

A Simple and Efficient Randomized Byzantine Agreement Algorithm*

Benny Chor

Brian A. Coan

MIT

August 22, 1984

Abstract: A new randomized Byzantine agreement algorithm is presented. This algorithm operates in a synchronous system of n processors, at most t of which can fail. The algorithm reaches agreement in $O(t/\log n)$ expected rounds and $O(n^2t/\log n)$ expected message bits independent of the distribution of processor failures. This performance is further improved to a constant expected number of rounds and $O(n^2)$ message bits if the distribution of processor failures is assumed to be uniform. In either event, the algorithm improves on the known lower bound on rounds for deterministic algorithms. Some other advantages of the algorithm are that it requires no cryptographic techniques, that the amount of local computation is small, and that the expected number of random bits used per processor is only one. It is argued that in many practical applications of Byzantine agreement, the randomized algorithm of this paper achieves superior performance.

Keywords: Byzantine agreement, fault tolerance, and randomized algorithms.

© 1984 Massachusetts Institute of Technology

*A preliminary version of this paper will appear in the Proceedings of the Fourth IEEE Symposium on Reliability in Distributed Software and Database Systems, 1984 [CC]. This research was supported by the NSF under Grants #DCR-8302391 and #MCS-8006938, by the U. S. Army Research Office under Contracts #DAAG29-79-C-0155 and #DAAG29-84-K-0058, and by the Advanced Research Projects Agency of the Department of Defense under Contract #N00014-83-K-0125.

1. Introduction

Byzantine agreement is the problem of reaching a consensus in a distributed system of n processors, at most t of which may fail in arbitrary, even malicious, ways. In a randomized Byzantine agreement algorithm each processor can use the results of local random coin tosses in the course of reaching agreement. In this paper we present a synchronous randomized Byzantine agreement algorithm that terminates in an expected $O(t/\log n)$ rounds and that works for any $n \geq 3t+1$. This contrasts with the lower bound of $t+1$ rounds for deterministic algorithms that was shown by Lynch and Fischer [LF]. Our algorithm is of significance because it is simple and efficient enough to be of practical use, because it performs better than any possible deterministic algorithm, because it operates for practical values of n and t , and because it does not require any cryptographic techniques.

Ben-Or [Be], Bracha and Toueg [BT], and Rabin [Ra] each have investigated the randomized Byzantine agreement problem. Their primary interest was asynchronous algorithms, but simple variants of their algorithms operate in the synchronous case. Because the topic of this paper is synchronous algorithms, we will discuss only the synchronous variants of their algorithms. A full comparison of the algorithms appears in Section 11; however, we give some advantages of our algorithm here. Our algorithm improves on Ben-Or's in that it operates efficiently for practical ratios of n and t , and it improves on Rabin's in that it does not require initial distribution of coin tosses by a trusted, failure-free dealer. In the synchronous case, the algorithm due to Bracha and Toueg is similar to the one due to Ben-Or. For simplicity, we will only make comparisons with Ben-Or's algorithm.

Bracha [Br] has recently devised a new randomized Byzantine agreement algorithm that terminates in an expected $O(\log n)$ rounds. His algorithm assumes a model of computation in which cryptographic techniques may be used to conceal information from malicious faulty processors. This is different from the model that we assume in this paper.

In evaluating a Byzantine agreement algorithm, it is often useful to consider the total number of processors that is required by the system in order to tolerate t processor failures. This motivates us to define the *redundancy* of a system of n processors as $(n-1)/t$. Lamport, Shostak, and Pease [LSP] have shown that no deterministic non-authenticated algorithm is possible unless the redundancy is at least 3. An easy extension of their proof shows that this same amount of redundancy is required for randomized algorithms. It is usually desirable to minimize the redundancy in a system in order to reduce the cost of hardware. It is always possible to increase the redundancy in a system and maintain correct operation of a Byzantine agreement algorithm; however, some of the algorithms that we will consider have lower bounds on redundancy that are higher than 3. The basic algorithm that we will present operates for any redundancy of 3 or more.

If we postulate the existence of a global reliable random coin toss, Byzantine agreement can be reached in a constant expected number of rounds using techniques developed by Ben-Or and by Rabin. Unfortunately, such a source of random coin tosses is unlikely to be available in a distributed system subject to Byzantine faults. A central component of the known randomized Byzantine agreement algorithms is the simulation of a global reliable coin toss using other methods. Ben-Or and Rabin each have techniques for producing global random coin tosses. Ben-Or uses a technique which works well when the redundancy is high but which has a low probability of producing a good toss when the redundancy is low. In his technique, each processor tosses a coin independently and broadcasts the result. If one result predominates by a sufficient margin then a coin toss has been produced; otherwise, no usable coin toss has been produced. A spread of sufficient size is likely only when the total number of processors is large relative to the number of faults. As a consequence, his algorithm either requires a large (exponential) number of rounds or a high amount of redundancy (equal to t). Rabin produces global coin tosses efficiently; however, he assumes a different, more powerful model of computation. His coin tosses are precomputed by a trusted dealer who splits the results of each coin toss so that $t + 1$ processors can determine the result, but t processors

have no information. The distributed, split coin tosses are an expensive resource that is partially consumed at each execution of the algorithm. In some applications, it may be unrealistic to assume that a trusted dealer exists. Rabin's algorithm is a good choice in those applications where his model seems realistic. The algorithm which we will describe will generate global coin tosses more efficiently than Ben-Or's algorithm but uses none of the extra machinery of Rabin's algorithm. In particular, no reliable, trusted dealer is required. One way in which our technique is inferior to theirs is that ours works only in a synchronous system while theirs work in both synchronous and asynchronous systems. There is something inherent in our technique that does not seem to generalize to the asynchronous case.

