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

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

▲Reprinted from ***Information Sciences,*** 180, Yukiko Yamauchi et al., Timer-based composition of fault-containing self-stabilizing protocols, 1802-1816, Copyright(2010), with permission from Elsevier.

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.

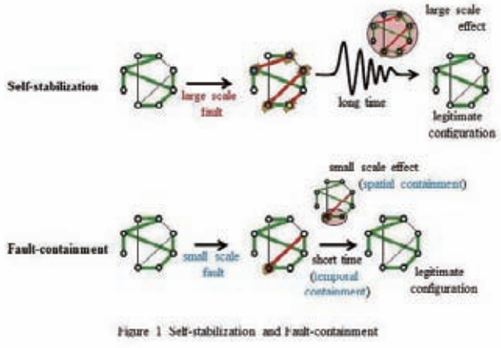

Figure 1 Self-stabilization and Fault-containment

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).

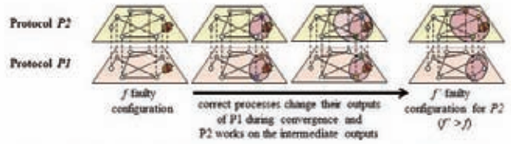

Figure 2 Composition based on parallel execution cannot preserve the fault-containment property

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-

chronization protocol for the composition.

Before presenting the RWFC-LNS technique, we give the definitions of an f-fault-containing protocol and the measures for the fault-containment property. A configuration is f-faulty configuration if a legitimate configuration can be obtained by changing the states of f processes. To measure the containment property in recovery from f-faulty configurations, the recovery time and the contamination radius are considered. The recovery time is the time required to reach a legitimate configuration and the contamination radius is the maximum distance to a contaminated process from its nearest faulty process, where a contaminated process is the process that changes its state during the recovery. A self-stabilizing protocol is f-fault-containing if the recovery time and the contamination radius for any f-faulty configuration depend on f (not the total number of processes).

We show the outline of the RWFC-LNS technique for the composition of two self-stabilizing fault-containing protocols P1 and P2. In the following, let P1 be an f1-fault-containing protocol with recovery time r1 and the contamination radius c1, and P2 be an f2-fault-containing protocol with recovery time r2 and the contamination radius c2. As described below, the RWFC-LNS technique generates an f'-fault containing protocol composed from P1 and P2 where f'= min{f1, f2}.

Starting from an f'-faulty configuration, a process finding inconsistency in P1 or P2 initiates a self-stabilizing local neighborhood synchronizer. The local neighborhood synchronizer can synchronize timers at processes in the neighborhood of the initiator. Using the synchronized timer, each process in the contamination radius of P1 and P2 first executes only P1 for r1 rounds (with stopping its execution of P2). Note that if a faulty process p has a correct process q as its neighbor, p or q can find the inconsistency between them in P1 or P2. During the first r1 rounds (or the recovery time of P1 for f1-faulty configurations (f1≥f')), these processes execute only P1, and P1 reaches a legitimate configuration. After that, these processes start their execution of P2 on the correct input from P1. These processes execute only P2 for the r2 rounds (or the recovery time of P2 for f2-faulty configurations (f2≥f')). Figure 3 illustrates the outline of the RWFC-LNS strategy.

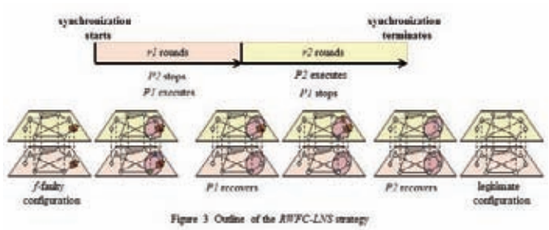

Figure 3 Outline of the RWFC-LNS strategy

The features of the local neighborhood synchronizer are summarized as follows. Once initiated, the local neighborhood synchronizer makes the processes within the distance max{c1, c2}+min{f1, f2}+1 set their timer values to r1+r2 and then decrease the values one by one. The decrement of the timer values is executed in a synchronized fashion; the processes keep difference of timer values between neighboring processes at most one. Even when started from any configuration, the self-stabilizing synchronizer reaches a legitimate configuration (where timer values are synchronously decreased) in O(r1+r2) rounds.

