Brief Announcement: Optimal Implementation of the Weakest Failure Detector for Solving Consensus*

Mikel Larrea
Univ. Pública de Navarra
31006 Pamplona, Spain
mikel.larrea@unavarra.es

Antonio Fernández
Universidad Rey Juan Carlos
28933 Móstoles, Spain
afernandez@acm.org

Sergio Arévalo
Universidad Rey Juan Carlos
28933 Móstoles, Spain
s.arevalo@escet.urjc.es

Unreliable failure detectors were introduced by Chandra and Toueg [2] as a mechanism that provides (possibly incorrect) information about process failures. They showed how unreliable failure detectors can be used to solve the Consensus problem in asynchronous systems. They also showed in [1] that one of the classes of failure detectors they defined, namely *Eventually Strong* ($\diamond\mathcal{S}$), is the weakest class allowing to solve Consensus[1].

This brief announcement presents a new algorithm implementing $\diamond\mathcal{S}$. Due to space limitation, the reader is referred to [4] for an in-depth presentation of the algorithm (system model, correctness proof, and performance analysis). Here, we present the general idea of the algorithm and compare it with other algorithms implementing unreliable failure detectors.

The algorithm works as follows. We have n processes, $p_1, \ldots, p_n$. Initially, process p_1 starts sending messages periodically to the rest of processes. The rest of processes initially *trust* p_1, and wait for its messages. If a process does not receive a message within some timeout period from its trusted process, then it suspects its trusted process and takes the next process as its new trusted process. If a process trusts itself, then it starts sending messages periodically to its successors. Otherwise, it just waits for periodical messages from its trusted process. If, at some point, a process receives a message from a process p_i such that p_i precedes its trusted process, then it will trust p_i again, increasing the value of its timeout period with respect to p_i.

With this algorithm, eventually all the correct processes will permanently trust the same correct process. This provides the eventual weak accuracy property required by $\diamond\mathcal{S}$. By simply suspecting the rest of processes, we obtain the strong completeness property required by $\diamond\mathcal{S}$.

Alg.	# of msgs	bit size of msgs	amount of information	eventual mon. degree
[2]	$n\mathcal{C}$	$\Theta(\log n)$	$\Theta(n\mathcal{C}\log n)$	$\mathcal{C}^2$
[3]	$2\mathcal{C}$	$\Theta(n)$	$\Theta(n\mathcal{C})$	$2\mathcal{C}$
LFA	$n-1$	$\Theta(\log n)$	$\Theta(n\log n)$	$\mathcal{C}-1$

Table 1: Comparing failure detector algorithms.

Our algorithm compares favorably with the algorithms proposed in [2] and [3] in terms of the number and size of the messages periodically sent and the total amount of information periodically exchanged. Since algorithms implementing failure detectors need not necessarily be periodic, we propose a new and (we believe) more adequate performance measure, which we call *eventual monitoring degree*. Informally, this measure counts the number of pairs of correct processes that will infinitely often communicate. We show that the proposed algorithm is optimal with respect to this measure. Table 1 summarizes the comparison, where $\mathcal{C}$ denotes the number of correct processes and LFA denotes the proposed algorithm.

1. REFERENCES

[1] T. D. Chandra, V. Hadzilacos, and S. Toueg. The weakest failure detector for solving consensus. *Journal of the ACM*, 43(4):685–722, July 1996.

[2] T. D. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, March 1996.

[3] M. Larrea, S. Arévalo, and A. Fernández. Efficient algorithms to implement unreliable failure detectors in partially synchronous systems. In *Proceedings of the 13th International Symposium on Distributed Computing (DISC'99), (Formerly WDAG)*, pages 34–48, September 1999.

[4] M. Larrea, A. Fernández, and S. Arévalo. Optimal implementation of the weakest failure detector for solving consensus. Technical Report, Universidad Pública de Navarra, Departamento de Matemática e Informática, January 2000. http://www.gsd.unavarra.es/pres/miembros/mikel/.

*Partially supported by the Spanish Research Council (contracts TIC99-0280-C02-02, TEL99-0582, and TIC98-1032-C03-01), and the Madrid Regional Research Council (contract CAM-07T/00112/1998).

[1] In fact, the *Eventually Weak* failure detector class, $\diamond\mathcal{W}$, is presented as the weakest one for solving Consensus. However, Chandra and Toueg have shown in [2] that $\diamond\mathcal{S}$ and $\diamond\mathcal{W}$ are equivalent.

PODC 2000 Portland Oregon