A contribution of this paper is our new technique for generating random coin tosses that (while they are not perfect) are of sufficient quality to permit our algorithm to make rapid progress. We now outline our technique for tossing coins. At each round, a small group of g processors is assigned the task of coin tossing. Each processor in this group tosses its own coin and broadcasts the result. The coin toss generated by this group is the majority of individual outcomes. If more than half the processors are faulty, they can bias the coin toss any way they want or cause some processors to see heads and others to see tails. But, if fewer than half are faulty, there is a sufficiently large probability that all correct coin tossers will have the same outcome (provided g is not too big). If they all happen to have the same outcome, the majority is determined regardless of what the faulty processors do. With no more than t faulty processors, there can not be more than $2t/g$ disjoint groups with a majority of faulty processors. After at most that many rounds, a group of coin tossers with a majority of correct processors will be reached. We will show that this leads to fast termination of the algorithm.

We now give a brief outline of the remainder of the paper. In Section 2 we define the synchronous randomized Byzantine agreement problem. In Section 3 we present the basic algorithm. In Section 4 we analyze the performance of the basic algorithm. In Section 5 we present an alternative

algorithm that doubles the speed of the basic algorithm at the cost of requiring twice the amount of redundancy. In Section 6 we explain how the performance of our algorithm improves if the actual number of faults is smaller than the bound on the number of faults, t . In Section 7 we show that it is possible to achieve coordinated termination (all processors decide in the same round) with high probability. In Section 8 we present an alternative analysis of the basic algorithm for the case where processor failures are assumed to be uniformly randomly distributed. In Section 9 we discuss the problem of reintegrating repaired processors into the algorithm. In Section 10 we explore an application of our algorithm—speeding up the distributed lottery agreement protocol of Broder and Dolev [BD]. In Section 11 we evaluate the basic algorithm by comparing it with the alternatives.

2. The Problem

A synchronous randomized Byzantine agreement algorithm is run by a distributed system of n processors, at most t of which may fail. The computation proceeds by the sending and receiving of messages. Message exchange takes place in a series of rounds over a network that is fully connected and reliable. At each round, a processor can toss coins as part of its computation. These coin tosses affect message generation and processor state. Correct processors toss fair coins and send messages according to their programs. Failed processors can send arbitrary messages.

Each processor starts the algorithm with an input value, v , from a fixed set of legal inputs, V . The goal is that after sending some messages each processor will produce an answer. There are two requirements on the answer. The *agreement* condition is that all correct processors produce the same answer. The *validity* condition is that if all correct processors start the algorithm with the same value, then this value will be the answer produced by the correct processors.

3. The Basic Algorithm

For simplicity of presentation, the algorithm given here is binary (reaches agreement on one bit). It can easily be extended to be multivalued (reach agreement on arbitrary values) using the

technique of Turpin and Coan [TC].

First, we give an informal description of the algorithm, then we give the code. The algorithm is organized as a series of epochs of message exchange. Each epoch consists of two rounds. The round structure is provided automatically by the synchronous communications network. In the presentation of the algorithm, epoch and round numbers are shown as the first two components of each message. These should be viewed as implicit. They are shown only to make it easier to discuss the algorithm.

The algorithm is parameterized by g , the group size; n , the number of processors; and t , the number of faults tolerated. It is assumed that $n \geq 3t + 1$. The parameter g is used to determine the number of processors in each group of coin tossers. The processors are divided into a maximal number of disjoint groups of g processors each. Any processors that are left over belong to no group. The groups are numbered from 1 to $\lfloor n/g \rfloor$.

In each epoch, the processors cooperate to perform a distributed coin toss. In epoch e , the group whose index is congruent to e modulo $\lfloor n/g \rfloor$ actively performs the coin toss. All processors attempt to observe the result. Each processor in the active group tosses a coin and broadcasts the result. A processor calculates the coin toss of the group as the majority of the coin tosses it receives from processors in the group. The value that processor P observes for a coin toss is defined to be this majority value seen by P . If a group contains a large number of faulty processors, then they can control the outcome of the toss or cause correct processors to observe inconsistent values. We will, however, show that for suitably selected group size, g , a large enough number of groups will toss sufficiently random coins for our purposes.

We describe the algorithm for the processor P . (All processors run the same code.) The variable `CURRENT` holds the value that processor P currently favors as the answer of the Byzantine agreement algorithm. At the start of the algorithm `CURRENT` is set to processor P 's input value. In the first round of each epoch, processor P broadcasts `CURRENT`. Based on the round 1 messages

received, processor P changes `CURRENT`. If it sees at least $n-t$ round 1 messages for some particular value, then it assigns that value to `CURRENT`; otherwise, it assigns the distinguished value “?” to `CURRENT`. In the second round of each epoch, processor P broadcasts `CURRENT` and (if required) the result of a coin toss. Based on the round 2 messages received, processor P either changes `CURRENT` or decides on an answer and exits the algorithm. Let `ANS` be the most frequent value (other than “?”) in round 2 messages received by P . Let `NUM` be the number of such messages. There are three cases depending on the value of `NUM`. If $\text{NUM} \geq n - t$ then processor P decides on the value `ANS` and exits the algorithm. If $n - t > \text{NUM} \geq t + 1$ then processor P assigns the value `ANS` to the variable `CURRENT` and continues the algorithm. If $t + 1 > \text{NUM}$ then processor P assigns the coin toss it received in the current round to the variable `CURRENT`. The selected coin toss is the one broadcast by the processor group whose index number is congruent to the current epoch number modulo $\lfloor n/g \rfloor$, the number of processor groups.

