Brief Announcement: Minimal System Conditions to Implement Unreliable Failure Detectors

Antonio Fernández
LADyR, GSyC
Universidad Rey Juan Carlos
28933 Móstoles, Spain
anto@gsyc.escet.urjc.es

Ernesto Jiménez
EUI
Universidad Politécnica de Madrid
28031 Madrid, Spain
ernes@eui.upm.es

Sergio Arévalo
LADyR, GSyC
Universidad Rey Juan Carlos
28933 Móstoles, Spain
sarevalo@gsyc.escet.urjc.es

Categories and Subject Descriptors: C.2.4 [Computer-Communication Networks]: Distributed Systems; C.4 [Performance of Systems]: Fault Tolerance

General Terms: Algorithms, reliability, theory.

Keywords: Unreliable failure detectors, partial synchrony, fault-tolerance.

In a tutorial at PODC 2002, Keidar and Rajsbaum [3] asked, among other questions, what are the weakest requirements on systems that allow to implement the different classes of failure detectors. In this brief announcement we explore the minimal system requirements to implement unreliable failure detectors [1]. A complete version of this brief announcement can be found in [2].

1. The System Model

We consider a system S formed by a finite set Π of $n > 1$ processes. We assume that processes can communicate among each other only by sending and receiving messages, and that every pair of processes is connected by a pair of directed links (with opposite directions). The execution of processes advances in steps, with an upper bound on the units of time that any non-faulty process takes to advance (execute) one step. A process can fail by permanently crashing. We say that a process is correct if it does not fail. We also assume that processes have clocks that can accurately measure intervals of time (it is not necessary that they are synchronized). We assume that the algorithms have no a priori knowledge of the number of failures that can occur.

We consider the following three types of links: *Lossy asynchronous*, *timely*, and *eventually timely*. Note that, timely links are special cases of eventually timely links.

2. Classes of Failure Detectors

We study four traditional classes: $\mathcal{P}$, $\mathcal{S}$, $\diamond\mathcal{P}$ and $\diamond\mathcal{S}$ [1], and two additional (perpetual) classes, $\mathcal{P}_4$ [4], and $\mathcal{S}'$, which are weak versions (they have weaker accuracy) of $\mathcal{P}$ and $\mathcal{S}$, respectively. As far as we know, the $\mathcal{S}'$ has never been previously proposed.

3. Classes of Systems

We consider two large classes of systems. Class Γ is formed by systems in which each of its links is either lossy asynchronous or timely. Class $\mathcal{E}$ is formed by systems in which each of its links is either lossy asynchronous or eventually timely. Note that $\Gamma \subset \mathcal{E}$. Then, the results of impossibily for $\mathcal{E}$ also apply to Γ, and the results of possibility for Γ also apply to $\mathcal{E}$. On these systems we define two properties:

Weak property: There is a correct process such that all correct processes can be reached from it with links that are not lossy asynchronous.

Strong property: All correct processes can be reached from all correct processes with links that are not lossy asynchronous.

4. Necessary Conditions to Implement Failure Detectors

Theorem 1: If $S \in \mathcal{E}$ is a system that does not satisfy the weak property and one single failure can happen, then no detector in $\diamond\mathcal{S}$, $\mathcal{S}'$, or $\mathcal{S}$ can be implemented.

Theorem 2: If $S \in \mathcal{E}$ is a system that does not satisfy the strong property and one single failure can happen, then no detector in $\diamond\mathcal{P}$, $\mathcal{P}_4$, or $\mathcal{P}$ can be implemented.

5. Algorithms to Implement Failure Detectors

For all systems in Γ (and, hence, for $\mathcal{E}$) we propose an algorithm that implements a failure detector of class $\diamond\mathcal{P}$ if the strong property is satisfied, and of class $\diamond\mathcal{S}$ if only the weak property is satisfied. For all systems in $\mathcal{E}$ we propose another algorithm that implements a failure detector of class $\mathcal{P}_4$ if the strong property is satisfied, and that implements a failure detector of the new detector class $\mathcal{S}'$ if only the weak property is satisfied. From the above theorems, we can say that these detectors run under minimal system conditions.

1. REFERENCES

[1] T. D. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *JACM*, 43(2):225-267, March 1996.

[2] A. Fernández, E. Jiménez and S. Arévalo. Minimal System Conditions to Implement Unreliable Failure Detectors. *RoSaC Technical Report,* 2005. http://gsyc.info/publicaciones/tr.

[3] I. Keidar and S. Rajsbaum. On the cost of fault-tolerant consensus when there are no faults. *Tutorial at PODC'2002,* Monterey, California, July 2002.

[4] M. Larrea. Brief announcement: Understanding perfect failure detectors. In *PODC'2002,* Monterey, California, July 2002.

PODC'05, July 17–20,2005,Las Vegas,Nevada,USA.
ACM 1-58113-994-2/05/0007.