Intern. J. Computer Math., Vol. 48, pp. 21–29
Reprints available directly from the publisher
Photocopying permitted by license only

AN ALGORITHM FOR IDENTIFICATION OF MALICIOUSLY FAULTY UNITS

FRANC NOVAK*

Institute Jožef Stefan, University of Ljubljana,
Jamova 39, 61111 Ljubljana, Slovenia

and

SANDI KLAVŽAR*

Department of Mathematics, PF Maribor,
Koroška c. 160, 62000 Maribor, Slovenia

(Received 17 March 1992)

The paper proposes an algorithm for identification of malicious units in a fully connected system. The assumed scenario is the same as in an execution of a full information gathering algorithm for Byzantine agreement without authentication. A criterion that can be used by a fault-free unit to identify a malicious unit is described. The exact bound of the maximum number of equal values that a unit which can be identified as malicious, sends to the other units, is calculated. It is proved that if a unit can be identified as malicious by sending some different messages to the other units, then this can be done in the two consecutive phases. These results form the core of the proposed algorithm.

KEY WORDS: Fault-tolerance, distributed systems, Byzantine agreement, diagnosis.

C.R. CATEGORIES: F.2.2, nonnumerical algorithms and problems.

1 INTRODUCTION

Since its introduction in 1980, [9], much attention has been given to the interactive consistency problem, also called the Byzantine Generals Problem, [7]. Requirements for the existence of a solution have been explored, [3]–[4], [6], [9], and numerous algorithms for reaching agreement in the presence of malicious faults have been derived [1], [2], [6]–[10], [12].

As Pease, Shostak, and Lamport in their original paper pointed out, an algorithm for reaching agreement need not reveal which units are faulty; it matters only that the fault-free units compute the same interactive consistency vector. Our objective is to explore the feasibility of identification of malicious units from the data generated by the message exchange process that takes place in the scenario in an execution of a full information gathering algorithm for the Byzantine Generals Problem solution

* Supported in part by the Research Council of Slovenia

without authentication. The approach differs from the work of Gupta and Ramakrishnan [5] which have addressed the problem of system level fault diagnosis in the presence of maliciously faulty units in terms of a generalized PMC model, [11]. In this model, the diagnosis is performed by an external observer from the outcomes of a single application of tests.

The paper is organized in the following way. In Section 2 we describe the assumed scenario for identification of maliciously faulty units. In Section 3 we suggest a criterion according to which a fault-free unit can identify a malicious unit. We introduce the notion of reliable (unreliable) messages and describe the conditions in which a fault-free unit can identify a malicious unit by the proposed criterion. In Section 4 we present an algorithm for identification of malicious units. Finally, in Section 5 some concluding remarks are drawn.

2 ASSUMED SCENARIO

Consider a point-to-point connected system composed of n independent units, of which no more than k are faulty. Each unit is able to communicate with any other only by means of two-party messages. The communication is assumed to be fail-safe. A unit which receives a message always knows the name of the unit that sends the message (i.e., a faulty unit cannot forge a different identity). The message exchange process is organized into phases. A phase is defined to be the interval of time in which each fault-free unit exchanges information with other units. It is assumed that the fault-free units are synchronized such that at every instant they are all executing the same phase. Each fault-free unit is able to decide when to terminate the current phase (i.e., missing messages due to omission faults are detected).

In the first phase the units exchange their private values; in phase i they exchange the information they obtained in phase $i - 1$. A typical message W is of the form: unit u_2 told unit u_1 that unit u_3 told unit u_2 that unit u_4 told unit $u_3 \ldots$ that unit u_i told unit u_{i-1} that X is u_i's private value. Hence a message consists of some value X and of a string indicating the sequence of units participating in its transfer. We shall refer to the length of the string as the *length of the message*. Likewise, the above message W is of length $i - 1$, also denoted by $|W| = i - 1$. Furthermore, we shall denote the value transmitted in the sequence of units $u_1u_2 \ldots u_i$ as $X = \sigma(u_1u_2 \ldots u_i)$. A message is passed only to the units that do not participate in it. Fischer and Lynch [3] showed that $k + 1$ phases are required both with and without authentication to assure interactive consistency.

