BIRS Workshop
Self-Stabilizing Distributed Systems
October 2 - October 7, 2004

ABSTRACTS
(in alphabetic order by speaker surname)

Speaker: **Uri Abraham** (Ben-Gurion University)
Title: *Self-Stabilizing TimeStamps*
Abstract:

Speaker: **Anish Arora** (Ohio State)
Title: *Self-Stabilizing Sensor Networks*
Abstract:

Speaker: **Alina Bejan** (University of Iowa)
Title: *Designing Self-Optimizing DHTs using Request Profiling*
Abstract:
Various studies on request patterns in P2P networks have confirmed the existence of the interest-based clusters. Some P2P networks that exhibit the small-world phenomenon contains clusters of peers that frequently communicate with one another, although this is not always true. However, the existence of interest-based clusters opens up the possibility of more efficient routing. In this paper we consider the problem of designing a *self-optimizing overlay network* and routing mechanisms to permit efficient location of resources by the *periodic profiling* of request patterns. Our self-optimization protocol uses selective replication of resources for restricting the sizes of the clusters, and proposes the deployment of inactive nodes for further reduction of the routing latency. The self-optimization protocol is demonstrated on the Chord network. It leads to a routing latency that scales with the size of the clusters.

Speaker: **Christian Boulinier** (LaRIA, CNRS University of Picardie)
Title: *When Graph Theory helps self-stabilization*
Abstract:
We propose a self-stabilizing and generic scheme to solve synchronization problems whose safety specification is defined from a local property. This scheme is the basis of several very efficient algorithms in terms of memory space for problems like asynchronous phase clock, local mutual exclusion, reader-writers and local group mutual exclusion. We show that all this algorithms use a phase clock whose minimum size in terms of number of states per processis at most equal to$C_g + T_g - 1$ where C_g is the length of a maximal cycle of a shortest maximalcycle basis if the graph contains cycles, 2 for tree networks, and T_g is the length of the longuest chorless cycle (hole) if the graph contains cycles, 2 for trees. In paticular, for asynchronous phase clock problem, our solution significantly improves all the best self-stabilizing solution in the literature which are quadratic in number of states. We also proposed a silent bounded generic algorithm which can be used to transform any serial system to a distributed one. This talk was presented in PODC 2004.

Speaker: **Olga Brukman** (Ben-Gurion University, Beer-Sheva, Israel)
Title: *Self Stabilizing Autonomic Recoverers*
Abstract:
We suggest to model software package flaws (bugs) by assuming eventual Byzantine behavior of the package. In particular, the package has been tested by the manufacturer for limited length scenarios when started in a predefined initial state; the behavior beyond the tested scenario may be Byzantine. Restarts (reboots) are useful for recovering such systems.

We suggest a general yet practical framework and paradigm, based on a theoretical foundation, for the monitoring and restarting of systems. An autonomic recoverer that monitors and restarts the system is proposed. The autonomic recoverer is designed to handle a task given specific task requirements in the form of predicates and actions. A DAG subsystem hierarchy structure is used by a consistency monitoring procedure in order to achieve gracious recovery. The existence and correct functionality of the autonomic recovery is guaranteed by the use of a kernel resident (anchor) process; the design of this process is self-stabilizing. The autonomic recoverer uses a new scheme for liveness assurance via on-line monitoring that complements known schemes for on-line safety ensurance.

Speaker: **Murat Demirbas** (MIT)
Title: *Designing and Implementing Self-stabilizing Algorithms for Wireless Sensor Networks*
Abstract:
The wireless sensor network regime imposes certain constraints on the algorithms, such as energy limitation of the nodes, lack of dedicated point-to-point channels between nodes, limitations on the bandwidth, the self-configurability requirement, and unreliable nature of the sensors and channels, that the algorithm designers have to deal with. On the other hand, this regime also offers new features, such as the atomic radio broadcast primitive and reasonably precise timers at each node, that the algorithm designers can exploit.

In my talk I will give samplers of self-stabilizing algorithms that satisfy the former constraints while exploiting the latter features. I will also talk about how these self-stabilizing algorithms can be compiled/refined into self-stabilizing implementations.

