研究課題別評価

1 研究課題名:

大規模分散アルゴリズム開発及び性能評価のツール構築

2 研究者氏名:

Xavier DÉFAGO

3 研究のねらい:

The objective of the research is to better understand the performance tradeoffs associated with fault-tolerant mechanisms for distributed systems. In particular, group communication protocols, such as Total Order Broadcast, are key factors in determining the performance of the system in the absence of failures. While failure-free executions constitute the common case, the occurrence of failures should not affect system performance too drastically, or else failures risk being perceived by the users, thus defeating the objective of masking them. The performance in the face of failures depends mostly on the ability of the system to detect failures promptly and accurately, but this is made difficult by an inherent tradeoff between these two measures. Thus the second objective is to provide a generic failure detection service, the speed and accuracy of which can be best tuned to the specific needs of each part of the entire distributed system.

4 研究成果:

In this research, we have made three major contributions.

Firstly, we have studied several group communication protocols, with a particular focus on the problem of Total Order Broadcast (also called Atomic Broadcast) because it is an important component for many kinds of practical systems, including distributed databases, distributed shared memory, highly-available replicated services, etc.

Secondly, we have proposed a novel failure detection method, called accrual failure detectors, with the ultimate goal of providing failure detection as a generic and highly-configurable service for distributed systems. Studying failure detection is also a prerequisite to understand the actual performance of group communication protocols in the face of failures.

Thirdly, we have developed a communication platform to provide a better support for the evaluation of distributed algorithms.

4.1 Group Communication and Distributed Agreement

Group communication and distributed agreement are basic primitives to maintain the cohesion of a distributed system, by allowing the participating nodes to agree on some common issues. A very practical instance of agreement problems is a communication primitive called *Total Order Broadcast* (also Atomic Broadcast). In short, this primitive allow any of the processes to broadcast

messages any time, but guarantees that all destinations will always see (and process) the messages in the same exact order. Among other things, Total Order Broadcast is a key component for supporting the replication of running programs and services (e.g., web or Grid services). Indeed, assuming that all replicas have the same initial state (i.e., value of variables, etc.), then Total Order Broadcast is used to issue requests to the service. Because of the guarantees offered by the primitive, all replicas perform the same actions in the same sequence, and thus their state change in exactly the same way. The benefit is that the replicas remain exact copies of each others, and thus the service can remain operational even after the crash of some of the replicas, that is, provided that the Total Order Broadcast can actually tolerate the crash of some of the processes.

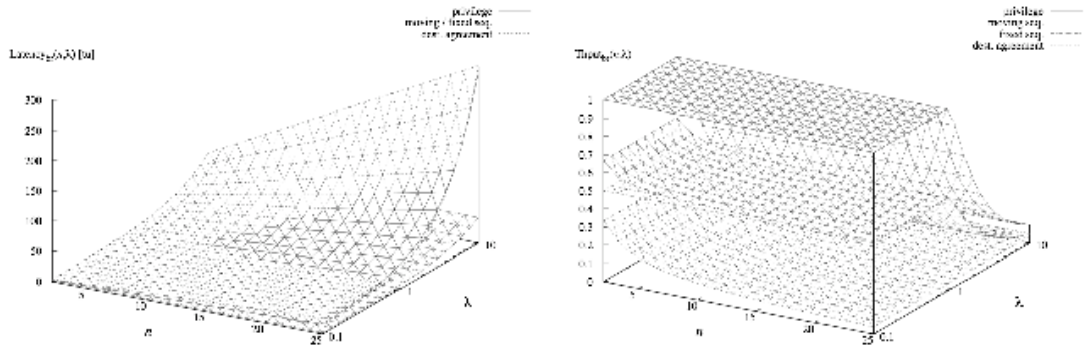

There exist many algorithms to solve Total Order Broadcast, most of which can tolerate failures. However, they do not offer exactly the same guarantees, and their respective performance can vary drastically. We have thus surveyed and analyzed about sixty different Total Order Broadcast algorithms [1], and identified five basic families and a total of thirteen subclasses. The basic families, defined on the decision process used to generate the delivery order, are called *fixed sequencer*, *moving sequencer*, *privilege-based*, *communication history*, and *destinations agreement*.