A unit is said to be *malicious* if it has at least one of the following two properties:

- it gives different private values to at least two other units in the first phase (i.e., the unit lies about its own private value);
- it generates at least one message such that its contents do not agree with the information obtained in the previous phase (i.e., the unit lies about the information obtained from other units).

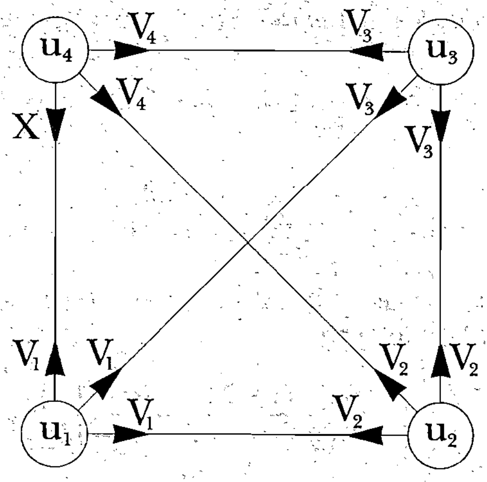

Figure 1 Messages exchanged in phase 1.

The Byzantine Generals Problem is known to be solvable for the fully connected system of n units, if and only if, $n \geq 3k + 1$, [3], [9]. It can easily be shown that the above condition is not sufficient for the problem of malicious unit identification.

Consider the case for $n = 4$ and $k = 1$. According to [9], two phases are sufficient to achieve interactive consistency.

Let the units u_1, u_2, and u_3 be fault-free, let the unit u_4 be malicious, and assume the messages exchanged in phase 1 as shown in Figure 1. In phase 2, u_1 may say that u_4 sent X, and u_4 may say that it sent V_4. The other processors can't tell which of the two is malicious. Notice that this example is independent of n and that the problem is inherently impossible to be solved under certain situations.

In the following we suggest a criterion by which a fault-free unit can identify a malicious unit and explore conditions under which the identification process is feasible.

3 CONDITION FOR IDENTIFICATION OF MALICIOUS UNITS

We say that a fault-free unit can *identify* a malicious unit by a *majority criterion*, if the number of equal values that the fault-free unit gets about the private value of any fault-free unit is greater than the maximum number of equal values that the fault-free unit gets about the private value of any malicious unit.

During the message exchange process, a unit receives reliable and unreliable messages. Messages generated such that only fault-free units participated in their distribution are considered to be *reliable*, otherwise they are *unreliable*. Note that all messages of length 1 are reliable.

LEMMA 1 *Assume that in a system with n units at least m units are fault-free, $m \geq 2$. Then the number of:*

(i) reliable messages of length $\leq i$, $1 \leq i \leq m-1$, from a specific fault-free unit to a specific fault-free unit is at least

$$\sum_{j=1}^{i} \frac{(m-2)!}{(m-j-1)!},$$

(ii) unreliable messages of length $\leq i$, $2 \leq i \leq m-1$, from a specific faulty unit to a specific fault-free unit is at most

$$\sum_{j=2}^{i} \left[\frac{(n-2)!}{(n-j-1)!} - \frac{(m-1)!}{(m-j)!} \right].$$

Proof (i) In a reliable message sequence of length i there are $i-1$ fault-free units between the source unit and the receiver. Since both the receiver and the source node are fault-free, we choose among at least $m-2$ fault-free units. Hence there are at least $\binom{m-2}{i-1}$ such sequences. As every permutation of a reliable message sequence is a reliable message sequence itself, the total number of reliable message sequences of length i is at least

$$\binom{m-2}{i-1}(i-1)! = \frac{(m-2)!}{(m-i-1)!}.$$

(ii) By a similar arguement as in (i) one sees easily that the number of all messages of length i from a unit to any other unit is equal to

$$\binom{n-2}{i-1}(i-1)! = \frac{(n-2)!}{(n-i-1)!},$$

while the number of all reliable messages of length i from a faulty unit to a fault-free unit is at least

$$\binom{m-1}{i-1}(i-1)! = \frac{(m-1)!}{(m-i)!}.$$

As a message is unreliable if and only if it is not reliable, the result follows. ■

