

비동기적 분산 시스템하에서 선출 문제와 합의 문제의 관련성 연구

김윤, 유선준, 김차영, 박성훈
남서울대학교 컴퓨터학과
e-mail : spark@nsu.ac.kr

Study on Relationship between Election and Consensus in Asynchronous Distributed Systems

Yoon Kim, Sun-Jun Yu, Cha-Young Kim, Sung-Hoon Park
Dept. of Computer Science, NamSeoul University

요 약

본 논문에서는 신뢰 할 수 없는 고장추적 장치로 구성 된 비동기적 분산 시스템 하에서 선출 (election) 문제와 합의(consensus) 문제의 관련성에 관하여 연구하고자 한다. 먼저 선출 문제는 합의 문제보다 더욱 어려운 문제임을 보인다. 합의 문제와는 대조적으로 선출 문제는 시스템 상에서 단 한 개의 노드가 죽은 경우에도 신뢰 할 수 없는 고장 추적 장치를 이용하여 선출 문제를 해결 할 수 없다. 보다 엄밀하게 표현하자면, 선출 문제를 해결하는데 필요한 가장 약한 고장 추적 장치는 완전한 고장 추적 장치이어야 하는 것으로, 이는 합의 문제를 해결하는데 필요한 가장 약한 고장 추적 장치보다 확실히 강한 것이다. 선출 문제가 합의 문제보다 어렵다는 것을 보이기 위해 본 논문에서는 축소(reduction) 프로토콜을 이용 한다.

1. Introduction

To elect a *Leader* (or Coordinator) in a distributed system, an *agreement* problem must be solved among a set of participating processes. This problem, called the *Election* problem, requires the participants to agree on only one leader in the systems. The Election problem is described as follows. At any time, there is at most one process that considers itself a *leader* and all other processes consider it as to be their only leader. If there is no leader, a leader is eventually elected.

Consensus and Election are similar problems in that they are both agreement problems. The so-called FLP impossibility result, which states that it is impossible to solve any non-trivial agreement in an asynchronous system even with a single crash failure, applies to both problems [1]. The starting point of this paper is the fundamental result of Chandra and Toueg[2], which states that Consensus is solvable in asynchronous systems with unreliable failure detectors.

An interesting question is then whether the Election problem can also be solved in asynchronous systems with unreliable failure detectors. The answer to this question is “No”, and this is not surprising because the Election problem has been considered harder than Consensus [3]. However, in contrast to initial intuition, the reason Election is harder than Consensus is not its *Liveness* condition. The difficulty in solving Election is actually its *Safety* condition (all the nodes connected the system never disagree on the leader when the nodes are in a state of normal operation). This condition requires precise knowledge about failures which unreliable failure detectors cannot provide.

The rest of the paper is organized as follows. In Section 2 we describe our system model. In Section 3 we define *Leader Election* and show that it is harder than *Consensus*. Finally, Section 4 summarizes the main contributions of this paper and discusses related and future work.

2. Model and Definitions

Our model of asynchronous computation with failure detection is the one described in [4]. In the following, we only recall some informal definitions and results that are needed in this paper.

2.1 Processes

We consider a distributed system composed of a finite set of processes $\Omega = \{p_1, p_2, \dots, p_n\}$ completely connected through a set of channels. Communication is by message passing, *asynchronous* and *reliable*. Processes fail by crashing; Byzantine failures are not considered.

Asynchrony means that there is no bound on communication delays or process relative speeds. A reliable channel ensures that a message, sent by a process p_i to a process p_j , is eventually received by p_j if p_i and p_j are correct (i.e. do not crash).

To simplify the presentation of the model, it is convenient to assume the existence of a discrete global clock. This is merely a fictional device inaccessible to processes. The range of clock ticks is the set of natural numbers. A history of a process $p_i \in \Omega$ is a sequence of events $h_i = e_i^0 \cdot e_i^1 \cdot e_i^2 \cdot \dots \cdot e_i^k$, where e_i^k denotes an event of process p_i occurred at time k . Histories of correct processes are infinite. If not infinite, the process history of p_i terminates with the event crash_i^k (process p_i crashes at time k). Processes can fail at any time, and we use f to denote the number of processes that may crash. we consider systems where at least one process correct (i.e. $f < |\Omega|$).

A failure detector is a distributed oracle which gives hints on failed processes. We consider algorithms that use failure detectors. An algorithm defines a set of runs, and a run of algorithm A using a failure detector D is a tuple $R = \langle F, H, I, S, T \rangle$: I is an initial configuration of A ; S is an infinite sequence of events of A (made of process histories); T is a list of increasing time values indicating when each event in S occurred; F is failure pattern that denotes the set $F(t)$ of processes that have crashed at any time t ; H is a failure detector history, which gives each process p and at any time t , a (possibly false) view $H(p, t)$ of the failure pattern: $H(p, t)$ denotes a set of processes, and $q \in H(p, t)$ means that process p suspects process q at time t .

