

CHRISTIAN-ALBRECHTS-UNIVERSITÄT ZU KIEL
Institut für Informatik und Praktische Mathematik

Prof. Dr. W.-P. de Roever
Martin Steffen, Immo Grabe



Verteilte Algorithmen

Wintersemester 2003/04

Serie 6

1. Dezember 2003

Thema : — Mittsemestertest! — (Aufgaben mit Lösungshinweisen)

Ausgabetermin: 1. Dezember 2003**Abgabe: 9. Dezember 2003**

Mittsemestertest: Diese Aufgabenserie ist von jedem *allein* zu bearbeiten! Die Aufgaben sind diesmal etwas breiter gestreut, d.h., auch 2 Aufgaben aus vorangegangenen Kapiteln sind mit aufgenommen. Bitte beachten Sie auch den Abgabetermin. Viel Erfolg!

Der Zettel behandelt *verteilte Einigung*, speziell die *randomisierte Variante*, und zwar in synchronen Netzen mit unzuverlässiger Kommunikation (Nachrichtenverlust).

Aufgabe 1 (Kantengewichte und Kantenidentifikatoren (4 Punkte)) Bearbeiten Sie Aufgabe 4.18 aus [?].

Aufgabe 2 (Byzantinismus (6 Punkte)) Bearbeiten Sie Aufgabe 6.21 über *EIGByz* aus [?].

Solution: The exercise is intended to get a feeling about how the *EIG*-algorithm for Byzantine failures works by playing around with scenarios, where the number of failed processors is large enough to undermine consensus.

1. *7 nodes, 2 failed processors, 2 rounds:* Since the algorithm's decision is based on *EIG-trees*, we first recapitulate this structure to fill it with values which lead to *disagreement*.

As usual for consensus algorithms, the basic idea is to send around values often enough to come to some conclusion based on majority. Since there might be Byzantine process failures, it is not enough to take the value received by a process “for the word of it” but one needs to hear a “second opinion” (and a third ...) about the value of a process. In other words, each process not just remembers the value heard *directly* from a process, but also what another process heard the value would be, and this level of indirection must be deep enough to filter out the cheater. What makes it hard is, that a liar may lie not only about his own value, but he can also lie about which value he knows about others, which creates confusions. This need of higher-order indirect information about

values, which forms the core of the exponential information gathering idea, makes the difference between the consensus problem in the presence of *stopping* failures and the situation here. When only stopping failures occur, needs no higher levels of indirection.¹

May it as it is, the information of one player “process i said value v ”, “process j said, that process k said that value v' ” ... is stored in a hierarchical data structure, the EIG-tree, whose depth is given by the number of rounds. In this exercise, we play $f = 2$ rounds which gives $f + 1 = 3$ level.² Each level in the tree, counted from the root, corresponds to the level of indirection of the information, so nodes, say, 2 and 23, correspond to “first-hand information” and “second-hand information” of some process about node 2. In the first round, located at level one in the tree, 2 directly gave its value, and in the second round resp. at the second tree level, 3 relayed the information that 2 possess some given value.

In our example, each tree —there are 7— consists of 7 nodes in the first level each of which possess 6 descendants in the next level, which form the leaves. To systematically approach construct a *disagreement* situation, let us consider the consequences of the 5 honest vs. 2 lying processes situation in terms of the given algorithm. If a sound process broadcast some value, say 0 in the first round, then only two of the values may arrive *forged* as second-level information at the processes. This means, at the second level of each tree, the leaves contain at least 4 original values about that process, which makes a majority of the 6 leaves.

According to the decision procedure of the *EIGByz*-algorithm and the decision tree, it means that the father-node of the 6 leaves is labelled with an *accurate* opinion about the picked process. Taking the 5 nodes of level 1, corresponding to sound processes, this means that in the decision tree³ at level one there a *5 correct labels*. And this for all sound processes.

To lead this to a contradiction, we therefore need *at least two different start values*. say *true* and *false*, in 3 : 2-ratio.⁴ To be concrete, let the processes 1, ..., 5 be sound and 6 and 7 the cheaters and let furthermore *false* be the initial value of 1, 2, and 3, and 5 and 6 are initially *true*.

As mentioned, this fixes 5 values at level 1 of the trees of all non-cheaters to the array:

$$\text{for all } T_i \text{ with } i = 1, \dots, 5: \quad \textit{false}, \textit{false}, \textit{false}, \textit{true}, \textit{true}, ?, ? \quad (1)$$

Now we need the cheaters to fill in the question marks of two different non-cheaters in such a way that they decide at their respective root node in a different way.

Consider, for instance, the cheater 6, respectively the tree node 6 at level 1 in the trees of the honest processes. In order than some of the sound processes considers 6 to have value *false* instead of the *?*, there must be a majority of leaves that has heard about

¹Since messages can be lost, one cannot directly base the decision on what one heard from another process, because one could have heard nothing, while other players could have received a dissenting vote from the failing (=stopping) processes and thus would fall back to the default vote. That was the idea of the various algorithms for stopping failures.

²We start with “zero-knowledge” before the first round.

³of course again for a sound process, only

⁴One could have guessed that from the start, too.

true. Indeed, *three nodes* suffice, since the cheater 7 may lie about what he had heard about 6, filling in a 4th *false*.

As, independant of the node 67, there are already 3 *false*'s in the 6 leaves, there can't be a majority of *true* for all sound trees, independant of any further cheating of 7, the dissenting vote can only be *the default value*. To achieve this, at least 3 sons of tree node 6 are marked with *true*. This means, 6 must send

To sum up concretly: The cheaters and non-cheater processes are numbered as above and possess the start values of Equation (??). The disagreement is constructed by having conflict between 1 and 2. Cheater 6 sends *false* to 1,2,3 and *true* to the rest in the first round; in the second round 7 *sends to 1 that it had heard that 6 is false*. This gives for the tree in process 1 that

61	<i>false</i>
62	<i>false</i>
63	<i>false</i>
64	<i>true</i>
65	<i>true</i>
(66)	—
67	<u><i>false</i></u>

Propagating this majority to the father 6 yields a *forth false* at level 1, and process 1 decides to *false*.

For process 2, the situation is harder since we need not just one *false* but we must avoid to assing another false to one of the ?, i.e., for node 6 and 7. The situation for the leaves 6*j* is already quite fixed, and the only place we still have freedom is the underline value for 67, which blocks the majority and leads to the default decision at node 6 in the tree for process 2. Remains the subtree for 7, where we also need to prevent a *false*-labelling. That's simple: 7 just sends *true* in the first round, and the other remaining cheater correctly states (for instance) that he had heard *true* from 7.⁵

61	<i>false</i>	71	<i>true</i>
62	<i>false</i>	72	<i>true</i>
63	<i>false</i>	73	<i>true</i>
64	<i>true</i>	74	<i>true</i>
65	<i>true</i>	75	<i>true</i>
(66)	—	76	<i>true</i>
67	<u><i>true</i></u>	(77)	—

Aufgabe 3 (Ablauf eines I/O Automaten (4 Punkte)) Bearbeiten Sie Aufgabe 8.4 aus [?].

Aufgabe 4 (Traces und fair traces (6 Punkte)) Bearbeiten Sie Aufgabe 8.12 aus [?].

⁵As we have made no assumptions so far on the behavior of 7 except that it sends a value of “6 is *false*” to 1 and “6 is *true*” to 2 in the second round, we have much freedom here.