

A LOWER BOUND FOR THE TIME TO ASSURE INTERACTIVE CONSISTENCY

Michael J. FISCHER and Nancy A. LYNCH

Laboratorium for Computer Science, MIT, Cambridge, MA 02139, U.S.A.

Received 18 September 1981; revised version received 9 February 1982

1. Introduction

The problem of 'assuring interactive consistency' is defined in [7]. It is assumed that there are n isolated processors, of which at most m are faulty. The processors can communicate by means of two-party messages, using a medium which is reliable and of negligible delay. The sender of a message is always identifiable by the receiver. Each processor p has a private value $\sigma(p)$. The problem is to devise an algorithm that will allow each processor p to compute a value for each processor r , such that (a) if p and r are nonfaulty, then p computes r 's private value $\sigma(r)$, and (b) all the nonfaulty processors compute the same value for each processor r .

It is shown in [7] that if $n < 3m + 1$, then there is no algorithm which assures interactive consistency. On the other hand, if $n \geq 3m + 1$, then an algorithm does exist. The algorithm presented in [7] uses $m + 1$ rounds of communication, and thus can be said to require 'time' $m + 1$. An obvious question is whether fewer rounds of communication suffice to solve the problem.

In this paper, we answer this question in the negative. That is, we show that any algorithm which assures interactive consistency in the presence of m faulty processors requires at least $m + 1$ rounds of communication.

The remainder of the paper is organized as follows. Section 2 contains motivation for our formal model and problem statement. Section 3 contains the notation and definitions. Section 4 contains a reduction of our set of allowable algorithms to a more restrictive set of 'uniform' algo-

rithms. Section 5 contains a restatement of the relevant results of [7]. Section 6 contains our main lower bound result. Section 7 contains an important open question.

The reader is urged to read [4] and [5] for discussion of the practical importance of assuring interactive consistency, and [7] for additional results not immediately relevant to this paper. Other related papers are [2] and [3].

2. Motivation for the definitions

A general model for solving the interactive consistency problem might consist of n processors (automata) communicating by means of n^2 one-way 'communication channels'. Each channel can be formalized as a shared variable which can be modified by exactly one processor and read by exactly one processor. (Such a formalization can be carried out, for example, within the model of [6].) The variables which each processor can modify are called its 'out-channels', while the variables it can read are called its 'in-channels'.

Each processor p starts with an arbitrary private value $\sigma(p)$. Execution of the system proceeds in synchronous 'rounds'; at each round, the following two steps occur:

Step 1. Each nonfaulty processor writes values ('sends messages') derived from its state into all of its out-channels, while each faulty processor writes arbitrary values into all of its out-channels.

Step 2. Each processor reads the values from all of its in-channels.

After some specified number, k , of rounds, each processor p outputs a vector of values, one for each processor r . These outputs are required to satisfy conditions (a) and (b) stated in Section 1.

If the only complexity measure of interest is the number of rounds, then we can assume without loss of generality that the messages sent by each nonfaulty processor p on the first round are all exactly equal to its private value $\sigma(p)$, and that the messages sent by each nonfaulty processor on subsequent rounds are all exactly the set of messages received from all processors on the previous round. That is, if there is any correct k -round algorithm, then there is a correct k -round algorithm in which exactly the messages described above are sent. This is so because (i) it is clear that the given information is the maximum nontrivial information which could be sent, (ii) it does not hurt to send nonfaulty processors the maximum information, since they can derive any needed information from the given maximum information, and (iii) it does not hurt to send faulty processors the maximum information, since it is assumed that the faulty processors can send arbitrary messages in any case, i.e., they could 'guess' any missing information.

In such a maximum-information algorithm, the output vector of each processor p is simply a function of the set of all values received by p at all rounds of the computation. (So far, this reduction is as in [7].) In addition, if p is nonfaulty, then the set of messages received by p at all rounds of the computation is determined by the set of messages received by p at the last round (since p sends messages to itself at each round containing the information p received at earlier rounds). Since the correctness conditions only involve the outputs of nonfaulty processors, it suffices to formalize the output of p as a function of the set of messages received by p at the last round of communication only.