From the above, given an f1-fault-containing protocol P1 and an f2-fault-containing protocol P2, RWFC-LNS provides a min{f1, f2}-fault-containing protocol. The contamination radius of the obtained protocol is O(max{c1, c2}+min{f1, f2}). The recovery time of the obtained protocol is O(r1+ r2).

Conclusions

We proposed a novel timer-based fault-containing composition technique RWFC-LNS for fault-containing self-stabilizing protocols that guarantees containment of the effect of faults during the recovery. The proposed composition technique facilitates the design of new fault-containing protocols and enhances reusability of existing fault-containing protocols, which reduces the burden of protocol designers.

The RWFC-LNS technique utilizes the temporal containment property of fault-containing protocols while the previous technique proposed by the authors [3] utilizes the spatial containment. To implement the timers at processes, we designed a local neighborhood synchronizer.

One of the most important applications of the proposed method is ring embedding in an arbitrary network. Ring is one of the most investigated networks in distributed computing and many fault-containing protocols have been designed for rings. The authors [4] proposed a ring embedding on an arbitrary rooted tree that preserves fault-containment property of protocols executed on the embedded ring. We can execute fault-containing protocols designed for rings on an arbitrary network by composing the ring embedding and fault-containing spanning tree construction by the proposed composition technique.

References

[1] S. Dolev, Self-stabilization, MIT Press, Cambridge, MA, 2000.

[2] S. Ghosh, A. Gupta, T. Herman, S.V. Pemmaraju, Fault-containing self-stabilizing algorithms, in Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing, 1996, pp. 45–54.

[3] Y. Yamauchi, S. Kamei, F. Ooshita, Y. Katayama, H. Kakugawa, T. Masuzawa, Composition of fault-containing protocols based on recovery waiting fault-containing composition framework, in Proceedings of the 8th International Symposium on Stabilization, Safety, and Security of Distributed Systems, 2006, pp. 516–532.

[4] Y. Yamauchi, T. Masuzawa, D. Bein, Preserving the fault-containment property of ring protocols executed on trees, The Computer Journal, 52 (4), 2009, pp. 483–498.

Direct Visualization of Secondary Structures of F-actin by Electron Cryomicroscopy

Paper in journals: this is the first page of a paper published in ***Nature***.

[*Nature*] 467, 724-728(2010)

LETTER

doi:10.1038/nature09372

Direct visualization of secondary structures of F-actin by electron cryomicroscopy

Takashi Fujii[1], Atsuko H. Iwane[1], Toshio Yanagida[1] & Keiichi Namba[1]

F-actin is a helical assembly of actin, which is a component of muscle fibres essential for contraction and has a crucial role in numerous cellular processes, such as the formation of lamellipodia and filopodia[1,2], as the most abundant component and regulator of cytoskeletons by dynamic assembly and disassembly (from G-actin to F-actin and vice versa). Actin is a ubiquitous protein and is involved in important biological functions, but the definitive high-resolution structure of F-actin remains unknown. Although a recent atomic model well reproduced X-ray fibre diffraction intensity data from a highly oriented liquid-crystalline sol specimen[3], its refinement without experimental phase information has certain limitations. Direct visualization of the structure by electron cryomicroscopy, however, has been difficult because it is relatively thin and flexible. Here we report the F-actin structure at 6.6 Å resolution, made obtainable by recent advances in electron cryomicroscopy. The density map clearly resolves all the secondary structures of G-actin, such as α-helices, β-structures and loops, and makes unambiguous modelling and refinement possible. Complex domain motions that open the nucleotide-binding pocket on F-actin formation, specific D-loop and terminal conformations, and relatively tight axial but markedly loose interprotofilament interactions hydrophilic in nature are revealed in the F-actin model, and all seem to be important for dynamic functions of actin.

Many of the atomic models of F-actin proposed over the years[3–9] were obtained by maximizing the agreement between experimental intensities of X-ray fibre diffraction and those calculated from a model. However, because the diffraction patterns are cylindrically averaged and layer lines are broadened in an arc owing to the finite disorientation, it is extremely difficult to obtain individual Fourier–Bessel components of the structure factors for three-dimensional (3D) density reconstruction. That is why these analyses have been done by building many models, calculating layer-line intensities and maximizing the agreement between observed and calculated intensities to find a best possible model, except in one case with tobacco mosaic virus[6]. Therefore, there is always some concern as to whether or not such models are unique.