LEMMA 2 *Assume that in a system consisting of n units m units are fault-free, $m \geq 2$. Let w be a faulty unit and let $t(w)$ be the maximum number of equal private values that w sends to the other units in the first phase. Then the maximum possible number of equal reliable messages of length $\leq i$, $2 \leq i \leq m-1$, from w to a specific fault-free* *unit u is equal to*

$$t(w)\sum_{j=1}^{i-1}\frac{(m-1)!}{(m-j)!}.$$

Proof WLOG assume $t(w) \leq m$ and let u be a fault-free unit. Clearly, the maximum is reached if in the first phase all $t(w)$ equal private values of w are sent to the fault-free units. If u in the first phase receives from w its most frequent private value then the maximum number of equal reliable messages of length $\leq i$ from w is equal to

$$1+(t(w)-1)\sum_{j=1}^{i-1}\frac{(m-1)!}{(m-j)!},$$

and if u in the first phase does not receive from w its most frequent private value then the maximum number of equal reliable messages of length $\leq i$ from w is equal to

$$t(w)\sum_{j=1}^{i-1}\frac{(m-1)!}{(m-j)!}. \qquad \blacksquare$$

Before proving our main result we need the following technical lemma.

LEMMA 3 *Let $n = m + k$, $k \geq 1$, $m \geq 2$. Then for $2 \leq i \leq m$ the following inequality holds*

$$\sum_{j=2}^{i}\frac{(n-2)!}{(n-j-1)!} \geq k\sum_{j=2}^{i}\frac{(m-2)!}{(m-j)!}+\frac{(m-2)!}{(m-i-1)!}.$$

Proof We observe that for n, m, k and i as defined, the following inequality is obviously true:

$$\frac{(m-2)!}{(m-i)!}+\frac{(n-2)!}{(n-i-1)!} \geq k\frac{(m-2)!}{(m-i)!}+\frac{(m-2)!}{(m-i-1)!}. \tag{1}$$

We prove the lemma by induction on i. For $i = 2$ we get $n \geq m + k$, which is clearly true. By expanding

$$\sum_{j=2}^{i}\frac{(n-2)!}{(n-j-1)!}=\sum_{j=2}^{i-1}\frac{(n-2)!}{(n-j-1)!}+\frac{(n-2)!}{(n-i-1)!},$$

we can use the induction hypothesis and get

$$\sum_{j=2}^{i}\frac{(n-2)!}{(n-j-1)!} \geq \frac{(m-2)!}{(m-i)!}+k\sum_{j=2}^{i-1}\frac{(m-2)!}{(m-j)!}+\frac{(n-2)!}{(n-i-1)!}.$$

From here and (1) we conclude

$$\sum_{j=2}^{i} \frac{(n-2)!}{(n-j-1)!} \geq k \sum_{j=2}^{i-1} \frac{(m-2)!}{(m-j)!} + k\,\frac{(m-2)!}{(m-i)!} + \frac{(m-2)!}{(m-i-1)!}. \qquad \blacksquare$$

Let $t(w)$ be the maximum number of equal private values that w sends to the other units in the first phase.

THEOREM 4 *Let a system consist of n units, $n \geq 3$, and assume that at most k units are malicious, $k \geq 1$.*

Then, for any unit w identified as faulty by a fault-free unit by the majority criterion: $t(w) < n - 2k$. Furthermore, two phases are sufficient for its identification.

Conversely, let $t(w) < n - 2k$. Then the number of equal private messages that a fault-free unit gets about a faulty unit w after phase 2 is smaller than $n - k - 1$, which ensures a fault-free unit to identify w.