Code for processor P with parameters g , n , and t :

```
1. procedure BYZANTINE_AGREEMENT(INPUT):
2.   CURRENT  $\leftarrow$  INPUT
3.   for  $e \leftarrow 1$  to  $\infty$  do
4.     broadcast ( $e, 1, \text{CURRENT}$ )
5.     receive ( $e, 1, *$ ) messages
6.     if for some  $v$  there are  $\geq n - t$  messages ( $e, 1, v$ )
7.       then CURRENT  $\leftarrow v$ 
8.       else CURRENT  $\leftarrow$  “?”
9.     if  $\text{GROUP}(P) \equiv e \pmod{\lfloor n/g \rfloor}$ 
10.      then TOSS  $\leftarrow$  TOSS_COIN()
11.      else TOSS  $\leftarrow$  0
12.     broadcast ( $e, 2, \text{CURRENT}, \text{TOSS}$ )
13.     receive ( $e, 2, *, *$ ) messages
14.     ANS  $\leftarrow$  the value  $v \neq \text{“?”}$  such that ( $e, 2, v, *$ ) messages are most frequent
15.     NUM  $\leftarrow$  number of occurrences of ( $e, 2, \text{ANS}, *$ ) messages
16.     if  $\text{NUM} \geq n - t$  then decide ANS
17.     elseif  $\text{NUM} \geq t + 1$  then CURRENT  $\leftarrow$  ANS
18.     else CURRENT  $\leftarrow$  majority toss from processor group  $x$  where  $x \equiv e \pmod{\lfloor n/g \rfloor}$ 
```

We make several remarks about the algorithm. `GROUP` is a procedure that takes as input a processor identity and returns the processor’s group number. `TOSS_COIN` is a procedure that takes

no argument and returns the result of a random coin toss (0 or 1). In message descriptions, “*” is a wild-card character that matches anything. In order to allow correct processors to stop sending after they decide, we adopt the convention that a processor that sends no message is assumed to vote for the value in the last message that it sent. Therefore, a processor that has decided may stop sending messages after the first round in which it broadcasts its decided value. This will either be the round in which it decides or the next following round. Other invalid or missing messages present no difficulty. A processor that sends such a message is faulty and therefore could have sent anything. The recipient of such a message may correctly replace it by any valid message.

Define *value* as a legal input to the algorithm, either 0 or 1. Specifically, “?” is not a value.

Lemma 1: *During each epoch, at most one value is sent in round 2 (step 12) messages by correct processors.*

Proof: Assume at least one value is sent in round 2 of epoch e . Say processor P sends value v . P has seen at least $n - t$ messages $(e, 1, v)$. At least $n - 2t$ of them are from correct processors and are therefore reliably sent to all processors in the system. All processors have at least $n - 2t$ messages $(e, 1, v)$ in a total of n messages. This leaves at most $2t < n - t$ messages for some other value. This is not enough to cause a correct processor to send a round 2 message for a value other than v . □

In Theorem 2 we prove that our algorithm never produces a wrong answer and we prove that in each epoch there is at least one coin-toss value that will terminate the algorithm. The analysis of the expected running time of the algorithm follows in Section 4.

Theorem 2: *The algorithm has the following three properties.*

Validity: *If value v is distributed as input to all correct processors, then all correct processors decide v in round 2 of epoch 1.*

Agreement: *Let e be the first epoch in which a correct processor decides. If correct processor P decides v in epoch e then by round 2 of epoch $e + 1$ all correct processors decide v .*

Termination: In any epoch, e , there is at least one value which (if it is adopted by all processors executing the assignment in step 18) will cause all correct processors to decide by round 2 of epoch $e + 1$.

Proof: We show that the algorithm satisfies the three conditions.

Validity: Assume that value v is distributed as input to all correct processors. All (at least $n - t$) correct processors broadcast v in rounds 1 and 2 of epoch 1. All correct processors assign v to ANS, set NUM to a value at least $n - t$, and therefore decide v in round 2 of epoch 1.

Agreement: For processor P to decide v in epoch e , it must be the case that P has seen at least $n - t$ messages $(e, 2, v, *)$. At least $n - 2t \geq t + 1$ of the messages are from correct processors and are therefore reliably sent to all processors in the system. No correct processor sends round 2 messages for any value other than v (by Lemma 1). Only failed processors send such messages. Because there are at most t failed processors, there are at most t round 2 messages for any value $v' \neq v$. All correct processors assign v to ANS, set NUM to a value at least $t + 1$, and either decide v in the current epoch or begin the following epoch with the variable CURRENT equal to v . All correct processors will therefore decide v by round 2 of epoch $e + 1$.