Using this classification, we have defined representative algorithms for each of the class, and compared their performance and scalability [3] in different network environment. We have also defined a novel replication technique using a variant of a destination agreement Total Order Broadcast algorithm [2].

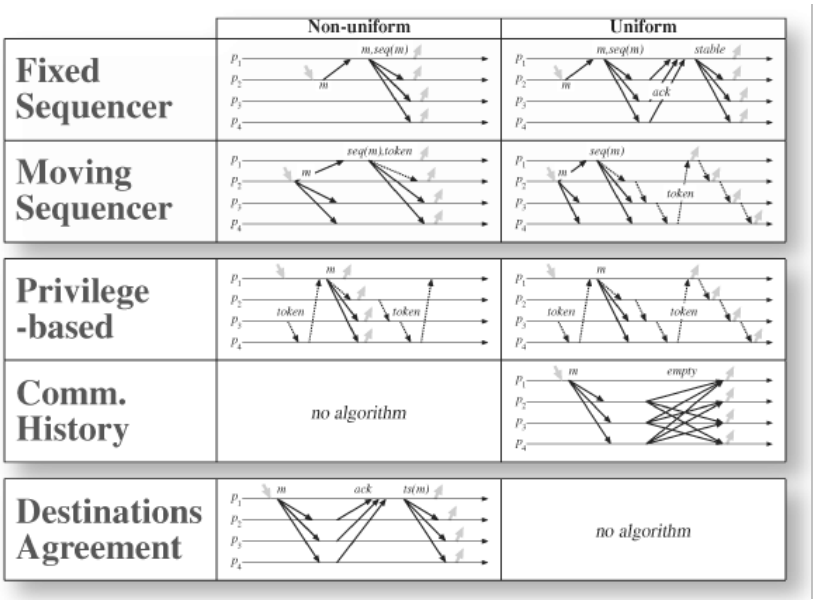

4.2 Accrual Failure Detectors

Failure detection plays an essential role in ensuring fault tolerance in distributed systems. Recently, many people have come to realize that failure detection ought to be provided as some form of generic service.

The performance of fault-tolerant distributed systems, and their ability to mask failures from the viewpoint of their users, depend greatly on the characteristics of the failure detection. This is especially true when it comes to group communication algorithms, such as the Total Order Broadcast discussed in the previous. When considering a complete system, it turns out that failure detection is required at many different levels, but often with very different performance requirements. Roughly speaking, the performance of a failure detector can be expressed by two measures: *detection latency* (i.e., how long it takes until a real failure is detected) and *accuracy* (i.e., how often a running process is erroneously suspected).

In conventional systems, there is always a tradeoff between detection latency (also called detection time) and accuracy, but different operations will benefit from very different settings. For instance, many Total Order Broadcast algorithms will have best overall efficiency with a failure detection latency in the order of a hundred milliseconds, even if the detection is often erroneous. In contrast, a global system reconfiguration will benefit from more accurate detection, even if this leads to a considerable detection latency. It is thus difficult to provide a failure detection scheme that simultaneously provides ideal performance for both.

To address this problem, we have developed the notion of accrual failure detector, promoting a clean decomposition of roles and formally establishing the link with the basic theory on failure detection [4]. Instead of a binary value (trust or suspect), accrual failure detectors associate to each process a real value representing a *suspicion level*. By establishing the clear link with the theory on failure detection, we have identified minimal properties whereby the failure detection

scheme can be used for solving distributed agreement problems, such as Total Order Broadcast and Consensus.

We have developed several implementations of accrual detectors and, in particular, a highly adaptive one called φ accrual failure detector [7]. We have conducted extensive performance measurements, comparing the performance of our φ accrual failure detector with that of other state-of-the-art failure detection schemes. Our results have shown that, for the same latency, our scheme could perform up to ten times more accurately in a local network. Our experiments on an intercontinental network (between Japan and Switzerland) have shown comparable performance in spite of the change of interaction scheme, thus effectively showing the practicality of our approach.