Technical advances in the electron cryomicroscopy (cryoEM) and image analysis of frozen hydrated specimens in recent years have allowed the structural analysis of helical assemblies of biological macromolecules, such as the bacterial flagellar filament, at near-atomic resolution by aligning and averaging tens of thousands of molecular images using helical symmetry[7–9]. CryoEM image analyses of F-actin have also been carried out but only up to 13 Å resolution[3,5,10], so the atomic model is of limited accuracy in studying the mechanisms of actin polymerization and depolymerization. Because F-actin is a flexible, ribbon-like filament with a diameter of 100 Å and a mass-per-length of 1.5 kDa $Å^{-1}$, the image contrast of frozen hydrated specimens is markedly lower than that of thicker tubular structures of the flagellar filament (230 Å, 11.0 kDa $Å^{-1}$) and tobacco mosaic virus (180 Å, 12.6 kDa $Å^{-1}$), making accurate image alignment and high-resolution analysis extremely difficult. However, recent technical advances[11] have allowed us to obtain a 3D density map of skeletal muscle F-actin at 6.6 Å resolution.

Key to this is our use of an in-column Ω-type energy filter and charge-couple-device (CCD) camera as well as a field-emission gun and a liquid-helium-cooled specimen stage. A remarkable ~5-fold gain in contrast has been achieved by energy filtering, controlling ice thickness and using a specimen temperature of 50 K instead of 4 K (ref. 11). We can now see the two-stranded helical features of F-actin even at small defocus levels (Fig. 1a). The CCD camera made the collections of high-quality images far more efficient, but to avoid undesirable removal of high-resolution contrast by its poor modulation transfer function[12] we used a magnification of 172,000 × (0.87 Å per pixel). We collected 490 such high-quality images in two days. We used a single-particle image analysis method using helical symmetry[11,13]. Because we fully automated the whole procedure including corrections for the contrast transfer function, the image analysis was completed within two days. The helical symmetry and the axial repeat distance were refined in the iterative process of image analysis and converged to a subunit rotation of −166.6°, corresponding to a helical symmetry of between 28/13 and 13/6 (subunits/turn), and an axial repeat of 27.6 Å. The resolution was 6.6 Å at a Fourier shell correlation of 0.143 (ref. 14; Supplementary Figs 1 and 2). Layer lines are visible out to 6.8 Å in the power spectrum of the 3D reconstruction (Fig. 1b; see also Supplementary Fig. 3). The variance map shows the reliability of the density map as well as the rigid helical backbone of the filament (Supplementary Fig. 4).

Because the resolution was better than that needed to identify individual secondary structures of actin, including loops and an extended amino-terminal chain previously unresolved in the crystal structure (Fig. 2 and Supplementary Movie 1), we were able to build a highly reliable atomic model of F-actin. The situation is analogous to protein crystallography at 3 Å resolution, where side chains can be identified and used to build a reliable atomic model. Although the flexibility of F-actin, especially in its helical order, has been debated over the years[10], the high-resolution map obtained by using over 97% of the collected images indicates that F-actin is not so flexible. Multireference alignment with reference 3D volumes of F-actin with different helical symmetries produced a narrow angular distribution (Fig. 1c), indicating that the twisting variability by thermal motion is less than ~1°. This also demonstrates the importance of careful cryospecimen preparation with an automated vitrification device (Methods).

To build a reliable atomic model of F-actin, we used FlexEM[15], which refines the atomic model by fitting it into the electron microscopy map by simulated annealing molecular dynamics with stereochemical and non-bonded interaction terms restrained. We used the crystal structure of uncomplexed actin[16] (Protein Data Bank ID, 1J6Z) as the model representing G-actin and divided it into four domains (Fig. 3), which have been conventionally named subdomains[17]. We call them domains because domains 1, 3 and 4 have well-defined hydrophobic cores and behave as rigid, independent units on conformational change from G-actin[18] to F-actin (Fig. 3a), as indicated by the small root-mean-squared displacements of Cα atoms (Supplementary Table 1). This in turn ensures the correctness and reliability of the present F-actin model. The model and the refinement process are shown in Fig. 2 and Supplementary Movie 2, respectively.

[1]Graduate School of Frontier Biosciences, Osaka University, 1-3 Yamadaoka, Suita, Osaka 565-0871, Japan.