n].\text{var}$  could not have reached  $K$ ), and is at most  $\frac{K+2n^2}{1-8n/T}$  by Lemma 8. Dividing by  $K$  shows

$$1 \leq \sum_t X_{nt}^2 \leq \frac{1 + 2n^2/K}{1 - 8n/T} = O\left(\frac{1 + O(1/\log n)}{1 - O(1/\log n)}\right)$$

which converges to 1 always.

Finally, for the third condition, we have already shown that  $\max_t |X_{nt}|$  converges to 0, and since  $X_{nt}^2 = |X_{nt}|^2$  for all  $n, t$ , we have that  $X_{nt}^2$  also converges to 0 for any fixed  $t$ ; boundedness follows.  $\square$

For each subtree  $s$  and time  $t_s$ , let  $D^s[t_s] = V^s[t_s].\text{total} - U^s[t_s].\text{total}$  be the difference between the generated votes in  $s$  at time  $t_s$  and the votes reported to the max register corresponding to  $s$  at time  $t_s$ . Let  $s_\ell$  and  $s_r$  be the left and right subtrees of  $s$ , and let  $t_{s_\ell}$  and  $t_{s_r}$  be the times at which the values added to produce  $U^s[t_s].\text{total}$  were read from these subtrees.

Then

$$\begin{aligned} D^s[t_s] &= V^s[t_s].\text{total} - U^s[t_s].\text{total} \\ &= (V^{s_\ell}[t_s].\text{total} + V^{s_r}[t_s].\text{total}) - (U^{s_\ell}[t_{s_\ell}].\text{total} + U^{s_r}[t_{s_r}].\text{total}) \\ &= (V^{s_\ell}[t_s].\text{total} + V^{s_r}[t_s].\text{total}) - (V^{s_\ell}[t_{s_\ell}].\text{total} + V^{s_r}[t_{s_r}].\text{total}) + (D^{s_\ell}[t_{s_\ell}] + D^{s_r}[t_{s_r}]) \\ &= (V^{s_\ell}[t_s].\text{total} - V^{s_\ell}[t_{s_\ell}].\text{total}) + (V^{s_r}[t_s].\text{total} - V^{s_r}[t_{s_r}].\text{total}) + D^{s_\ell}[t_{s_\ell}] + D^{s_r}[t_{s_r}]. \end{aligned}$$

Iterating this recurrence gives

$$D^s[t_s] = \sum_{s'} \left( V^{s'}[t_{\text{parent}(s')}].\text{total} - V^{s'}[t_{s'}].\text{total} \right),$$

where  $s'$  ranges over all proper subtrees of  $s$ .

To bound this sum, we consider each (horizontal) layer of the tree separately, and observe that the missing interval of votes  $\left( V^{s'}[t_{\text{parent}(s')}].\text{total} - V^{s'}[t_{s'}].\text{total} \right)$  for each subtree  $s$  in layer  $h$  consists of at most  $2^h$  votes by each of at most  $2^h$  processes. For each process  $p_i$  individually, the variance of its  $2^h$  heaviest votes, using Lemma 8, is at most  $2^h \left( 1 + (4/T) \sum_{j=1}^{n_i} w_j^2 \right)$ , so if we sum the total variance of at most  $2^h$  votes from all processes, we get at most

$$2^h \left( n^2 + (4/T) \sum_{i=1}^n \sum_{j=1}^{n_i} w_j^2 \right) \leq 2^h \left( n^2 + \frac{K + 2n^2}{1 - 8n/T} \right),$$

using Lemma 8.

We would like to use this bound on the total variance across all missing intervals to show that the sum of the total votes across all missing intervals is not too large. Intuitively, if we can apply a bound to the total variance on a particular interval, we expect the Azuma-Hoeffding inequality to do this for us. But there is a complication in that the total variance for an interval may depend in a complicated way on the actions taken by the adversary during the interval. So instead, we attack the problem indirectly, by adopting a different characterization of the relevant intervals of votes and letting the adversary choose between them to obtain the actual intervals that contributed to  $D^{\text{root}}[t]$ .