3. Notation and definitions

If A is any alphabet, $i, j \in \mathbb{N} \cup \{0\}$, $i \geq j$, we use $A^{i:j}$ to denote the set of strings of symbols in A , of length at least i and at most j .

Let P be the set of processors, $|P| = n$, and let m be an upper bound on the number of faulty processors. Fix V to be the domain of values on which the processors wish to reach agreement. Assume $\{0, 1\} \subseteq V$.

For any $k \in \mathbb{N} \cup \{0\}$, let \mathcal{Q}^k denote the set of mappings from P^k into V . (An element of \mathcal{Q}^k is intended to represent a set of messages which a processor could receive at the last round of a computation.)

A k -round algorithm A (for P) is a set $\{F_p: p \in P\}$ of functions, where $F_p: \mathcal{Q}^k \times P \rightarrow V$. A is uniform if $F_p = F_q$ for all $p, q \in P$.

A k -round scenario (for P with m faults) is a mapping $\sigma: P^{1:k+1} \rightarrow V$, such that $|T_\sigma| \geq n - m$, where T_σ (the set of truth-tellers) $= \{q \in P: \sigma(wqp) = \sigma(wq) \text{ for all } p \in P \text{ and all } w \in P^{0:k-1}\}$. Intuitively, $\sigma(p_1 p_2 \dots p_i)$ is intended to represent the value in V which p_{i-1} told p_i that p_{i-2} told p_{i-1} that ... that p_1 told p_2 was p_1 's private value; as a special case, $\sigma(p)$ represents p 's private value. (Note that this definition reverses the direction of the string arguments in the [7] definition.) Let L_σ (the set of liars) denote $P - T_\sigma$. If $\sigma: P^{1:k+1} \rightarrow V$ and $p \in P$, then p 's view of σ is the map $\sigma_p \in \mathcal{Q}^k$ given by $\sigma_p(w) = \sigma(wp)$. Let $\mathcal{V}_p^k = \{\sigma_p: \sigma \text{ is a } k\text{-round scenario (for } P \text{ with } m \text{ faults) and } p \in T_\sigma\}$. That is, \mathcal{V}_p^k is the set of possible views for p when p is a truth-teller in a k -round scenario.

Let $A = \{F_p: p \in P\}$ be a k -round algorithm. Then A assures interactive consistency (for P with m faults) provided that for each k -round scenario σ (for P with m faults) the following two conditions hold:

- (a) *Validity.* $F_p(\sigma_p, r) = \sigma(r)$ for all $p, r \in T_\sigma$,
- (b) *Agreement.* $F_p(\sigma_p, r) = F_q(\sigma_q, r)$ for all $p, q \in T_\sigma$ and all $r \in P$.

4. Reduction to uniform algorithms

In this section, we show that it suffices to restrict attention to uniform algorithms.

Lemma 1. Assume $n \geq 2m + 1$. Let $A = \{F_p: p \in P\}$ be a k -round algorithm which assures interactive consistency. Then $F_p(\alpha, r) = F_q(\alpha, r)$ for all $p, q, r \in P$ and all $\alpha \in \mathcal{V}_p^k \cap \mathcal{V}_q^k$.

Proof. Let $p, q, r \in P, \alpha \in \mathcal{V}_p^k \cap \mathcal{V}_q^k$. Then there are k -round scenarios σ and τ such that $p \in T_\sigma, q \in T_\tau$ and $\alpha = \sigma_p = \tau_q$. Let $s \in T_\sigma \cap T_\tau$. (Such an s is guaranteed to exist because $n \geq 2m + 1$.)

Modify only the last round of σ and τ to obtain new k -round scenarios σ' and τ' , as follows. Let $\sigma'(ws) = \sigma(wp)$ for all $w \in P^k$, and let $\sigma'(x) = \sigma(x)$ otherwise. Similarly, let $\tau'(wq) = \tau(wq)$ for all $w \in P^k$, and let $\tau'(x) = \tau(x)$ otherwise. It is easy to check that σ' and τ' are scenarios (since $T_{\sigma'} \supseteq T_\sigma$ and $T_{\tau'} \supseteq T_\tau$), that $\{p, s\} \subseteq T_{\sigma'}, \{q, s\} \subseteq T_{\tau'}$ and that $\alpha = \sigma'_p = \sigma'_s = \tau'_q = \tau'_s$. Thus, $F_p(\alpha, r) = F_s(\alpha, r) = F_q(\alpha, r)$, by the agreement property. \square

Theorem 2. Assume $n \geq 2m + 1$. If there is a k -round algorithm which assures interactive consistency, then there is a k -round uniform algorithm which assures interactive consistency.