Proof Let v be a fault-free unit and assume that v identifies the unit w in phase i. Clearly, $i \geq 2$ since nothing can be done after phase 1. We may also assume that $i \leq m - 1$ as no reliable messages are exchanged between fault-free units in phase m and subsequent phases.

Then, according to the assumed criterion, the minimum number of equal values that the unit v gets about the private value of any fault-free unit is greater than the maximum number of equal values that the unit v could possibly get about the private value of the unit w. Hence due to Lemma 1 (i) and (ii) and Lemma 2 the following holds:

$$\sum_{j=2}^{i} \frac{(m-2)!}{(m-j-1)!} > t(w) \sum_{j=1}^{i-1} \frac{(m-1)!}{(m-j)!} + \sum_{j=2}^{i} \left[\frac{(n-2)!}{(n-i-1)!} - \frac{(m-1)!}{(m-j)!} \right].$$

Applying Lemma 3 we get

$$\sum_{j=2}^{i} \frac{(m-2)!}{(m-j)!} > t(w) \sum_{j=1}^{i-1} \frac{(m-1)!}{(m-j)!} + k \sum_{j=2}^{i} \frac{(m-2)!}{(m-j)!} - (m-1) \sum_{j=2}^{i} \frac{(m-2)!}{(m-j)!},$$

and from here by rearranging

$$(m-k) \sum_{j=2}^{i} \frac{(m-2)!}{(m-j)!} > t(w) \sum_{j=1}^{i-1} \frac{(m-1)!}{(m-j)!}.$$

From the last inequality we get $t(w) < m - k = n - 2k$. Furthermore, the inequality shows that the upper bound for $t(w)$ is obtained after the first two phases, ($i = 2$).

To prove the converse, we claim that a fault-free unit v is able to identify the unit w as malicious, if the number of equal values which the unit v receives in phase 2 about the private value of the unit w is smaller than $n - k - 1$.

Let u be a fault-free unit, $u \neq v$. Since $t(u) = n - 1$, v gets at most k unreliable messages about the private value of u in phase 2. Hence, v receives in phase 2 at least

$n - k - 1$ equal values of the private value of the unit u. Let a malicious unit w send $t(w)$ equal values of its private value to the other units in the first phase. In the second phase, let the remaining $k - 1$ malicious units lie about the received private value from the unit u, by claiming that they received w's most frequent private value in the first phase. In this way, the fault-free unit v may receive at most $t(w) + (k - 1)$ equal values of the private value of the unit w. Since $t(w) < n - 2k$ is assumed, then

$$t(w) + (k - 1) < n - 2k + k - 1 = n - k - 1.$$

Hence, a fault-free unit v can distinguish a malicious unit w from a fault-free unit u on the basis of the number of equal private values and the proof is complete. ■

4 ALGORITHM FOR IDENTIFICATION OF MALICIOUS UNITS

In previous section we described a criterion for identification of a malicious unit in the first two phases of message exchange. However, we can use the same idea to detect units which behave maliciously in later phases. In phase i we apply the same criterion for each message of length $i - 2$ generated by a given unit. Hereby, the units that are already involved in the message and the units that have already been identified as faulty should be excluded from the further exchange of the message. Let us collect all such units at the level of unit u in the set $F(u)$.

Notice that if $|W| = 0$ then W is the empty sequence of units. Let procedure *False* be a procedure for detecting "stupid" mistakes, such as for example a missing message or receiving two messages containing the same sequence of units. Finally, if A is a multiset, let the function *Majority (A)* return the number of occurrences of the most frequent element in the multiset A.