Termination: By Lemma 1, there is at most one value, say v , that is broadcast by correct processors in round 1 of epoch e . Any correct processor that does not execute the assignment in step 18 must have ANS = v because it is not possible to get more than t votes for any other value. The value v (if adopted by all processors executing the assignment in step 18 of epoch e) will cause all correct processors to decide v by round 2 of epoch $e + 1$. This is because all correct processors will start epoch $e + 1$ with the value v assigned to the variable CURRENT. \square

Based on the termination property shown in Theorem 2, we make the following definitions: Define a *good* value in epoch e as one which (if it is adopted by all processors executing the assignment in step 18) will cause all correct processors to decide by round 2 of epoch $e + 1$. Define

a *good* coin toss in epoch e as one which distributes in epoch e a good value to all correct processors. Define a *bad* value or coin toss as one which is not good.

We are now in a position to explain why randomization is necessary for the correct operation of our algorithm. The termination part of Theorem 2 guarantees that in any epoch, say e , there will be at least one good value. This value can, however, be determined by the failed processors in round 1 of epoch e . If the coin toss were replaced with some predetermined value, then the faulty processors could always cause the good value to differ from this predetermined value. In our scheme, this strategy is unavailable to the faulty processors because the coin toss is not performed until round 2 of epoch e —after the good value has already been fixed.

4. Analysis of the Algorithm

In this section we analyze the computational resources used by the basic algorithm. At the system level, we calculate the expected number of rounds needed to reach agreement and the expected number of message bits sent. At the processor level, we calculate the amount of internal memory, the number of computation steps, and the expected number of random bits used by each processor. In particular, we show that the expected number of rounds to reach agreement is $O(t/\log n)$ and the expected number of random bits used by each processor is bounded above by 1. The analysis is worst case in the sense that it allows faulty processors to behave maliciously and holds for any distribution of them, as long as the redundancy is at least 3.

We define some terminology and notation. Let $g = 2m + 1$ denote the number of processors in each group. (This relation between g and m holds for the rest of the paper.) Let c_e be the coin toss generated at epoch e . We say that c_e is *fair* if each correct processor observes the same coin toss and this coin toss is unbiased and independent of previous coin tosses. Denote by p_e the probability that the coin toss c_e is good (will cause termination), and let $q_e = 1 - p_e$.

If there are at most m faulty processors among the coin tossers of epoch e (a majority of the coin tossers are correct), then

$$\frac{1}{2^m} \leq \Pr(c_e \text{ fair}).$$

By Theorem 2 and the definition of good coin toss,

$$\Pr(c_e \text{ good} \mid c_e \text{ fair}) = \frac{1}{2},$$

So, the conditional probability that c_e is good given that at most m coin tossers are faulty, is at least $1/2^{m+1}$.

Theorem 3: For group size $g = \log n$, the expected number of rounds to reach agreement, r , is bounded by

$$r < \frac{4t}{\log n} + 2\sqrt{n} + o(1) = O\left(\frac{t}{\log n}\right).$$

Proof: By the definition of a good coin toss, agreement is reached by all correct processors at most one epoch after a good coin toss is achieved. Every epoch consists of two rounds. Therefore, if Exp denotes the expected number of epochs until a good coin toss is achieved, then

$$r = 2 \cdot Exp + 2.$$

The expected number of epochs to get a good coin toss is

$$Exp = \sum_{e=1}^{\infty} e \cdot \Pr(c_e \text{ is the first good coin toss})$$

Since fair coin tosses are independent, we get

$$\Pr(c_e \text{ is the first good coin toss}) = q_1 q_2 \cdots q_{e-1} (1 - q_e),$$

so

$$\begin{aligned}
 Exp &= \sum_{e=1}^{\infty} e \cdot q_1 q_2 \cdots q_{e-1} (1 - q_e) \\
 &= (1 - q_1) + \sum_{e=2}^{\infty} (e \cdot q_1 q_2 \cdots q_{e-1} - e \cdot q_1 q_2 \cdots q_e) \\
 &= 1 + \sum_{e=1}^{\infty} q_1 q_2 \cdots q_e.
 \end{aligned}$$

For any specific e , an adversary can *block* the coin toss generated at epoch e by assigning $m + 1$ faulty coin tossers, thus making $q_e = 1$. However, with a limited supply of faulty processors, such blocking can not be repeated indefinitely.

We start by analyzing the first $\lfloor n/g \rfloor$ coin tosses, which are performed by disjoint groups of processors. To do this we calculate an upper bound on $\sum_{e=1}^{\lfloor n/g \rfloor} q_1 q_2 \cdots q_e$. If $q_i < q_{i+1}$ for some i , then the sum can be strictly increased by exchanging q_i and q_{i+1} . Therefore, we assume that the sequence of q_i 's is non-increasing. There can be at most $\lfloor t/(m + 1) \rfloor$ groups with a majority of faulty processors. These groups are blocked (the probability of getting a bad coin toss is 1). After the blocked groups, all groups must have at least $m + 1$ good processors and therefore at most $1 - 1/2^{m+1}$ probability of producing a bad toss. We permit the adversary an infinite number of such groups and calculate

$$\sum_{e=1}^{\lfloor n/g \rfloor} q_1 q_2 \cdots q_e \leq \frac{t}{m + 1} + 2^{m+1}$$