2.2 Failure detector classes

Failure detectors are abstractly characterized by *completeness* and *accuracy* properties [4]. Completeness characterizes the degree to which crashed processes are permanently suspected by correct processes. Accuracy restricts the false suspicions that a process can make.

Two completeness properties have been identified. *Strong Completeness*, i.e. there is a time after which every process that crashes is permanently suspected by every correct process, and *Weak Completeness*, i.e. there is a time after which every process that crashes is permanently

suspected by some correct process.

Four accuracy properties have been identified. *Strong Accuracy*, i.e. no process is never suspected before it crashes. *Weak Accuracy*, i.e. some correct process is never suspected. *Eventual Strong Accuracy* i.e. there is a time after which correct processes are not suspected by any correct process; and *Eventual Weak Accuracy*, i.e. there is a time after which some correct process is never suspected by any correct process. A failure detector class is a set of failure detectors characterized by the same completeness and the same accuracy properties (Figure 1).

For example, the failure detector class P , called *Perfect Failure Detector*, is the set of failure detectors characterized by Strong Completeness and Strong Accuracy. Failure detectors characterized by Strong Accuracy are reliable: no false suspicions are made. Otherwise, they are unreliable

Completeness	Accuracy			
	Strong	Weak	◇Strong	◇Weak
Strong	P	S	◇ P	◇ S
Weak	Q	W	◇ Q	◇ W

Fig.1. Failure detector classes

For example, failure detectors of S , called Strong Failure Detector, are *unreliable*, whereas the failure detectors of P are *reliable*.

2.3 Reducibility and transformation

An algorithm A solves a problem B if every run of A satisfies the specification of B . A problem B is said to be *solvable* with a class C if there is an algorithm which solves B using any failure detector of C . A problem B_1 is said to be reducible to a problem B_2 with class C , if any algorithm that solves B_2 with C can be transformed to solve B_1 with C . If B_1 is not reducible to B_2 , we say that B_1 is *harder than* B_2 .

A failure detector class C_1 is said to be *stronger than* a class C_2 , (written $C_1 \geq C_2$), if there is an algorithm which, using any failure detector of C_1 , can emulate a failure detector of C_2 . Hence if C_1 is stronger than C_2 and a problem B is solvable with C_2 , then B is solvable with C_1 . The following relations are obvious: $P \geq Q$, $P \geq S$, $\diamond P \geq \diamond Q$, $\diamond P \geq \diamond S$, $S \geq W$, $\diamond S \geq \diamond W$, $Q \geq W$, and $\diamond Q \geq \diamond W$. As it has been shown that any failure detector with *Weak Completeness* can be transformed into a failure detector with *Strong Completeness* [4], we also have the following relations: $Q \geq P$, $\diamond Q \geq \diamond P$, $W \geq S$ and $\diamond W \geq \diamond S$. Classes S and $\diamond P$ are incomparable.

2.4 Consensus

In the *Consensus* problem (or simply Consensus), every participant *proposes* an input value, and correct participant

must eventually *decide* on some common output value [5]. Consensus is specified by the following conditions. *Agreement*: no two correct participant decide different values; *Uniform-Validity*: if a participant decides v , then v must have been proposed by some participant; *Termination*: every correct participant eventually decide. Chandra and Toueg have stated that *Consensus* is solvable with $\diamond P$ or S [4].

3. Election is harder than consensus

In this section, we show that the Election problem is not solvable in asynchronous systems with unreliable failure detectors. This impossibility result holds even with the assumption that at most one process may crash. Hence Election is harder than Consensus.

3.1 The Election Problem

The proof of the impossibility of Consensus in [1] assumes that it is impossible for a process to determine whether another process has crashed, or is just very slow. This assumption is widely cited as the “reason” for the impossibility result. There are other problems that cannot be solved in asynchronous systems with crash failures for the same intuitive reason that Consensus cannot be solved. Some of these problems can be solved with a weak failure detector; however, some cannot. In particular, the Election problem cannot be solved if a crashed process cannot be distinguished from a slow process.

The Election Problem is specified by the following two properties. *Safety*: All processes connected the system never disagree on a *leader* when the nodes are in a state of normal operation. *Liveness*: All processes should eventually progress to be in a state in which all processes connected to the system agree to the *only one* leader. An *election protocol* is a protocol that generates runs that satisfies the Election specification.

3.2 Impossibility of solving Election Problem

Though $\diamond P$ or S are sufficient to solve Consensus, it is not sufficient to solve Election. Therefore the Election problem is strictly harder than the Consensus problem since even when assuming a single crash, unreliable failure detectors are not strong enough to solve election. In this section, we show that Strong Accuracy is necessary for solving Election, and it is sufficient for solving Election.

Theorem 1 *If $f > 0$, Election can not be solved with either $\diamond P$ or S .*

PROOF. Consider a failure detector D of $\diamond P$. We assume for a contradiction that there exists a deterministic election protocol E that can be combined with a failure detector D such that $E + D$ is also an election protocol.

