Timer-based Composition of Fault-containing Self-stabilizing Protocols
Paper in journals: this is the first page of a paper published in ***Information Sciences***.
[***Information Sciences***] **180,** 1802-1816(2010)


Information Sciences 180 (2010) 1802–1816

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

ELSEVIER

Timer-based composition of fault-containing self-stabilizing protocols

Yukiko Yamauchi [a,*], Sayaka Kamei [b,1], Fukuhito Ooshita [c,2], Yoshiaki Katayama [d,3], Hirotsugu Kakugawa [c,2], Toshimitsu Masuzawa [c,2]

[a] *Graduate School of Information Science, Nara Institute of Science and Technology, 8916-5 Takayama, Ikoma, Nara 630-0192, Japan*
[b] *Department of Information Engineering, Graduate School of Engineering, Hiroshima University, 1-4-1 Kagamiyama, Higashi Hiroshima, Hiroshima 739-8527, Japan*
[c] *Graduate School of Information Science and Technology, Osaka University, 1-5, Yamadaoka, Suita, Osaka 565-0871, Japan*
[d] *Graduate School of Computer Science and Engineering, Nagoya Institute of Technology, Gokiso-cho, Showa-ku, Nagoya 466-8555, Japan*

ARTICLE INFO

Article history:
Received 21 November 2008
Received in revised form 1 October 2009
Accepted 4 October 2009

Keywords:
Distributed system
Fault tolerance
Self-adaptability
Self-stabilization
Fault-containment
Hierarchical composition
Timer
Synchronizer

ABSTRACT

One of the desired properties of distributed systems is self-adaptability against faults. Self-stabilizing protocols provide autonomous recovery from any finite number of transient faults. However, in practice, catastrophic faults rarely occur, while small-scale faults are more likely to occur. Fault-containing self-stabilizing protocols promise not only self-stabilization but also containment of the effect of small-scale faults, *i.e.*, they promise quick recovery and small effect for small-scale faults. Hierarchical composition of self-stabilizing protocols is expected to ease the design of new self-stabilizing protocols. However existing composition techniques for self-stabilizing protocols cannot preserve the fault-containment property of source protocols. In this paper, we propose a novel timer-based hierarchical composition of fault-containing self-stabilizing protocols that preserves the fault-containment property of source protocols. To implement timers, we propose a local neighborhood synchronizer that synchronizes limited number of processes during a short time after a fault without involving the entire network into the synchronization. The proposed composition technique facilitates the design of new fault-containing self-stabilizing protocols and enhances the reusability of existing fault-containing self-stabilizing protocols.

1. Introduction

A distributed system consists of processes communicating with each other by communication links. It is expected to add useful properties to distributed systems, *e.g.*, performance, availability, and scalability. Large scale networks that consist of a large number of processes have became popular such as the Internet and peer-to-peer networks. As the number of processes in a distributed system grows, the distributed system becomes more prone to faults. The effect of faults may spread over the entire network due to the communication among processes and the whole system may be disrupted. This is the reason why fault-tolerance is the major concern when we design distributed systems.

Self-stabilization provides autonomous adaptability against any finite number of *transient faults* (*e.g.*, memory crash at processes). Starting from an arbitrary initial configuration, a self-stabilizing protocol converges to a legitimate configuration where the protocol satisfies its specification. A self-stabilizing protocol guarantees autonomous adaptability by considering

* Corresponding author. Tel.: +81 743 72 5253; fax: +81 743 72 5259.
E-mail addresses: y-yamauchi@is.naist.jp (Y. Yamauchi), s-kamei@se.hiroshima-u.ac.jp (S. Kamei), f-oosita@ist.osaka-u.ac.jp (F. Ooshita), katayama@nitech.ac.jp (Y. Katayama), kakugawa@ist.osaka-u.ac.jp (H. Kakugawa), masuzawa@ist.osaka-u.ac.jp (T. Masuzawa).
[1] Tel.: +81 82 424 7685; fax: +81 82 422 7195.
[2] Tel.: +81 6 6879 4118; fax: +81 6 6879 4119.
[3] Tel.: +81 52 735 5576; fax: +81 52 735 5465.

0020-0255/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.ins.2009.10.003

The following is a comment on the published paper shown on the preceding page.