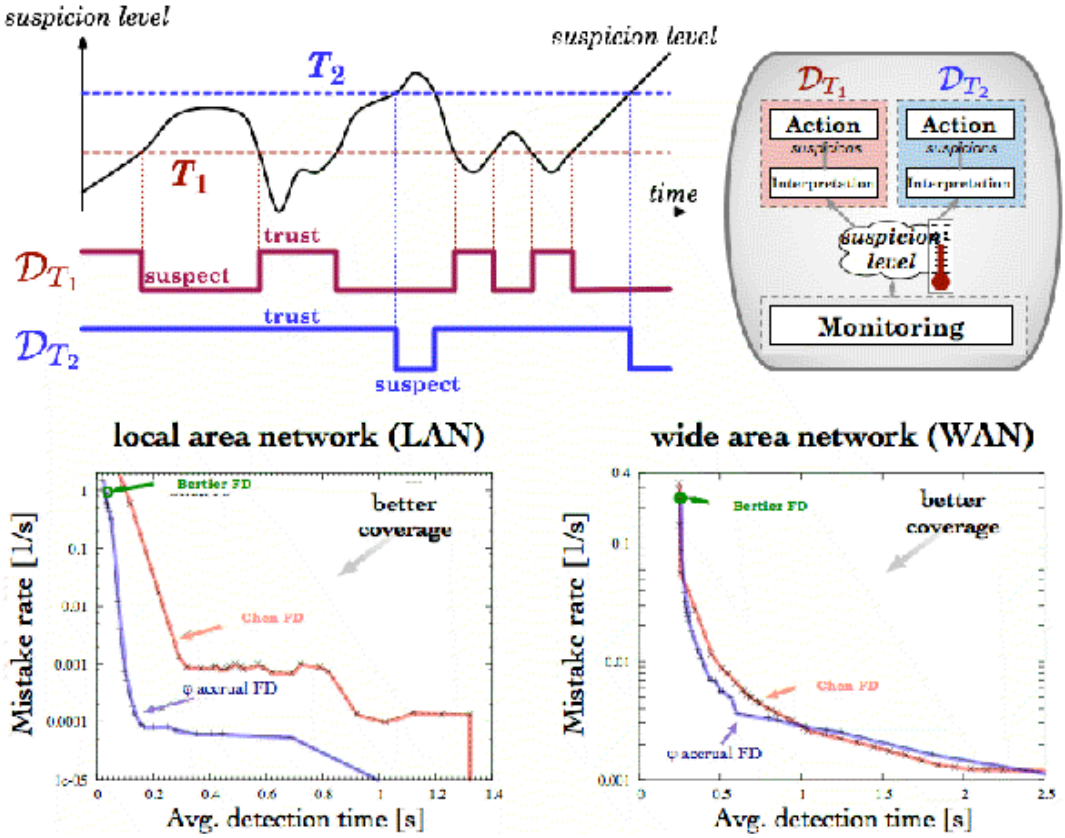

4.3 NekoLS Prototyping Platform

To conduct our experiments more efficiently, we have developed a communication platform called NekoLS. This platform is an extension of an earlier system of us, called Neko, that allowed an easier development of distributed algorithms, and with which the same code can be executed either in a real network environment or on a single machine, within a simulated network. We have made many improvements of this system, the most notable of which is a seamless integration with the SSFNet project; a network simulator aimed for describing large and complex network topologies.

For describing algorithms, the Neko platform is based on a simple layered architecture. While this choice was good for describing simple protocols, it turned out that even moderately complex protocols were very difficult to design elegantly, due for a large part to the difficulty to prevent deadlocks in the protocol. To address this issue, we have been working on a novel mechanism for the composition of micro-protocols.

5 自己評価:

Our research focused on ensuring fault-tolerance in large-scale distributed systems. In particular, we examined three fundamental aspects in relation to this goal: fault-tolerant agreement, failure detection, evaluation/prototyping.

Fault-tolerant agreement is a fundamental component in distributed systems. In particular, we have studied closely the problem of Total Order Broadcast (i.e., broadcast with the guarantee that all messages are received by all destinations in the same global order). We have analyzed about 60 algorithms; identifying their exact guarantees, and defining classes of algorithms [1]. Based on this classification, we have analyzed the performance tradeoffs of these algorithms [3]. In parallel, we have used our understanding of the mechanisms for reaching agreement and developed a novel replication algorithm called semi-passive replication, that is uses less resources than active replication schemes while being more robust than passive ones [2].