After $\lfloor n/g \rfloor$ groups, we cycle back through the coin tossers, and so

$$\begin{aligned}
 \sum_{e=1}^{\infty} q_1 q_2 \cdots q_e &= \left(\sum_{e=1}^{\lfloor n/g \rfloor} q_1 q_2 \cdots q_e \right) \left(1 + q_1 q_2 \cdots q_{\lfloor n/g \rfloor} + (q_1 q_2 \cdots q_{\lfloor n/g \rfloor})^2 + \dots \right) \\
 &\leq \left(\frac{t}{m + 1} + 2^{m+1} \right) \left(1 + \sum_{i=1}^{\infty} (q_1 q_2 \cdots q_{\lfloor n/g \rfloor})^i \right)
 \end{aligned}$$

With t faulty processors, at most $\lfloor t/(m + 1) \rfloor$ groups of coin tossers can be blocked. Therefore

$$\begin{aligned}
 q_1 q_2 \cdots q_{\lfloor n/g \rfloor} &\leq \left(1 - \frac{1}{2^{m+1}} \right)^{\lfloor n/g \rfloor - \lfloor t/(m+1) \rfloor} \\
 &< \left(1 - \frac{1}{2^{m+1}} \right)^{(n-2t)/g}
 \end{aligned}$$

Using the bound $t < n/3$, we have

$$q_1 q_2 \cdots q_{\lfloor n/g \rfloor} < \left(1 - \frac{1}{2^{m+1}}\right)^{n/3g}.$$

Because $g = \log n$ and $m < \frac{1}{2} \log n$, we get ($e = 2.718+$ in the next two equations)

$$\left(1 - \frac{1}{2^{m+1}}\right)^{n/3g} < \left(1 - \frac{1}{\sqrt{n}}\right)^{n/3 \log n} < \left(\frac{1}{e}\right)^{\sqrt{n}/3 \log n}.$$

The last bound implies

$$\sum_{i=1}^{\infty} (q_1 q_2 \cdots q_{\lfloor n/g \rfloor})^i < 2 \cdot \left(\frac{1}{e}\right)^{\sqrt{n}/3 \log n},$$

and therefore

$$\left(\frac{t}{m+1} + 2^{m+1}\right) \sum_{i=1}^{\infty} (q_1 q_2 \cdots q_{\lfloor n/g \rfloor})^i = o(1).$$

Combining all these bounds together, we get

$$Exp < \frac{2t}{\log n} + \sqrt{n} + o(1) = O\left(\frac{t}{\log n}\right).$$

The expected number of rounds is therefore bounded above by $4t/\log n + 2\sqrt{n} + o(1)$. \square

We remark that by setting $m = (1 - \epsilon) \log n$, we have

$$\frac{t}{m+1} + 2^{m+1} = \frac{t}{(1 - \epsilon) \log n} + n^{1-\epsilon},$$

so the constant 1 in the $O(t/\log n)$ expression for the expected number of epochs can be achieved asymptotically. Our analysis is somewhat loose in that the adversary can reuse faulty processors. For explicit values of n , t , and g the exact value of Exp can be found by direct computation. In Table 1 we list some values of Exp for small, practical systems.

We calculate the expected number of message bits sent by the basic algorithm. Individual messages have a constant size. In each round there are $O(n^2)$ messages sent. By Theorem 3, the expected number of rounds is $O(t/\log n)$. Therefore, over the course of the algorithm, an expected $O(n^2 t/\log n)$ message bits are sent.

Expected Number of Tries to Produce a Good Coin Toss							
Processors	Faults	Group Size	$E(\text{Tosses})$	Processors	Faults	Group Size	$E(\text{Tosses})$
4	1	1	3.2	55	18	5	8.4
7	2	3	4.0	58	19	5	8.8
10	3	3	4.4	61	20	5	9.0
13	4	3	4.7	64	21	5	9.4
16	5	3	5.1	67	22	7	8.9
19	6	5	8.0	70	23	7	8.6
22	7	5	6.2	73	24	7	9.0
25	8	5	5.7	76	25	7	9.5
28	9	5	6.6	79	26	7	9.3
31	10	5	6.3	82	27	7	9.7
34	11	5	7.1	85	28	7	9.7
37	12	5	6.8	88	29	7	10.0
40	13	5	6.9	91	30	7	10.0
43	14	5	7.5	94	31	7	10.4
46	15	5	7.5	97	32	7	10.7
49	16	5	8.0	100	33	7	10.8
52	17	5	8.1	103	34	7	11.1

Table 1: Coin Tossing Efficiency

In randomized algorithms, the number of random bits used by each processor is important. Current physical devices for producing random bits are rather slow. If a large number of random bits are required, then pseudo-random number generators are often used. Plumstead [P] showed that the fast, linear congruence generators are not secure. After seeing a few outcomes, an adversary can predict the remaining tosses (thus allowing the faulty processors to block all future coin tosses). Secure pseudo-random number generators, based on cryptographic techniques, are known to exist under certain intractability assumptions (see [BM] and [ACGS]). However, they require a lot of computation, so we are better off if we can avoid using them altogether. A surprising result is the number of random bits used by our algorithm. The expected number of times we cycle through all groups of coin tossers is bounded above by 1. At each cycle one random bit is used by each processor. Therefore, the expected number of coins tossed by each processor is bounded above by 1. Slow physical generators are good enough then, and it is not necessary to resort to pseudo-random number generators.

