Self-Stabilization Workshop

Shing-Tsaan Huang
Program Chair
National Central University, Taiwan

Ted Herman
Publication Chair
University of Iowa

Abstract

The DSN Workshop on Self-Stabilization's programme includes fifteen research presentations. The main areas in the programme are network protocols, sensor networks, distributed algorithms, methods for analysis of self-stabilization, distributed system fault tolerance, and techniques used in the construction of systems that self-stabilize.

1. Workshop Goals and Areas

Self-stabilization contrasts with other approaches to dependable systems: instead of entirely masking failures by replication or safe storage, or using backward recovery to restore system state, self-stabilization is pure forward recovery, tolerating any degree of transient failure. Though self-stabilization has been a standard algorithmic topic of distributed systems, there is now renewed interest due to emerging concerns of scalability and more extreme distribution of system control. There is also recognition that system components become more valuable if they are individually self-stabilizing.

The Workshop on Self-stabilizing Systems has been a biennial workshop. The workshop has been a high-quality research forum held in collaboration with other major conferences in related areas: ACM STOC (1995), ACM PODC (1997), IEEE ICDCS (1999), DISC (2001), and many workshop results have been subsequently published in journals (including special issues devoted to self-stabilization). Meetings of researchers interested in self-stabilization have been stimulating and useful. At workshops, new adaptations of stabilization have been proposed, new technical advances in algorithms and analysis have been discussed, and perhaps most important, new problem areas and applications suited to the stabilization approach became known to a wider community. This workshop continued these themes with high-quality research presentations and open discussion of the latest results and most exciting new directions. In addition to the fifteen technical presentations summarized below, an invited presentation was given by Anish Arora.

2. Summary of Presentations

We start by summarizing some topical contributions, inspired by biological and emerging sensor network areas. In "Self-stabilizing pulse synchronization inspired by biological pacemaker networks" (Ariel Daliot, Danny Dolev, Hanna Parnas), the problem of achieving synchronization of periodic pulse events of nodes is investigated. It is a phenomenon of many distributed biological systems, and this paper shows how to deal with Byzantine behavior of the processes. The presentation "A pursuer-evader game for sensor networks" (Murat Demirbas, Anish Arora, Mohamed Gouda) researches several approaches to tunable, energy-efficient and self-stabilizing pursuer strategies in a pursuer-evader game in sensor networks. This is a new area where self-stabilization is expected to have a large impact. One approach is evader-centric; another is pursuer-centric; and a hybrid approach is tunable between these two poles. The sensor network theme continues in "Collision-free communication in sensor networks" (Sandeep S. Kulkarni, Umamaheswaran Arumugam) where the problem of diffusing information in a wireless sensor/actuator network while avoiding collision is solved. By using a pattern of staged time-division multiplexing in a network, messages propagate without colliding.

The following papers, presented at the workshop, study network or distributed system protocols. In "A stabilizing solution to the stable paths problem" (Jorge A. Cobb, Mohamed G. Gouda, Ravi Musunuri) the issue is routing stability, which is a problem in routing protocols such as BGP. A technique of diffusing computation enables nodes to select paths so as to avoid instabilities. "Route preserving stabilization" (Colette Johnen, Sébastien Tixeuil) presents a self-stabilizing shortest path spanning tree algorithm for semi-uniform systems with a distributed de-

Proceedings of the 2003 International Conference on Dependable Systems and Networks (DSN'03)
0-7695-1959-8/03 $17.00 (c) 2003 IEEE

mon. The algorithm is adaptive to edge cost changes: it is loop free, meaning that any such change will maintain a tree, and route preserving, meaning that every message sent to the root is eventually received even if continuous cost changes occur. Shlomi Dolev and Elad Schiller continue work from the 2001 workshop in "Self-stabilizing group communication in directed networks", which generalizes previous work that considered undirected communication links. The algorithms in this presentation cover several algorithm variations, including token circulation in a virtual ring and spanning tree structures.

Distributed system issues motivate much of the foundational work on self-stabilization. "A framework of safe stabilization" (Sukumar Ghosh, Alina Bejan) examines the intersection of self-stabilization and system safety (a core concern in dependable systems). Under certain conditions, systems can move to a known safe state, even if full restoration of legitimate behavior is not possible. The paper "Self-stabilizing atomicity refinement allowing neighborhood concurrency" (Sébastien Cantarell, Ajoy K. Datta) considers a problem of resource sharing in a network model. By using a self-stabilizing transformer from a high-level model to a more realistic one, the solution allows more concurrency, while hiding low-level activity from higher layers. Resource allocation is also a subject in "A new self-stabilizing k-out-of-ℓ exclusion algorithm on rings" (Ajoy K. Datta, Rachid Hadid, Vincent Villain), which proposes a new metric for the evaluation of ℓ-exclusion. "A self-stabilizing algorithm for token circulation by using edge-tokens" (Shing-Tsaan Huang, Su-Shen Hung) presents a uniform token-circulation algorithm using edge tokens. The proposed work is better in some respect (either uniformity or the stabilization time) than all previous algorithms for the same problem and model. The algorithm uses a uniform technique to fairly traverse a token around a tree with $O(n)$ convergence time.

Many of the papers already mentioned contribute to the store of techniques known to stabilize distributed systems. Some research presentations focus especially on methods for constructing self-stabilizing systems and algorithms. From the presentation "A general method for constructing self-stabilizing algorithms with logarithmic time complexity" (Felix C. Gärtner, Henning Pagnia) we see how the convergence time complexity of self-stabilization can be reduced, and fault containment improved, by carefully partitioning system components, structuring them into a hierarchy, and using an overlay network to restrict communication. "Lyapunov analysis of neural network stability in an adaptive flight control system" (Sampath Yerramalla, Bojan Cukic, Martin Mladenovski, Edgar Fuller) emphasizes the link between system adaptivity and self-stabilization, featuring a case study. The Lyapunov function chosen here is the network's quantization error, which must decrease for asymptotic stability to hold.

Theoretical studies are represented by several presentations. "Self-stabilizing algorithms for $\{k\}$-domination" (Martin Gairing, Stephen T. Hedetniemi, Petter Kristiansen, Alice A. McRae) develops an algorithm to achieve the graph-theoretic $\{k\}$-dominating function (also giving a linear-time algorithm for a $\{2\}$-dominating function). The presentation "A method for evaluating efficiency of protocols on the asynchronous shared-state model" (Yoshihiro Nakaminami, Toshimitsu Masuzawa, Ted Herman) is a technical result; its contribution is a characterization of algorithms for which the asynchronous complexity can be derived from the (more readily calculated) synchronous time complexity. "An improved snap-stabilizing PIF algorithm" (Lélia Blin, Alain Cournier, Vincent Villain) removes the assumption of knowing the network size from a snap-stabilizing solution to Propagation of Information with Feedback in a network of processes.

In addition to these prepared presentations, the workshop programme also makes time for a general discussion of future directions and open problems. This is an opportunity to exchange ideas and suggest new challenges and opportunities for the unique type of fault tolerance that self-stabilization offers. The workshop papers are published by Springer in the Lecture Notes on Computer Science (LNCS) series.

Proceedings of the 2003 International Conference on Dependable Systems and Networks (DSN'03)
0-7695-1959-8/03 $17.00 (c) 2003 IEEE