When analyzing the performance of agreement problems, it turned out that their performance in the face of computer failures is completely dependent on the performance of an underlying failure detection mechanism. While most systems use simple timeouts mechanisms and detect failures within a few minutes, our research has lead to failure detection that can detect failures in less than a second, even with a world-wide system. In particular, we have developed several failure detection mechanisms of which we have analyzed the performance in local, as well as wide area networks. One such mechanism is the ϕ accrual failure detector [5]. Realizing that fault-tolerant systems actually need to rely on several failure detectors with different performance guarantees, we have defined formally the notion of accrual failure detectors, about which we could derive interesting properties regarding the actual quality of service (i.e., the performance) of the failure detection [4]. These results will provide the basis for building a generic failure detection and monitoring system for large-scale systems.

The NekoLS platform has been very useful for running our experiments and performance analyses, thus meeting its initially intended purpose. In addition, with the experience obtained when building the tool and running the experiments, we have been able to find a much better approach to express distributed protocols in general, so that they can be reusable. That approach consists of an advanced protocol composition mechanism. We have obtained very interesting results about this recently, that are currently under submission to a major international conference.

To sum up, this research project has been very productive in terms of scientific results, with many important contributions to the field of fault-tolerant distributed systems. Most of these

contributions have been published in major international conferences or journals.

6 研究総括の見解:

Defago 氏の研究は，大規模分散システムの高信頼，耐故障性に関するものである．このようなシステムでは，それを構成する計算ノードやネットワークが故障することは日常的であり，このような故障に際してもシステム全体が動作し続けることが重要である．Defago 氏は，最も基本的な全順序ブロードキャスト通信プロトコルのためのアルゴリズムの詳細な検討と分類を行い，新しいアルゴリズムを得ている．また，誤り検出方式に関しては，従来よりも高性能で応用の広い Accrual 誤り検出器を発明した．さらに，これらの成果を実験的に検証するためのシミュレーション環境 NekoLS の開発を行い，これらの方式の有効性を確認している．これらは，国際的評価の高い論文誌や専門家会議の招待講演などで発表されている．高く評価できる研究である．

7 主な論文等:

International Journals

1. X. Défago, A. Schiper, and P. Urbán, Total order broadcast and multicast algorithms: Taxonomy and survey, *ACM Computing Surveys*, 36(4):372-421, December 2004. ACM Press.
2. X. Défago and A. Schiper, Semi-passive replication and Lazy Consensus. *Journal of Parallel and Distributed Computing*, 64(12):1380-1398, December 2004. Elsevier.
3. X. Défago, A. Schiper, and P. Urbán. Comparative performance analysis of ordering strategies in atomic broadcast algorithms. *IEICE Trans. on Information and Systems*, Vol.E86-D, No.12, pp.2698-2709, December 2003.

Refereed International Conferences

4. X. Défago, P. Urbán, N. Hayashibara, T. Katayama. Definition and specification of accrual failure detectors. In *Proc. IEEE/IFIP Intl. Conf. on Dependable Systems and Networks,* pp. 206-215, June 2005. IEEE CS Press.
5. N. Hayashibara, X. Défago, R. Yared, and T. Katayama. The ϕ accrual failure detector. In *Proc. 23rd IEEE Intl. Symp. on Reliable Distributed Systems*, pp. 66-78, October 2004. IEEE CS Press.
6. M. Wiesmann, X. Défago, and A. Schiper. Group communication based on standard interfaces. In *Proc. 2nd IEEE Intl. Symp. on Network Computing and Applications*, pp.140-147, April 2003.

(and many other...)

Important Invited Presentations

2005 年 12 月 5 日 "*Failure Detection in Distributed Systems: Retrospective and recent advances*." **Tutorial**. 6th Intl. Conf. on Parallel and Distributed Computing, Applications and Technologies.

2005 年 7 月 2 日 "*Revisiting Failure Detection for Grid Systems.*." **Invited talk.** 48th meeting IFIP working group 10.4 (dependable computing & fault-tolerance).