Brief Announcement: Early Decision Despite General Process Omission Failures

Fabrice Le Fessant
Microsoft Research lab
Cambridge, CB3 0FB
UK

Philippe Raipin Parvédy
IRISA, Campus Beaulieu
35042 Rennes Cedex
France

Michel Raynal
IRISA, Campus Beaulieu
35042 Rennes Cedex
France

In the consensus problem, each process proposes a value and the non-faulty processes have to decide (termination) on the same value (agreement) that has to be one of the proposed values (validity). Two versions of the consensus problem are usually distinguished. They differ in the statement of the agreement property. In the non-uniform version (*consensus*), agreement is only required on the non-faulty processes (this means that faulty processes are allowed to decide values different from the value decided by the non-faulty processes). In the uniform version (*uniform consensus*), agreement concerns all the processes that decide, no two processes (be them faulty or not) can decide differently.

Synchronous systems provide upper bounds on processing time and communication delays. Those bounds allow consensus problems to be solved. Basically, a distributed synchronous consensus protocol proceeds by successive rounds. During each round the processes exchange values until they attain a round during which they can conclude that they have converged to the same value. Consensus protocols in synchronous systems are characterized by the failure model they address, the maximal number of processes they allow to be faulty, and the maximal number of rounds they need to attain a decision. Three failure models have been investigated: the *crash* model where processes can fail by prematurely halting, the *omission* model where processes can fail by halting or omitting to send or receive messages they should, and the *Byzantine* model where processes can fail by exhibiting arbitrary behavior. These failure models are of increasing severity. Let n and t denote the total number of processes and the maximum number of processes that can be faulty, respectively. We have:

- Crash failure model: both consensus and uniform consensus can be solved for any value of t (i.e., for $t < n$),
- Omission failure model: consensus can be solved for any t, while uniform consensus requires $t < n/2$,
- Byzantine failure model: consensus can be solved provided that $t < n/3$. (Uniform consensus is meaningless in that model.)

It has been shown that $t + 1$ is a lower bound on the number of rounds to solve any of these agreement problems in presence of up to t faulty processes. So, these agreement problems can be seen as "equivalent" when we consider the worst case time complexity lower bound. The interested reader can consult [3] for a survey on synchronous consensus.

The $t+1$ lower bound on the number of rounds is actually due to worst case scenarios. But, although failures do occur, they are rare in practice. So, from both theoretical and practical points of view, an attractive approach is to look for *early-stopping* consensus protocols, i.e., protocols whose number of rounds depends on f, the number of processes that are actually faulty during a given execution ($0 \leq f \leq t$). The following lower bounds have been established for the crash failure model: consensus can be solved in $f+1$ rounds; uniform consensus can be solved in $\min(f+2, t+1)$ rounds.

While (1) consensus has received a lot of attention in both the crash failure model and the Byzantine failure model, and (2) uniform consensus has (though more recently) received attention in the crash failure model, uniform consensus has received very few attention in the omission failure model. We have recently addressed uniform consensus in the general omission failure model where a process can fail by not sending a message, not receiving a message or crashing. Early deciding uniform consensus protocols resilient to up to t processes exhibiting general omission failures are presented in [1, 2]. These protocols, which assume $t < n/2$ (a necessary requirement), are optimal both in the number of processes (i.e., $2t < n$) and in round complexity (as each requires the processes to execute at most $\min(f+2, t+1)$ rounds). The interested reader can consult the previous references that describe the protocols, prove their correction and optimality, and show that the constraint $2t < n$ is no longer necessary when process failures are limited to be crashes or send omissions (this shows that the "majority of correct processes" assumption is required only to cope with receive omission failures). Interestingly, those results can be extended to the non-blocking atomic commit problem.

PODC'03, July 13–16, 2003, Boston, Massachusetts, USA.

1. REFERENCES

[1] Le Fessant F., The Complexity of Early-Deciding in Unreliable Networks. *Research Report #TR-2003-23*, Microsoft Research, Cambridge, UK, April 2003.

[2] Raipin Parvédy Ph. and Raynal M., Optimal Early Stopping Uniform Consensus in Synchronous Systems with Process Omission Failures. *Research Report #1509*, 20 pages, IRISA, France, January 2003. www.irisa.fr/bibli/publi/pi/2003/1509/1509.html.

[3] Raynal M., Consensus in Synchronous Systems: a Concise Guided Tour. *9th IEEE Pacific Rim Int. Symp. on Dependable Comp.*, Tsukuba (Japan), pp. 221-228, 2002.