We will use the following extended version of the classic Azuma-Hoeffding inequality from [10]:

**Lemma 13** ([10, Theorem 4.5]). *Let  $\{S_0, \mathcal{F}_0\}$ ,  $0 \leq i \leq n$  be a zero-mean martingale with difference sequence  $\{X_i\}$ . Let  $w_i$  be measurable  $\mathcal{F}_{i-1}$ , and suppose that for all  $i$ ,  $|X_i| \leq w_i$  with probability 1; and that there exists a bound  $W$  such that  $\sum_{i=1}^n w_i^2 \leq W$  with probability 1. Then for any  $\lambda > 0$ ,*

$$\Pr[S_n \geq \lambda] \leq e^{-\lambda^2/2W}. \quad (7)$$

Fix an adversary strategy. For each subtree  $s$ , let  $X_1^s, X_2^s, \dots$  be the sequence of votes generated in  $s$ . For each  $s, t$ , and  $W$ , let  $Y_i^{tsW} = X_i^s$  if (a) at least  $t$  votes have been generated by all processes before  $X_i^s$  is generated, and (b)  $\sum_{j < i} (Y_j^{tsW})^2 + (X_i^s)^2 \leq W$ ; otherwise let  $Y_i^{tsW}$  be 0. If we let  $\mathcal{F}_i$  be generated by all votes preceding  $X_i^s$ , then the events (a) and (b) are measurable  $\mathcal{F}_i$ , so  $\{Y_i^{tsW}, \mathcal{F}_i\}$  forms a martingale. Furthermore, since only the sign of  $Y_i^{tsW}$  is unpredictable, we can define  $w_i = (Y_i^{tsW})^2$  for the purposes of Lemma 13, and from (b) we have  $\sum w_i^2 \leq W$  always. It follows that, for any  $c > 0$ ,

$$\Pr \left[ \sum_i Y_i^{tsW} \geq \sqrt{2cW \ln n} \right] \leq e^{-c \ln n} = n^{-c}.$$

There are polynomially many choices for  $t$ ,  $s$ , and  $W$ ; taking a union bound over all such choices shows that, for  $c$  sufficiently large, with high probability  $\sum_i Y_i^{tsW}$  is bounded by  $\sqrt{2cW \ln n}$  for all such intervals.

We can use this to show:

**Lemma 14.** *For any adversary strategy and sufficiently large  $n$ , with probability  $1 - o(1)$ , it holds that at all times  $t$ ,*

$$|V^{\text{root}}[t].\text{total} - U^{\text{root}}[t].\text{total}| \leq 6n\sqrt{\log_2 n}.$$

*Proof.* We are trying to bound  $D^{\text{root}}[t] = \sum_s (V^s[t_{\text{parent}(s)}] - V^s[t_s])$ , where  $s$  ranges over all proper subtrees of the tree and for each  $s$  of size  $2^h$ , the interval  $(t_s, t_{\text{parent}(s)})$  includes at most  $2^h$  votes for each process.

Suppose that for each  $t, s, W$ , it holds that  $Y^{tsW} \leq \sqrt{9W \ln n}$ . By the preceding argument, each such event fails with probability at most  $n^{-9/2}$ . There are  $O(n^2 \log n)$  choices for  $t$ ,  $O(n)$  choices for  $s$ , and  $O(n^2 \log n)$  choices for  $W$ , so taking a union bound over all choices of  $Y^{tsW}$  not occurring shows that this event occurs with probability  $O(n^{-1/2} \log^2 n) = o(1)$ .

If all  $Y^{tsW}$  are bounded, then it holds deterministically that

$$\begin{aligned} \sum_s (V^s[t_{\text{parent}(s)}] - V^s[t_s]) &= \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} (V^s[t_{\text{parent}(s)}] - V^s[t_s]) \\ &= \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} Y^{tsW_s} \\ &\leq \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} \sqrt{2cW_s \ln n} \\ &= \sqrt{2c \ln n} \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} \sqrt{W_s}, \end{aligned}$$

where  $W_s$  is the total variance of the votes generated by  $s$  in the interval  $(t_s, t_{\text{parent}(s)})$ .

Note that this inequality does not depend on analyzing the interaction between voting and when processes read and write the max registers. For the purposes of computing the total offset we are effectively allowing the adversary to choose what intervals it includes retrospectively, after carrying out whatever strategy it likes for maximizing the probability that any particular values  $Y^{tsW}$  are too big.