Proof. Let $A = \{F_p : p \in P\}$ be a k -round algorithm which assures interactive consistency. Define $F : \mathcal{U}^k \times P \rightarrow V$ as follows. Let

$$F(\alpha, r) = \begin{cases} F_p(\alpha, r) & \text{if } \alpha \in \mathcal{V}_p^k, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 1 shows that this definition is consistent. Then the algorithm which uses F for all processors is a k -round algorithm which assures interactive consistency. \square

5. Earlier results

In this section, we state the two relevant results from [7].

Theorem 3. Assume $n < 3m + 1$. Then there is no algorithm which assures interactive consistency.

Proof. The proof can be found in [7]. \square

Theorem 4. Assume $n \geq 3m + 1$. Then there is an $(m + 1)$ -round uniform algorithm which assures interactive consistency.

Proof. The proof can be found in [7]. \square

6. Lower bound

In this section, we present our main result.

Theorem 5. If $k \leq m$, then there is no k -round algorithm which assures interactive consistency.

Proof. The theorem is easily seen to be true if $m = 0$, so assume that $m \geq 1$. Assume that the theorem is false: that $k \leq m$ and there is a k -round algorithm $A = \{F_p : p \in P\}$ which assures interactive consistency. By Theorem 3 we know that $n \geq 3m + 1$. By Theorem 2 we can assume that A is uniform, i.e., that $F_p = F$ for all $p \in P$.

Define a relation \sim on \mathcal{U}^k as follows: Let $\alpha \sim \beta$ provided there exists a k -round scenario σ and $p, q \in T_\sigma$ for which $\alpha = \sigma_p$ and $\beta = \sigma_q$. Let \equiv be the smallest equivalence relation containing \sim . By the agreement property, we have the following.

Fact 1. $F(\alpha, r) = F(\beta, r)$ for all $\alpha, \beta \in \mathcal{U}^k$ with $\alpha \equiv \beta$, and all $r \in P$.

For each $v \in V, w \in P^k$, let $\gamma_v(w) = v$. By the validity property, we have the following.

Fact 2. $F(\gamma_v, r) = v$ for all $r \in P$ and all $v \in V$.

Define an arbitrary total order on P , let $N = n^k$, and let $\ell : P^k \rightarrow \{1, \dots, N\}$ be a bijection corresponding to lexicographic order on the strings in P^k . That is, if $v, w \in P^k, 0 \leq i \leq k - 1, p, q \in P, v = r_1 \dots r_i p x, w = r_1 \dots r_i q y$ and $p < q$, then $\ell(v) < \ell(w)$.

For $1 \leq a \leq N + 1$, define $\alpha_a : P^k \rightarrow \{0, 1\}$ by

$$\alpha_a(w) = \begin{cases} 0 & \text{if } \ell(w) < a, \\ 1 & \text{otherwise.} \end{cases}$$

Note that $\alpha_1 = \gamma_1$ and $\alpha_{N+1} = \gamma_0$.

We claim that $\alpha_a \sim \alpha_{a+1}$ for all $a, 1 \leq a \leq N$. If so, then $\gamma_1 = \alpha_1 \sim \alpha_2 \sim \dots \sim \alpha_{N+1} = \gamma_0$, so that $\gamma_1 \equiv \gamma_0$. Fix any $r \in P$. By Fact 1, $F(\gamma_1, r) = F(\gamma_0, r)$. However, by Fact 2, $F(\gamma_1, r) = 1$ and $F(\gamma_0, r) = 0$. This provides the needed contradiction.