Now, the algorithm executed by unit u can be expressed in the following way:

```
$F(u) \leftarrow \{u\}$; $phase \leftarrow 1$
while ($phase \leq k + 1$) and ($n > 2k + 1$) do
  begin
    {message exchange}
    foreach ($W$; $|W| = phase - 1$) do in parallel
      foreach ($v \neq u$, $v \notin W$) do in parallel
        $\sigma(vuW)$
      end foreach
    end foreach
    {detection of "stupid" errors}
    foreach ($W$; $|W| = phase$) do in parallel
      if $\sigma(uW) = False$ then
        $W \leftarrow vW'$; $Remove(v)$
      end if
    end foreach
    {majority criterion for the last two phases}
    if $phase \geq 2$ then
```

```
        foreach ($W$; $|W| = phase - 1$) do in parallel
            $W \leftarrow vW'$
            if $v \notin F(u)$ then
                $s \leftarrow Majority\ (\{\sigma(uzW);\ z \notin F(u),\ z \notin W\} \cup \{\sigma(uW)\})$
                if $s < n - k - 1$ then $Remove(v)$
            end if
        end foreach
    end if
    $n \leftarrow n - 1$; $phase \leftarrow phase + 1$
end while
```

The procedure *Remove* is simple:

```
procedure $Remove(v)$
    $F(u) \leftarrow F(u) \cup \{v\}$
    $n \leftarrow n - 1$; $k \leftarrow k - 1$

        end
```

5 CONCLUDING REMARKS

We have presented an algorithm which can be used by a fault-free unit to identify malicious units during the execution of a full information gathering algorithm for Byzantine agreement without authentication. Notice that the upper bound of the number of malicious units is assumed to be known to a fault-free unit. From this point of view, the approach is similar to the results related to the t-fault diagnosability of the PMC model, [11]. However, the diagnosis is performed in a distributed way without global observer. System fault-free units can exchange their local diagnoses (i.e., each of them can report the identified malicious units to the other units) in the next execution of the algorithm for Byzantine agreement. Hence the two algorithms can run in parallel such that current execution of the algorithm for Byzantine agreement is also used to distribute the results of prior executed algorithm for identification of maliciously faulty units.

References

[1] D. Dolev, The Byzantine Generals Strike Again, *J. Algorithms* **3** (1982), 14–30.

[2] D. Dolev and H. R. Strong, Authenticated Algorithms for Byzantine Agreement, *SIAM J. Comput.* **12** (1983), 656–666.

[3] M. Fischer and N. Lynch, A Lower Bound for the Time to Assure Interactive Consistency, *Inform. Processing Lett.* **14** (1982), 183–186.

[4] M. Fischer, N. Lynch, and M. Merritt, Easy impossibility proofs for distributed consensus problems, *Distributed Computing* **1** (1986), 26–39.

[5] R. Gupta, I. V. Ramakrishnan, System-Level Fault Diagnosis in Malicious Environments, *Proc. 17th Int. Symp. on Fault Tolerant Computing* IEEE Computer Society Publications, (1987), 184–189.

[6] V. Hadzilacos, Connectivity requirements for Byzantine agreement under restricted types of failures, *Distributed Computing* **1** (1987), 95–103.

[7] L. Lamport, R. Shostak, and M. Pease, The Byzantine Generals Problem, *ACM Trans. Program. Lang. Syst.* **4** (1982), 382–401.

[8] N. Lynch, M. Fischer, and R. Fowler, A Simple and Efficient Byzantine Generals Algorithm, *Proc. 2nd Symp. on Reliable Dist. Soft. and Data Base Syst.* (1982), 46–52.

[9] M. Pease, R. Shostak, and L. Lamport, Reaching Agreement in the Presence of Faults, *J. ACM* **27** (1980), 228–234.

[10] K. J. Perry and S. Toueg, Distributed Agreement in the Presence of Processor and Communication Faults, *IEEE Trans. Software Eng.* **SE-12** (1986), 477–482.

[11] F. P. Preparata, G. Metze, and R. T. Chien, On the connection assignment problem of diagnosable systems, *IEEE Trans. Electronic Comput.* **EC-16** (1967), 848–854.

[12] T. K. Srikanth and S. Toueg, Simulating authenticated broadcasts to derive simple fault-tolerant algorithms, *Distributed Computing* **1** (1987), 80–94.