Because each process  $i$  in  $s$  generates at most  $2^h$  votes, and each such vote has variance at most  $w_{n_i}^2$ , we have

$$W_s \leq 2^h \sum_{i \in s} w_{n_i}^2.$$

Furthermore, the subtrees at any fixed level  $h$  partition the set of processes, so applying Lemma 10 gives

$$\sum_{s, |s|=2^h} W_s \leq \sum_{s, |s|=2^h} 2^h \sum_{i \in s} w_{n_i}^2 = 2^h \sum_{s, |s|=2^h} \sum_{i \in s} w_{n_i}^2 = 2^h \sum_i w_{n_i}^2 \leq 2^h \left( n + \frac{4K + 8n^2}{T - 8n} \right).$$

By concavity of square root,  $\sum \sqrt{x_i}$  is maximized for non-negative  $x_i$  constrained by a fixed bound on  $\sum x_i$  by setting all  $x_i$  equal. Setting all  $n/2^h$  values  $W_s$  equal gives

$$\begin{aligned} W_s &\leq \frac{2^h}{n} \cdot 2^h \left( n + \frac{4K + 8n^2}{T - 8n} \right) \\ &= 2^{2h} \left( 1 + \frac{4K + 8n^2}{Tn - 8n^2} \right) \end{aligned}$$

and thus

$$\sqrt{W_s} \leq 2^h \sqrt{1 + \frac{4K + 8n^2}{Tn - 8n^2}}$$

which gives the bound

$$\begin{aligned} \sqrt{9 \ln n} \sum_{h=0}^{\log_2 n - 1} \sum_{s, |s|=2^h} \sqrt{W_s} &\leq \sqrt{9 \ln n} \sum_{h=0}^{\log_2 n - 1} 2^h \sqrt{1 + \frac{4K + 8n^2}{Tn - 8n^2}} \\ &= \sqrt{9 \ln n} \left(1 + \frac{4K + 8n^2}{Tn - 8n^2}\right)^{\log_2 n - 1} \sum_{h=0}^{\log_2 n - 1} 2^h \\ &= 3(n-1) \sqrt{\ln n} \sqrt{1 + \frac{4K + 8n^2}{Tn - 8n^2}} \\ &= 3(n-1) \sqrt{\ln n} \sqrt{1 + \frac{\log_2 n + 2}{\log_2 n - 2}} \\ &\leq 6n \sqrt{\log_2 n}, \end{aligned}$$

when  $n$  is sufficiently large.

The last step uses the fact that for  $K = n^2 \log n$  and  $T = 4n \log n$ , the value under the radical converges to 2 in the limit, and  $3\sqrt{2/\ln 2} < 6$ .  $\square$

For the last step, we need to show that the **extra votes** that arrive after  $K$  variance has been accumulated are not enough to push  $V^{\text{root}}[t]$  close to the origin. For this, we use Kolmogorov's inequality, a martingale analogue to Chebyshev's inequality, which says that if we are given a zero-mean martingale  $\{S_i, \mathcal{F}_i\}$  with bounded variance, then

$$\Pr[\exists i \leq n : |S_i| \geq \lambda] \leq \frac{\lambda^2}{\text{Var}[S_n]}.$$

Consider the martingale  $S_1, S_2, \dots$  where  $S_i$  is the sum of the first  $i$  votes after  $V^{\text{root}}[i]$ .var first passes  $K$ . Then from Lemma 8,

$$\text{Var}[S_i] \leq \frac{K + 2n^2}{1 - 8n/T} - K = \frac{8K(n/T) + 2n^2}{1 - 8n/T} = \frac{8n^2 + 2n^2}{1 - 2/(\log_2 n)} = O(n^2).$$

So for any fixed  $c$ , the probability that  $|S_i|$  exceeds  $cK$  for any  $i$  is  $O(1/\log n) = o(1)$ .

Putting these results together, from Lemma 12, we have that the total common vote  $V^{\text{root}}[\tau_n]$ .total converges in distribution to  $N(0, 1)$  when scaled by  $\sqrt{K} = n\sqrt{\log_2 n}$ . In particular, for any fixed constant  $c$ , there is a constant probability  $p_c > 0$  that for sufficiently large  $n$ ,  $\Pr[V^{\text{root}}[\tau_n] \geq cn\sqrt{\log_2 n}] \geq p_c$ .