The complexity of the internal computation is that of counting small numbers. The amount of internal space required by our algorithm is small. Only $\log n$ memory bits for counting NUM are required. Reintegration of faulty processors is easy because no long histories need be stored (see Section 9).

5. An Alternative Algorithm

Our basic algorithm is resilient to $t < n/3$ faults and uses two rounds per epoch. If the number of faulty processors, t , is bound by $t < n/6$, then one round per epoch suffices. Thus the expected number of rounds is cut by a factor of 2. The code for this case uses two thresholds to which NUM, the number of supporters for the current majority value, are compared. This two-threshold scheme is an adaptation of one by Rabin [Ra]. If NUM falls between the two thresholds, the coin toss is used to determine the value of the variable CURRENT in the next round. With the two thresholds further than t apart, one of the two possible outcomes of the coin toss is good. The generation of coin tosses is done in the same way as in our basic algorithm. This makes the analysis of the expected number of coin tosses needed to get a good one identical to the analysis in Section 4.

Code for processor P with parameters g , n , and t :

```
1.  procedure BYZANTINE_AGREEMENT(INPUT):
2.  CURRENT  $\leftarrow$  INPUT
3.  for  $e \leftarrow 1$  to  $\infty$  do
4.      if GROUP( $P$ )  $\equiv e \pmod{\lfloor n/g \rfloor}$ 
5.          then TOSS  $\leftarrow$  TOSS_COIN()
6.          else TOSS  $\leftarrow$  0
7.      broadcast ( $e$ , CURRENT, TOSS)
8.      receive ( $e$ , *, *) messages
9.      ANS  $\leftarrow$  the value  $v$  such that ( $e$ ,  $v$ , *) messages are most frequent
10.     NUM  $\leftarrow$  number of occurrences of ( $e$ , ANS, *) messages
11.     if NUM  $\geq n - t$  then decide ANS
12.     elseif NUM  $\geq n - 2t$  then CURRENT  $\leftarrow$  ANS
13.     elseif NUM  $< n - 3t$  then CURRENT  $\leftarrow$  0
14.     else
15.         COIN  $\leftarrow$  majority toss from processor group  $x$  where  $x \equiv e \pmod{\lfloor n/g \rfloor}$ 
16.         if COIN = 0 then CURRENT  $\leftarrow$  0
17.         elseif COIN = 1 then CURRENT  $\leftarrow$  ANS
```

6. Early Stopping

Our algorithm is resilient to t faults, but the actual number of faulty processors, f , might be smaller than the upper bound t . A desirable property of any Byzantine agreement algorithm is that agreement be reached early in this case. Dolev, Rieschuck, and Strong [DRS] have studied early termination for deterministic Byzantine agreement algorithms. From the analysis in Section 4, it follows that the expected number of rounds to reach agreement in the presence of f faults is bounded above by $4f/\log n + 2\sqrt{n} + o(1)$. Thus early stopping is automatically achieved. Furthermore, for the range $\sqrt{n} \log n \leq f < n/3$, agreement is reached in $O(f/\log n)$ rounds.

7. Coordinated Termination

One disadvantage of our algorithm is that even though all processors start the algorithm in exactly the same round, they might be off by one epoch when they terminate. In this section we show that a minor modification of the algorithm yields an almost certain coordinated termination, namely all correct processors halt at exactly the same round with overwhelming probability. This is done without violating the agreement, validity, and termination requirements, which will still be achieved with probability 1. The expected running time is changed only by a small multiplicative constant.

The modification is quite simple. We know that the expected number of epochs to reach agreement in the basic algorithm is at most $2t/\log n + \sqrt{n} + o(1)$ and that the tail probability for the number of epochs converges rapidly. For example, the probability that more than $3t/\log n$ epochs will be needed is no greater than $(1/e)^{\sqrt{n}/3 \log n}$. In the modified algorithm, each correct processor will just delay the transition to a “halt” state until epoch $3t/\log n$ (in case it made its decision before that epoch) and behave as before otherwise. This guarantees coordinated termination by epoch $3t/\log n$ with probability at least $1 - (1/e)^{\sqrt{n}/3 \log n}$. Greater confidence of coordinated termination can be achieved at the cost of more rounds.

It remains to be seen whether Byzantine agreement in $O(t/\log n)$ expected time and probability 1 of coordinated termination can be achieved. The deterministic lower bound implies only $t + 1$ worst case lower bound for any such algorithm.

8. Uniformly Distributed Processor Failures

In the previous analysis we have assumed that an adversary controls both the selection of which processors fail and the behavior of the failed processors. It does this in the way which will cause our algorithm the most difficulty. An alternative assumption is that the faulty processors are randomly distributed.

In the analysis in this section we assume that the distribution of processor faults is uniform. That is, each of the $\binom{n}{t}$ ways of distributing t faults among the n processors is equally likely. We retain the assumption that failed processors are under the control of the adversary and therefore behave in a way which will cause our algorithm the most difficulty. In Theorem 4 we show that for $g = 1$ and uniform processor faults our algorithm terminates in a constant expected number of rounds. Because $g = 1$, each group of coin tossers consists of a single processor.