Speaker: **Shlomi Dolev** (Ben Gurion University, Israel)
Title: *A Way to Lead the Self* Initiatives*
Abstract:

Speaker: **Mohamed G. Gouda** (The University of Texas at Austin)
Title: *Sentries and Sleepers in Sensor Networks*
Abstract:
A sensor is a battery-operated small computer with an antenna and a sensing board that can sense magnetism, sound, heat, etc. Sensors in a network can use their antennas to communicate in a wireless fashion by broadcasting messages over radio frequency to neighboring sensors in the same network. In order to lengthen the relatively short lifetime of sensor batteries, each sensor in a network can be replaced by a group of n sensors, for some $n \geq 2$. The group of n sensors act as one sensor, whose lifetime is about n times that of a regular sensor as follows. For a time period, only one sensor in the group, called *sentry*, stays awake and performs all the tasks assigned to the group, while the remaining sensors, called *sleepers*, go to sleep to save their batteries. At the beginning of the next time period, the sleepers wake up, then all the sensors in the group elect a new sentry for the next time period, and the cycle repeats. In this paper, we describe a practical protocol that can be used by a group of sensors to elect a new sentry at the beginning of each time period. Our protocol, unlike earlier protocols, is based on the assumption that the sensors in a group are perfectly identical (e.g. they do not have unique identifiers; rather each of them has the same group identifier). This feature makes our protocol resilient against any attack by an adversary

sensor in the group that may lie about its own identity to be elected a sentry over and over, and keep the legitimate sensors in the group asleep for a long time.

Speaker: **Ted Herman** (University of Iowa)
Title: *Self-Stabilization and Fault-Containment of Clock Synchronization*
Abstract:
The Internet uses the NTP protocol to keep time. NTP is tailored to the behavior of the Internet and its model of communication. In wireless, ad hoc networks, the Internet model can be inappropriate and the needs for synchronization have different motivation than typical Internet applications. Moreover, ad hoc networks are dynamic, often with less reliable hardware than used by most Internet infrastructure, so new clock synchronization protocols are of practical interest, and self-stabilization is a desired property. This talk describes ongoing work in implementing a clock synchronization protocol guided by principles from self-stabilization and fault containment.

Speaker: **Shing-Tsaan Huang** (National Central University, TAIWAN)
Title: *A memory-efficient, self-stabilizing algorithm for constructing spanning trees.*
Abstract:
This talk discusses a memory-efficient SS algorithm for constructing spanning trees in an asynchronous way. The proposed algorithm is semi-uniform. There is a unique special node r. Every node except r has the same behavior. The constructed tree is rooted at r. Conventional algorithms use a level variable d for detecting and breaking-down cycles, where d has space complexity of the system size. Due to variable d, the conventional algorithms cannot allow the system expand into a large one. Johnen solved the problem in 1997; her solution uses 16 states and 2 pointers per node. Here, our algorithm improves to 12 states and 1 pointer per node. The pointer for a node in both algorithms points to one of the neighbors for the node, and hence has the space complexity of the degree of the node, instead of the system size.

Speaker: **Yoshiaki Katayama** (Nagoya Institute of Technology)
Title: *A yet another self-stabilizing algorithm for load balancing on rooted trees*
Abstract:
We propose a self-stabilizing load balancing algorithm on rooted tree networks. Given arbitrary distribution of integral loads over the processors, the algorithm eventually converges to a 1-balancing configuration where each processor has an integral load such that the difference of loads between any two processors are at most one. The remarkable feature of the algorithm is that it uses only "pure local" information: the behavior of each process depends only on its own load and the loads of its neighbors.

Speaker: **Fredrik Manne** (University of Bergen, currently on sabbatical leave at Sandia National Laboratories and University of New Mexico)
Title: *Efficient Generic Multi-Stage Self-Stabilizing Algorithms for Trees*
Abstract:
We define a general framework for developing efficient multi-stage self-stabilizing algorithms for trees. This is achieved by combining multiple self-stabilizing algorithms (stages) in such a way that the local stabilized values resulting from execution of one stage drive the predicates of subsequent stages. Independently each of our stages has a moves complexity of $\Theta(n2)$. Contrary to what one might expect, the moves complexity of any self-stabilizing algorithm made up of k of our stages is significantly less than the expected multiplicative combined moves complexity of $O(n^{2k})$, coming in at only $O(n^{k+1})$. This provides an improvement of several orders of magnitude over a number of previously published self-stabilizing algorithms.