It remains to prove the claim. Fix $a, 1 \leq a \leq N$, and choose r_1, \dots, r_k so that $\ell(r_1 \dots r_k) = a$. By assumption, $n - k \geq n - m \geq 3m + 1 - m = 2m + 1$

> 2 , so that there exist two distinct participants, r_{k+1} and r_{k+2} , in $P - \{r_1, \dots, r_k\}$. Assume without loss of generality that $r_{k+1} > r_{k+2}$ in the total order on P . We construct a k -round scenario σ with $L_\sigma \subseteq \{r_1, \dots, r_k\}$, in which $\sigma_{r_{k+1}} = \alpha_a$ and $\sigma_{r_{k+2}} = \alpha_{a+1}$. Let

$$\sigma(w) = \begin{cases} 0 & \text{if } w = r_1 \dots r_i p x, \\ & \text{where } 0 \leq i \leq k, p \in P, x \in P^{0:k-i} \\ & \text{and } p < r_{i+1}, \\ 1 & \text{otherwise.} \end{cases}$$

We show that σ is a k -round scenario, with $L_\sigma \subseteq \{r_1, \dots, r_k\}$. Let $q \in P - \{r_1, \dots, r_k\}$, $p \in P$ and $w \in P^{0:k-1}$. We must show that $\sigma(wq) = \sigma(wqp)$. Now, $|wq| \leq k$, so that wq is of the form $r_1 \dots r_i s x$, where $i \leq k-1$, $s \in P$, and $s \neq r_{i+1}$. Then

$$\begin{aligned} \sigma(wq) &= \begin{cases} 0 & \text{if } s < r_{i+1}, \\ 1 & \text{otherwise,} \end{cases} \\ &= \sigma(r_1 \dots r_i s x p) = \sigma(wqp). \end{aligned}$$

Next, we show that $\sigma_{r_{k+1}} = \alpha_a$. Let $w \in P^k$. Then

$$\begin{aligned} \alpha_a(w) &= \begin{cases} 0 & \text{if } \ell(w) < a, \\ 1 & \text{otherwise,} \end{cases} \\ &= \sigma(wr_{k+1}) = \sigma_{r_{k+1}}(w). \end{aligned}$$

Finally, we show that $\sigma_{r_{k+2}} = \alpha_{a+1}$. Let $w \in P^k$. Then

$$\begin{aligned} \alpha_{a+1}(w) &= \begin{cases} 0 & \text{if } \ell(w) \leq a, \\ 1 & \text{otherwise,} \end{cases} \\ &= \sigma(wr_{k+2}) \quad \text{since } r_{k+2} < r_{k+1}, \\ &= \sigma_{r_{k+2}}(w). \quad \square \end{aligned}$$

Note that Theorem 4 provides an upper bound on both the number of processors and the number of rounds. Thus, it demonstrates that both the lower bounds of Theorems 3 and 5 are tight.

7. Open question

The most important question remaining involves the amount of communication and storage needed to assure interactive consistency. The algorithm in [7], which uses the minimum possible number of rounds, involves sending enormous amounts of information - approximately n^{m+2} values in v . We would like to know if this amount can be reduced, say to an amount polynomial in n and m (using either the minimum number, $m+1$, of rounds, or perhaps a larger number of rounds). An algorithm using such a reduced amount of communication might be of considerably more practical value than the current algorithm.

Acknowledgement

The authors thank Leslie Lamport, Michael Merritt and Eugene Stark for helpful discussions and suggestions about this manuscript.

References

- [1] D. Davies and J. Wakerly, Synchronization and matching in redundant systems, *IEEE Trans. Comput.* C-27 6 (1978) 531-539.
- [2] D. Dolev, The Byzantine generals strike again, *J. Algorithms* 3 (1982).
- [3] D. Dolev, Unanimity in an unknown and unreliable environment, 22nd Annual Symp. on Foundations of Computer Science (1981) 159-168.
- [4] L. Lamport, Using time instead of timeout for fault-tolerant distributed systems, *J. Algorithms* 3 (1) (1982).
- [5] L. Lamport, R. Shostak and M. Pease, The Byzantine generals problem, *ACM Trans. Programming Languages, to appear.*
- [6] N.A. Lynch and M.J. Fischer, On describing the behavior and implementation of distributed systems, *Theoret. Comput. Sci.* 13 (1981) 17-43.
- [7] M. Pease, R. Shostak and L. Lamport, Reaching agreement in the presence of faults, *J. ACM* 27 (2) (1980) 228-234.