Timer-based Composition of Fault-containing Self-stabilizing Protocols

MASUZAWA Toshimitsu and KAKUGAWA Hirotsugu
(Graduate School of Information Science and Technology)

Introduction

Large scale distributed systems have been developed recently. As the number of processes in a distributed system grows, the distributed system becomes more prone to faults. Self-stabilization [1] provides autonomous adaptability to any finite number and any kind of transient faults (e.g., memory soft error at processes). Even when a distributed system is corrupted by transient faults and gets into an unexpected configuration, a self-stabilizing system can autonomously recover its desired behavior without any external intervention. Although self-stabilization promises autonomous adaptability to any scale of transient faults, the adaptability to small scale faults is more important in practice. This is because catastrophic faults rarely occur in practical environments and small scale faults are more likely to occur. Nevertheless, self-stabilization promises nothing during the recovery and the effect of small scale faults can spread over the entire network. A self-stabilizing system can be contaminated entirely even by small scale faults while we expect that the system can recover quickly with small effect from small scale faults.

A fault-containing self-stabilizing protocol [2] promises self-stabilization against large scale faults and fault-containment against small scale faults (Figure 1). Starting from any configuration corrupted by f processes or less, an f-fault-containing protocol reaches a legitimate configuration with small effect and in short time, i.e., both the recovery time and the number of processes affected by the faults are proportional to f or less. So, the fault-containment property improves the adaptability of self-stabilization to small scale faults.

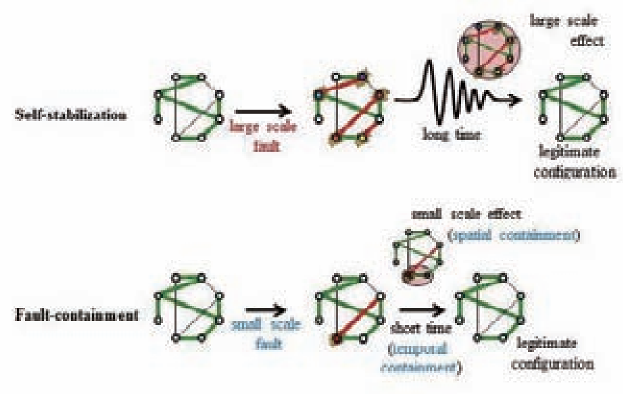

Figure 1 Self-stabilization and Fault-containment

Hierarchical composition of protocols facilitates the design of new protocols. In hierarchical composition of two (or more) protocols, the output of one protocol (called the lower protocol) is used as the input to the other (called the upper protocol), and the obtained protocol provides the output of the upper protocol for the input to the lower protocol.

Hierarchical composition of protocols is commonly used to relieve the difficulty in designing self-stabilizing protocols. Different from composition of classical (or non-self-stabilizing) protocols, protocol composition based on sequential execution of protocols is impossible for self-stabilizing protocols. Instead, the adaptability to any configuration of self-stabilizing protocols allows us to composite protocols based on parallel execution of protocols. However, the composition technique cannot preserve the fault-containment property of the source protocols. This is because the parallel execution of the source protocols allows the upper protocol to execute its actions before stabilization of the lower protocol, that is, the upper protocol can work on an incorrect intermediate output of the lower protocol (Figure 2).

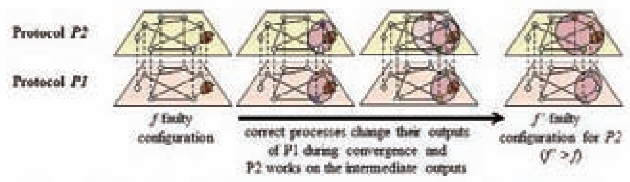

Figure 2 Composition based on parallel execution cannot preserve the fault-containment property

Containment-preserving composition of self-stabilizing protocols

In this paper, we propose, as a novel composition technique of fault-containing self-stabilizing protocols, a containment-preserving composition technique, RWFC-LNS (Recovery Waiting Fault-containing Composition with the Local Neighborhood Synchronizer). The composition technique follows a general strategy, RWFC strategy, which was previously proposed by the authors [3]. The RWFC strategy realizes containment-preserving composition of self-stabilizing protocols by forcing the upper protocol to stop its execution until the lower protocol completes the recovery from a faulty configuration. This strategy guarantees that the upper protocol always works on the correct input from the lower protocol. Therefore, the upper protocol can recover from a faulty configuration with keeping its fault-containment property. The key to implementation of the RWCF strategy is how the waiting at the upper protocol is realized. In the RWFC-LNS technique, the waiting at the upper protocol is realized using a synchronized timer at each process. Since we consider asynchronous systems, we design self-stabilizing syn-