Speaker: **Toshimitsu Masuzawa** (Osaka University)
Title: *Self-Stabilizing Link-Coloring Resilient to Byzantine Faults*
Abstract:
This paper considers self-stabilizing link-coloring resilient to (permanent) Byzantine faults.

First, we consider the problem in tree networks. We present a protocol under the central daemon that uses $\Delta+1$ colors where Δ is the maximum degree in the network. This protocol achieves fault containment with radius of two against the Byzantine faults. Moreover, we show necessity of $\Delta+1$ colors and the central daemon to achieve fault containment with a constant radius.

Second, we consider the problem in arbitrary networks. We present a protocol that uses $2\Delta - 1$ colors and achieves fault containment with radius of one against the Byzantine faults.

Speaker: **Mikhail Nesterenko** (Kent State University)
Title: *Secure Location Verification and Stabilization*
Abstract:
Secure location verification is a recently stated problem that has a number of practical applications. The problem requires a wireless sensor network to confirm that a potentially malicious *prover* is located in a designated area. The original solution to the problem, as well as solutions to related problems, exploits the difference between propagation speeds of radio and sound waves to estimate the position of the prover. In this talk I discuss a solution that leverages the broadcast nature of the radio signal emitted by the prover and the distributed topology of the network. The solution is versatile and deals with provers using either omni-directional or directional propagation of radio signals without requiring any special hardware besides a radio transceiver. I estimate the bounds on the number of sensors required to protect the areas of various shapes and extend our solution to handle complex radio signal propagation, optimize sensor placement and operate without precise topology information.

To successfully carry out location verification, the sensors have to acquire system and geometric information. I address the issues of how this information can be effectively gathered and how the system can stabilize from faults and state corruptions.

Speaker: **Elad Schiller** (Ben Gurion Universiy, Israel.)
Title: *Self-Stabilizing Group Communication*
Abstract:
Group communication services are becoming widely accepted as useful building blocks for the construction of fault-tolerant distributed systems and communication networks. Group communication systems enable processors that share a collective interest to identify themselves as a single logical communication endpoint.

Group communication services, like other middleware services, are long-lived, on-line tasks. Therefore, it is very unlikely that a transient fault will never occur. Self-stabilizing group communication services can automatically recover following the occurrence of unexpected faults.

We describe algorithms for self-stabilizing group communication services for three different settings; asynchronous system, directed networks, and ad-hoc networks.

Speaker: **Sébastien Tixeuil** (LRI - CNRS and INRIA Grand Large)
Title: *On Self-stabilization and Wireless Sensor Networks*
Abstract:

Speaker: **Shmuel Zaks** (Department of Computer Science, Technion, Haifa, Israel)
Title: *On design problems in ATM and optical networks*
Abstract:
Problems in design of ATM and optical networks will be described, with emphasis on open problems, and on problems related to the theme of self stabilization.

Speaker: **Chen Zhang** (University of Iowa)
Title: *Reliable-stabilizing PIF in tree networks*
Abstract:
We introduce the concept of reliable-stabilization, and give a reliable-stabilizing algorithm for propagation

of information with feedback problem(PIF) in tree networks. Our algorithm guarantees that from an arbitrary initial configuration, every propagation wave will reach all nodes, given that only single transient faults will occur after the first propagation wave starts.

Speaker: **Hongwei Zhang** (The Ohio State University)
Title: *Continuous Fault-containment and Local Stabilization in Path-vector Routing*
Abstract:
We present CBGP, a path-vector routing protocol that locally contains continuously occurring faults and locally stabilizes. Local containment enables CBGP to protect distant nodes from being affected by faults. Local stabilization enables CBGP to stabilize the network within time depending on the perturbation size after faults stop occurring. In CBGP, the distance to which the state of a node propagates is proportional to the time the state lasts. These properties are achieved by reacting to new fault only after first containing the response to the previous fault. Thus, the distance to which the corrected state of a node is propagated is proportional to the amount of time for which that state remains valid. In addition to analytically proving these properties, we evaluate CBGP by simulating Internet-type networks with up to 75 autonomous systems; we observe that CBGP reduces the number of fault-affected nodes by a factor of 71 and the network convergence time by a factor of 9.2 when compared with BGP.