Let  $c$  be 7. Then with probability  $p_7 - o(1)$ , all of the following events occur:

1. The common vote  $V^{\text{root}}[\tau_n]$ .total exceeds  $7n\sqrt{\log_2 n}$ ;
2. For any  $i$ , the next  $i$  votes have sum  $o(n\sqrt{\log_2 n})$ ;
3. The vote  $U^{\text{root}}[t]$ .total observed by any process differs from  $V^{\text{root}}[t]$ .total by at most  $6n\sqrt{\log_2 n}$ .

If this occurs, then every process observes, for some  $t$ ,  $U^{\text{root}}[t]$ .total  $\geq 7n\sqrt{\log_2 n} - 6n\sqrt{\log_2 n} - o(n\sqrt{\log n}) > 0$ . In other words, all processes return the same value +1 with constant probability for sufficiently large  $n$ . By symmetry, the same is true for  $-1$ . We have constructed a weak shared coin with constant agreement probability.

**Theorem 4.** *Algorithm 1 implements a weak shared coin with constant bias, message complexity  $O(n^2 \log^2 n)$ , and with a bound of  $O(n \log^3 n)$  on the number of messages sent and received by any one process.*

*Proof.* We have just shown that Algorithm 1 implements a weak shared coin with constant bias, and from Corollary 1 we know that the maximum number of votes generated by any single process is  $O(n \log^2 n)$ . Because each process communicates with a subtree of  $2^h$  other processes every  $2^{-h}$  votes, each level of the tree contributes  $\Theta(1)$  amortized outgoing messages and incoming responses per vote, for a total of  $\Theta(\log n)$  messages per vote, or  $O(n \log^3 n)$  messages altogether.

In addition, we must count messages received and sent by a process  $p$  as part of the max register implementation. Here for each process  $q$  in  $p$ 's level- $h$  subtree,  $p$  may incur  $O(1)$  messages every  $2^h$  votes done by  $q$ . Each such process  $q$  performs at most  $O(n \log^2 n)$  votes, and there are  $2^h$  such  $q$ , so  $p$  incurs a total of  $O(n \log^2 n)$  votes from its level- $h$  subtree. Summing over all  $\log n$  levels gives the same  $O(n \log^3 n)$  bound on messages as for max-register operations initiated by  $p$ .

This gives the bound of  $O(n \log^3 n)$  messages per process. Applying the same reasoning to the bound on the total number of votes from Corollary 2 gives the bound of  $O(n^2 \log^2 n)$  on total message complexity.  $\square$

## B From a Shared Coin to Consensus

For the last step, we must show how to convert a message-efficient weak shared coin into message-efficient consensus. Typically, a shared coin is used in the context of a larger framework in which faster processes do not use the coin (to ensure validity) while slower processes use the shared coin (to ensure agreement after each round with constant probability). A direct approach is to drop the shared coin into an existing message-passing agreement protocol like Ben-Or's [14]. However, as recently observed by Aguilera and Toueg [2], this can produce bad outcomes when the failure bound is large, because the adversary may be able to predict the value of the shared coin before committing to the value of the faster processes. Instead, we adapt a shared-memory consensus protocol, due to Chandra [18], which, like many shared-memory consensus protocols has the **early binding** property identified by Aguilera and Toueg as necessary to avoid this problem.

Chandra's protocol uses two long arrays of bits to track the speed of processes with preference 0 or 1. The mechanism of the protocol is similar to previous protocols of Chor, Israeli, and Li [20] and Aspnes and Herlihy [9]: if a process observes that the other team has advanced beyond it, it adopts that value, and if it observes that all processes with different preferences are two or more rounds behind, it decides on its current preference secure in the knowledge that they will switch sides before they catch up. The arrays of bits effectively function as a max register, so it is natural to replace them with two max registers  $m[0]$  and  $m[1]$ , initially set to 0, implemented as in Section 4. Pseudocode for the resulting algorithm is given in Algorithm 2.<sup>3</sup>

**Theorem 5.** *Let  $\text{SharedCoin}_r$ , for each  $r$ , be a shared coin protocol with constant agreement parameter, individual message complexity  $T_1(n)$ , and total message complexity  $T(n)$ . Then Algorithm 2 implements a consensus protocol with expected individual message complexity  $O(T_1(n) + n)$  and total message complexity  $O(T(n) + n^2)$ .*

*Proof.* Validity follows from the fact that, if all processes start with the same preference  $p$ , then  $m[1 - p]$  never advances beyond zero and no process ever changes its preference.

