Tight Bounds for k-Set Agreement with Limited-Scope Failure Detectors (Brief Announcement)

www.cs.brown.edu/people/mph/detector.pdf

Maurice Herlihy and Lucia Draque Penso
Computer Science Department
Brown University
Providence, RI 02912

In the k-set agreement problem [1], each process in a group starts with a private input value, communicates with the others, and then halts after choosing a private output value. Each process is required to choose some process's input, and at most k distinct values may be chosen.

We consider this problem in an asynchronous message-passing system of $n+1$ processes, of which at most f may fail by halting. A *failure detector* [5] is an unreliable oracle that informs each process when it suspects other processes to have failed. In practice, however, a process can typically detect some failures more easily than others. For example, timeouts may reliably detect failures on the same local area network, but less reliably over a wide-area network. A *limited-scope* failure detector model, the set of processes encompasses sets $X_0, \ldots, X_{q-1}$, such that some correct process in X_i is never suspected by any process in X_i. Failure detectors in class S_x satisfy this property all the time, while failure detectors in $\diamond S_x$ satisfy it eventually.

In a classic paper, Mostéfaoui and Raynal [4] give two algorithms for the k-set agreement task in asynchronous message-passing models augmented with failure detectors from the S_x or $\diamond S_x$ classes, subject to the restrictions that $q = 1$. They conjectured that their algorithms for each model are optimal.

This paper makes the following contributions. We give the first ever lower bounds for k-set agreement protocols employing failure detectors from classes S_x and $\diamond S_x$. For S_x, we show that any k-set agreement protocol must satisfy

$$f < k + x - q,$$

while for $\diamond S_x$,

$$f < min(\left\lfloor \frac{n+1}{2} \right\rfloor, k + x - q).$$

Our proof employs mechanisms from Combinatorial Topology [2, 3]. These methods have been successful in other models, this is the first time such methods have been applied to failure detectors.

For the class S_x, our lower bound implies that the elegant *TWA-based* protocol of Mostéfaoui and Raynal [4] is optimal for the case where $q = 1$. For the class $\diamond S_x$, we give a novel protocol that matches our lower bound, and exceeds the fault-tolerance of the Mostéfaoui and Raynal algorithm for this model. Our protocol has an unexpectedly simple structure: it alternates the TWA-based protocol with a novel *convergence detection* protocol that halts when it detects that an earlier iteration of TWA has succeeded. Our protocol thus disproves the conjecture of Mostéfaoui and Raynal that their algorithm for the $\diamond S_x$ model is optimal.

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

1. REFERENCES

[1] Soma Chaudhuri. Agreement is harder than consensus: set consensus problems in totally asynchronous systems. In *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, pages 311–234, August 1990.

[2] Maurice Herlihy, Sergio Rajsbaum, and Mark R. Tuttle. Unifying synchronous and asynchronous message-passing models. In *Proceedings of the seventeenth annual ACM symposium on Principles of distributed computing*, pages 133–142. ACM Press, 1998.

[3] Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *Journal of the ACM (JACM)*, 46(6):858–923, 1999.

[4] Achour Mostéfaoui and Michel Raynal. k-set agreement with limited accuracy failure detectors. In *Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing*, pages 143–152. ACM Press, 2000.

[5] Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for solving consensus. *Journal of the ACM (JACM)*, 43(4):685–722, 1996.