Consider an algorithm A combined with $E + D$ which solves Election and a run $R = \langle F, H_D, I, S, T \rangle$ of A . We assume that only two processes P_i and P_j are correct and all messages from them is delayed until after t in R .

Consider that P_i is a leader at time (R, k) . At time (R, k_1) where $(k + t) > k_1 > k$, the process P_j falsely suspects other process P_i in some run. At time (R, k_2) where $k_2 > k_1$, P_j considers itself a leader by delaying the receipt of all messages sent by P_i until k_3 , where $(k + t) > k_3 > k_1$. Thus in (R, k_3) both P_i and P_j consider themselves the leader, violating the assumption that A is an election protocol.

But after a time t , all the processes except P_i and P_j are suspected. Hence there is a time after which every process that crashes is permanently suspected by every correct process. So H_D satisfies Strong Completeness. Consider Accuracy. After a time t , P_i and P_j are never suspected in H_D . Hence H_D satisfies Eventual Strong Accuracy. This is a contradiction. \square

Theorem 2 *A weakest failure detector to solve Election is the Perfect Failure Detector.*

Proof: It is shown in [3] that a failure detector satisfying Strong Accuracy and Strong Completeness can be used to implement a Perfect Failure Detector. Strong Accuracy has processes never suspect a correct process: suspicions are never false. Every correct process always detects a leader failure only when the leader crashes using a Perfect Failure Detector. After an election is started, the problem of electing only one process as a leader is a kind of consensus problem; hence this problem is easily solved with a Strong Failure Detector that is less strong than Perfect Failure detectors. That means that every correct process eventually gets into the state in which it considers only one process to be a leader. Therefore a Perfect Failure Detector is the weakest failure detector that is sufficient to solve Election. \square

4. Concluding Remarks

The importance of this work is in extending the applicability field of the results of Chandra and Toueg [4] on solving problems in asynchronous system (with crash failures and reliable channels) augmented with unreliable failure detectors. The applicability of these results to problems other than Consensus has been discussed in [2,5,6,7,8]. To our knowledge, it is however the first time that Election problems are discussed in asynchronous systems with unreliable failure detectors.

We are not the first to show that there are problems harder than Consensus. The first such result that we are aware of is [9] in which the authors show that Non-Blocking Atomic Commitment (NB-AC) cannot be implemented with the weakest failure detector that can implement Consensus. This problem arises when transactions update data in a distributed system and the termination of transactions should be coordinated among

all participants if data consistency is to be preserved even in the presence of failures [10]. It resembles the Election problem in that NB-AC is harder than Consensus, but it does show that a failure detector weaker than a Perfect Failure Detector is strong enough to solve the problem. Hence, NB-AC appears to be harder than Consensus and easier than Election.

We believe that there are problems harder than Election as well. One can define failure detectors that are stronger than a Perfect Failure Detector. For example, we can define a failure detector that is not only perfect but also guarantees that a failure of a process is detected only after all messages that it has sent have been received by the detecting process. This failure detector is required by some problems, including the non-blocking version of the asynchronous Primary-Backup problem [10].

References

- [1] M. Fischer, N. Lynch, and M. Paterson. Impossibility of Distributed Consensus with One Faulty Process. *Journal of the ACM*, pages 374-382. (32) 1985.
- [2] T. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. Technical Report, Department of computer Science, Cornell Univ., 1994.
- [3] D. Dolev and R. Strong. A Simple Model For Agreement in Distributed Systems. In *Fault-Tolerant Distributed computing*, pages 42-59. B. Simons and A. Spector ed, Springer Verlag (LNCS 448, 1987).
- [4] T. Chandra, V. Hadzilacos and S. Toueg. The Weakest Failure Detector for Solving Consensus. *Proceedings of the 11th ACM Symposium on Principles of Distributed Computing*, pages 147-158. ACM press, 1992.
- [5] R. Guerraoui and A. Schiper. Transaction model vs Virtual Synchrony model: bridging the gap. In *Distributed Systems: From Theory to Practice*, pages 121-132. K. Birman, F. Mattern and A. Schiper ed, Springer Verlag (LNCS 938), 1995.
- [6] V. Hadzilacos. On the relationship between the atomic commitment and consensus problems. In *Fault-Tolerant Distributed Computing*, pages 201-208. B. Simons and A. Spector ed, Springer Verlag (LNCS 448), 1987.
- [7] L. Sabel and K. Marzullo. Election vs. Consensus in Asynchronous Systems. Technical Report TR95-1488, cornell Univ, 1995.
- [8] A. Schiper and A. Sandoz. Primary Partition "Virtually-Synchronous Communication" harder than consensus. In *Proceedings of the 8th Workshop on Distributed Algorithms*, 1994.
- [9] Rachid Guerraoui. Revisiting the relationship between non-blocking atomic commitment and consensus. In *Proceedings of the 10th International Workshop on Distributed Algorithms*, Springer Verlag (LNCS 857), 1996.
- [10] P.A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison Wesley, 1987.