Theorem 4: *If $g = 1$ and processor faults are uniformly distributed, then the expected number of rounds until the last correct processor decides is at most 8.*

Proof: The expected number of rounds is $2c + 2$ where c is the expected number of coin tosses to get a good toss. We show that the expected number of coin tosses is at most 3. The probability that the coin toss of a correct processor is good is at least $\frac{1}{2}$ (by Theorem 2 and the fairness of the coin toss). The probability that processor P is correct is at least $\frac{2}{3}$ because $n \geq 3t + 1$. Therefore, the probability that the coin toss of processor P is good is at least $\frac{1}{3}$. The conditional probability that the coin toss of processor P is good given that previous coin tosses have been bad is at least as great as the unconditional probability (at least $\frac{1}{3}$) because the previous bad coin tosses increase

the conditional likelihood that processor P is a correct processor. Therefore, the expected number of coin tosses is at most the mean of a geometric random variable with parameter $\frac{1}{3}$. This is 3. \square

We calculate the expected number of message bits sent by the algorithm. Individual messages have a constant size. On each round, there are $O(n^2)$ messages sent. By Theorem 4, the expected number of rounds is 8. Therefore, over the course of the algorithm, an expected $O(n^2)$ message bits are sent.

Termination in a constant expected number of rounds is an attractive feature of the algorithm; however, this feature is not unique to randomized algorithms. Under the same assumption of uniform randomly distributed processor failures, a deterministic algorithm due to Reischuk [Re] also terminates in a constant expected number of rounds.

9. Reintegration of Failed Processors

It is possible for a processor that fails and is subsequently repaired to rejoin our Byzantine agreement algorithm. We assume that such a processor loses its local memory and that it runs special recovery code after it is repaired. In order to rejoin the algorithm a processor needs to replace its lost state information. This is easy with our algorithm because the amount of state information is small. It is also important that correct processors do not record the identities of known faulty processors. Reintegrating failed processors permits the algorithm to tolerate a larger number of failures as long as at most t occur simultaneously. In this section we describe how a repaired processor rejoins the algorithm.

During an epoch, there are several times at which a repaired processor can rejoin the algorithm. We describe one. The repaired processor simply begins the epoch with step 13, receiving round 2 messages. A processor that fails and is subsequently repaired is counted as a failed processor only from the epoch in which it fails until two epochs after it recovers. After that it is considered to be a correct processor. The constraint is that at any time there are no more than t failed processors.

A processor that attempts to rejoin the algorithm after the correct processors have decided or as they are deciding may not see enough messages in order to decide. To solve this problem, we adopt the rule that a processor that sees $t + 1$ or more silent processors in the epoch after it attempts to rejoin the algorithm will conclude that a decision has been reached. It then broadcasts a query and decides on any value that it receives from at least $t + 1$ processors. It is easy to see that a processor that rejoins the algorithm as a decision is being reached can not be tricked into making a wrong decision. There will not be enough votes for any incorrect value because only faulty processors will cast such votes after a correct processor has decided.

10. Distributed Lottery Agreement

The goal of a distributed lottery agreement protocol is to produce an unbiased coin toss on which all correct processors agree. The *a priori* probabilities of both outcomes should be equal, regardless of the behavior of the faulty processors. Broder and Dolev [BD] present a cryptographic distributed lottery agreement protocol which is resilient to $t < n/2$ faults and takes $3t + 3$ rounds. They also show a $t + 1$ *worst case* lower bound for any distributed lottery agreement algorithm. In this section we show that for $t < n/3$, our Byzantine agreement algorithm can be used to speed up the Broder and Dolev distributed lottery agreement protocol, resulting in $O(t/\log n)$ expected time for agreeing on the value of the coin toss.

The protocol of Broder and Dolev is semi-deterministic in the sense that there is only an initial randomized stage (in which each processor tosses its own local coin, and picks a key for a public-key cryptosystem). After this initial stage, there are three phases. In each phase, every processor sends one value, and n deterministic Byzantine agreements on each of these values are processed concurrently. At the end of the third phase, the value of the coin toss is known to all processors. A critical requirement is that the second (resp. third) phase not start before the results of all n Byzantine agreements from the first (resp. second) phase are determined. Both encryption (in a

way which is crucial to the algorithm) and authentication are used. Authentication is needed only for reaching Byzantine agreement in the case that $n/3 \leq t$. If $t < n/3$, then authentication is not necessary.

Replacing deterministic Byzantine agreement by our probabilistic algorithm, we get a faster method for distributed lottery agreement. A property of our algorithm is that if any correct processor has decided on some value, this value is determined and will be agreed upon by all correct processors no later than the next epoch. To use our algorithm in the distributed lottery agreement, the following rule is imposed: After a processor decides on all n values of the current phase, it waits one epoch and only then sends its value for the next phase. By waiting one epoch, it is guaranteed that when one correct processor tries to start the new phase, all correct processors have already decided on all values of the last phase. If a processor does not receive any value from some sender one round after the processor has sent its own value, the processor assumes that sender is faulty, and uses some default value as its initial value in the agreement on that sender's value. The epoch counting is continued from the first epoch, so that all processor refer to the same group of coin tossers throughout the three phases. We have shown that our Byzantine agreement algorithm can be used as a subroutine in the Broder and Dolev protocol. To see that we really speed things up, use the rapid convergence of the tail probabilities. One can show that the expected time for finishing n concurrent Byzantine agreements is no greater than twice the expected time for one Byzantine agreement. Hence we get

Theorem 5: *Distributed lottery agreement can be done in expected time $O(t/\log n)$ for $t < n/3$.*

Theorem 5 shows that the worst case lower bound of Broder and Dolev can not be extended to the average case.