Agreement is more subtle, although the ability to read the max register atomically makes the proof easier than the original proof in [18]. Suppose that some process returns  $p$ . Then this process previously read  $m[1 - p] \leq r - 2$  after writing  $m[p] \geq r$ . At the time of this read, no other process has yet written  $r - 1$  or greater to  $m[1 - p]$ , so any other process with preference  $1 - p$  has not yet executed the write at the start of round  $r - 1$ . It follows that each such process sees  $m[p] \geq (r - 1) + 1$  and adopts  $p$  as its new

<sup>3</sup>The presentation here is influenced by the simplified version of Chandra's original protocol given in [4].

```

1  $p \leftarrow$  input
2 for  $r \leftarrow 1 \dots \infty$  do
3   WriteMax( $m[p], r$ )
4    $r' \leftarrow$  ReadMax( $m[1 - p]$ )
5   if  $r' \geq r + 1$  then
6     |  $p' \leftarrow 1 - p$ 
7   else if  $r' = r$  then
8     |  $p' \leftarrow$  SharedCoin $_r$ ()
9   else if  $r' = r - 1$  then
10    |  $p' \leftarrow p$ 
11  else
12    | return  $p$ 
13   $r'' \leftarrow$  ReadMax( $m[p]$ )
14  if  $r'' < r + 1$  then
15    |  $p \leftarrow p'$ 

```

**Algorithm 2:** Consensus protocol

preference before reaching round  $r$ . Since all processes entering round  $r$  agree on  $p$ , they all decide  $p$  after at most one more round.

For termination, it cannot be the case that processes with different preferences both avoid the shared coin in round  $r$ , because processes with both preferences  $p$  would need to read  $m[1 - p] < r$  after writing  $m[p] = r$ , violating the specification of the max register. Suppose that at least one process with preference  $p$  does not execute the shared coin in round  $r$ , and processes with preference  $1 - p$  do. For this to occur, the adversary must show the preference- $p$  process  $m[1 - p] \leq r - 1$  before any preference- $(1 - p)$  process writes  $r$  to  $m[1 - p]$ , and thus before any process (with either preference) initiates SharedCoin $_r$ . It follows that SharedCoin $_r$  returns  $p$  to all processes with constant probability, producing agreement in round  $r$  and termination after at most two more rounds.

For a shared coin protocol with agreement parameter  $\delta$ , we have a probability of at least  $\delta$  per round that the protocol reaches agreement in that round. This gives an expected number of rounds of at most  $1/\delta + 2 = O(1)$ . Each involves at most three max-register operations for each process, imposing a cost of  $O(n)$  messages on that process and  $O(1)$  messages on each other process, plus an invocation of the shared coin. The expected individual message complexity is thus  $O(n) + T_1(n) = O(T_1(n) + n)$  as claimed. A similar argument shows the bound for total message complexity.  $\square$

Using Algorithm 1 for the shared coin,  $T(n)$  dominates, giving that each process sends and receives  $O(n \log^3 n)$  messages on average.

## C Variant for General Failure Parameter $t$

In this section, we describe a variant of the consensus protocol which adapts to the number of failures  $t$ . The basic idea is to reduce message complexity by “deputizing” a set of  $2t + 1$  processes to run the consensus protocol described in Section B and produce an output value, which they broadcast to all other participants.

**Description.** For this, we fix processes  $p_1, \dots, p_{2t+1}$  to be the group of processes running the consensus protocol, which we call the *deputies*. When executing an instance of the protocol, each process first sends a *Start* message to the deputies. If the process is a deputy, it waits to receive *Start* notifications from  $n - t$  processes. Upon receiving these notifications, the process runs the algorithm from Section B, where the only participants are processes  $p_1, \dots, p_{2t+1}$ . Upon completing this protocol, each deputy broadcasts a  $\langle \text{Result}, \text{value} \rangle$  message to all processes, and returns the decided value. If the process is not a deputy, then it simply waits for a *Result* message from one of the deputies, and returns the corresponding value.

**Analysis.** This construction guarantees the following properties. The proof follows from the argument in Section B.

**Theorem 6.** *Let  $n, t > 0$  be parameters such that  $t < n$ . The algorithm described above implements randomized consensus using  $O(nt + t^2 \log^3 t)$  expected total messages, and  $O(n + t \log^3 t)$  expected messages per process.*