11. Evaluation

In this section we evaluate the extent to which our algorithm is of practical importance. We

do this by comparing our algorithm with three alternatives: Ben-Or's randomized synchronous algorithm, Rabin's randomized synchronous algorithm, and a good deterministic algorithm.

When the redundancy, r , of a system of processors is $\Omega(t)$, Ben-Or's algorithm terminates in a constant expected number of rounds. For practical systems, however, it is desirable to operate at a lower redundancy in order to minimize the cost of computer hardware. The ideal value is $r = 3$. Unfortunately, for any r that is $O(1)$, Ben-Or's algorithm requires an exponential number of rounds. Compared to Ben-Or's algorithm, ours is more efficient for practical amounts of system redundancy.

Rabin's algorithm terminates in a constant expected number of rounds; however, it requires more resources than our algorithm. In particular, it requires a trusted dealer that distributes random coin tosses before the start of the algorithm. The underlying mechanism is authentication and Shamir's [S] shared secret. If this cost seems small, then Rabin's algorithm would be the choice. On the other hand, if the cost seems high, then our algorithm would be the choice. We believe that in practical systems, it is often unrealistic to assume the existence of a trusted dealer.

The most practical deterministic algorithm is due to Lynch, Fischer, and Fowler [LFF]. There are trade-offs between their algorithm and ours. The principal advantages of the deterministic algorithm are that it uses a fixed number of rounds and that all processors decide at the same round. The principal advantages of our algorithm are that the expected number of rounds is small, that the expected number of message bits is small, and that only two rounds are required if the input to all processors is the same. Our algorithm can not ensure synchronous termination; however, using the techniques of Section 7, this property can be achieved with high probability.

We conclude that for a practical algorithm, one would choose either the deterministic algorithm of Lynch, Fischer, and Fowler or our randomized algorithm. The deterministic algorithm would be chosen if the extra synchronization that it provides were important to the particular application. Otherwise, our randomized algorithm would be chosen. If it seems realistic to assume that processor

faults are uniformly distributed, then the randomized algorithm is especially attractive because its expected running time is constant.

Acknowledgment

We would like to thank Noga Alon, Jim Burns, Oded Goldreich, Leonid Levin, Bruce Lindsay, Nancy Lynch, and William Weihl for many helpful discussions.

References

- [ACGS] W. Alexi, B. Chor, O. Goldreich, and C. Schnorr, "RSA/Rabin Bits Are $\frac{1}{2} + \frac{1}{\text{poly}(\log N)}$ Secure," *Proc. 25th Annual Symposium on Foundations of Computer Science* (1984), to appear.
- [BD] A. Broder and D. Dolev, "Flipping Coins in Many Pockets (Byzantine Agreement on Uniformly Random Values)," *Proc. 25th Annual Symposium on Foundations of Computer Science* (1984), to appear.
- [Be] M. Ben-Or, "Another Advantage of Free Choice: Completely Asynchronous Agreement Protocols," *Proc. 2nd Annual ACM Symposium on Principles of Distributed Computing* (1983).
- [BM] M. Blum and S. Micali, "How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits," *SIAM Journal on Computing*, to appear.
- [Br] G. Bracha, "A Randomized Byzantine Agreement Algorithm with an $O(\log n)$ Expected Rounds," manuscript (1984).
- [BT] G. Bracha and S. Toueg, "Resilient Consensus Protocols," *Proc. 2nd Annual ACM Symposium on Principles of Distributed Computing* (1983).
- [CC] B. Chor and B. Coan, "A Simple and Efficient Randomized Byzantine Agreement Algorithm," *Proc. 4th Symposium on Reliability in Distributed Software and Database Systems* (1984), to appear.

- [DRS] D. Dolev, R. Reischuk, and H. R. Strong, "Eventual is Earlier than Immediate," *Proc. 23rd Annual Symposium on Foundations of Computer Science* (1982).
- [LF] N. Lynch and M. Fischer, "A Lower Bound for the Time to Assure Interactive Consistency," *Information Processing Letters* 14(4) (1982).
- [LFF] N. Lynch, M. Fischer, and R. Fowler, "A Simple and Efficient Byzantine Generals Algorithm," *Proc. 2nd Symposium on Reliability in Distributed Software and Database Systems* (1982).
- [LSP] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems* 4(3) (1982).
- [P] J. Plumstead, "Inferring a Sequence Generated by a Linear Congruence," *Proc. 23rd Annual Symposium on Foundations of Computer Science* (1982).
- [Ra] M. Rabin, "Randomized Byzantine Generals," *Proc. 24th Annual Symposium on Foundations of Computer Science* (1983).
- [Re] R. Reischuk, "A New Solution for the Byzantine Generals Problem," IBM-RJ-3673 (1982).
- [S] A. Shamir, "How to Share a Secret," *Communications of the ACM* 22(11) (1979).
- [TC] R. Turpin and B. Coan, "Extending Binary Byzantine Agreement to Multivalued Byzantine Agreement," *Information Processing Letters* 18(2) (1984).