Proceedings of

IWIN 2003

7th International Workshop on Interconnection Networks

June 16-17, 2003, Umea University, Sweden

The IWIN workshop series

IWIN 2003 is the 7th edition of the series of International Workshops on Interconnection Networks, holding every two years. After Marseilles (France) in 1991-95, Prague (Czech Rep.) in 1997, Schwerin (Germany) in 1999, and Vall de Nuria (Spain) in 2001, IWIN 2003 holds in Umea, a small city in the northern half of Sweden.

The goal of IWIN is to gather people interested by any structural and/or algorithmic aspect of Interconnection Networks. Topics of interest include, but are not limited to: information dissemination (broadcasting, gossiping, etc.), routing and scheduling, models for interconnection networks, dynamic graphs, distributed computing, mobile computing, network optimization, and structural properties of networks, with application to: fixed or wireless telecommunication systems, ad hoc or mobile networks, tightly coupled parallel computers, the World Wide Web, Internet, GRID computing, peer-to-peer systems, the Small World phenomenon, parallel architectures design, etc.

Presentations at IWIN 2003 were selected according to their relevance to the IWIN topics, and to their significance.

Program Committee:

Pierre Fraigniaud, CNRS, Univ. Paris-Sud, France
Jop Sibeyn, Halle University, Germany

Local Arrangements Chair:

Jop Sibeyn, Halle University, Germany

Conference Chair:

Pierre Fraigniaud, CNRS, Univ. Paris-Sud, France

Steering Committee:

Jean-Claude Bermond, INRIA Sophia-Antipolis.
Robert Cypher, Johns Hopkins Univ.
Miguel Fiol, UPC Barcelona.
Luisa Gargano, Univ. di Salerno.
Miro Kraetzl, Defense Science and Technology Organization.
Ludek Kucera, Charles Univ., Prague.
Roger Labahn, Univ. Rostock.
Tom Leighton, MIT.
Art Liestman, SFU, Vancouver.
Burkhard Monien, Univ. Paderborn.
Jarda Opatrny, Concordia Univ., Montreal.
David Peleg, Weizmann Institute.
Claudine Peyrat, Univ. of Nice.
Dominique Sotteau, CNRS, Univ. Paris-Sud (on leave)
Pavel Tvrdik, Czech Technical Univ., Prague
Eli Upfal, Brown Univ.
Imrich Vrto , Slovak Academy of Sciences, Bratislava

IWIN 2003 Program

	Monday, June 16	Tuesday, June 17
09:00		**Arie Koster** The Potential of Relaying in Cellular Wireless Computer Networks
09:30		**Tomas Johansson** On Clustering Algorithms in Ad Hoc Networks
10:00		**David Coudert** Traffic Grooming in Unidirectional WDM Ring Networks
10:30	**Opening**	Coffee break
11:00	**Johan Nykvist** On Disputes in Policy Based Path-Vector Routing	**Pino Persiano** Deterministic Truthful Approximation Schemes for Scheduling Related Machines
11:30	**Harald Gropp** Configurations as Interconnection Networks	**Alex Shvartsman** Atomic object services for mobile and dynamic networks
12:00	Lunch	Lunch
14:00	**Andrzej Pelc** Exploration and rendezvous problems in networks	
14:30	**Masafumi Yamashita** Random Walks on Finite Graphs using Local Topological Information	**Open Problems Session**
15:00	**Jop Sibeyn** Online and offline routing in meshes	
15:30	Coffee break	Coffee break
16:00	**Francesc Comellas** The Spectra of Deterministic Small-World Graphs	**Giuseppe Prencipe** The black hole search problem
16:30	**Jarda Opatrny** All-to-all, 2-hop Virtual Path Layout in Meshes and Chordal Rings	**Dariusz Kowalski** Performing Tasks in Asynchronous Environment
17:00		**End of Workshop**
20:00	**IWIN Workshop dinner**	

The Spectra of Deterministic Small-World Graphs

By: Francesc Comellas

Affiliation: Departament de Matemàtica Aplicada IV
Universitat Politècnica de Catalunya
Avda. Canal Olímpic s/n, 08860 Castelldefels, Barcelona, Catalonia, Spain

Contact: `http://www-mat.upc.es/~comellas/`
`email: comellas@mat.upc.es`

Abstract. Small-world graphs have a relatively low diameter, with respect to an structured graph with a similar number of vertices and edges, and a large clustering (vertices have many mutual neighbors) in relation to a random graph also with the same number of vertices and edges. Small-world graphs are of considerable interest because they model important real life networks as the WWW, transportation and communication networks, or even biological and social networks. Most previous research on small-world graphs has used simulations, probabilistic techniques, and random replacements of edges to study the main parameters and behaviour of these graphs. However, a deterministic approach has been considered in [F. Comellas, J. Ozón and J.G. Peters, Inf. Proc. Lett., 76, 83 (2000)] by starting with a circulant graph with n vertices, large diameter and high clustering and choosing a small number h of evenly-spaced vertices to be *hubs*, and interconnecting the hubs. With this approach we obtain deterministic small-world graphs (including regular graphs), that match former simulations. Small-world graphs occur when approximately 1% of the vertices are chosen to be hubs. Then the diameter is approximately 20% of the value in the original circulant graph, but 95% of the clustering remains.

On the other hand, the knowledge of the spectrum of a graph is important for the relation which the eigenvalues and their multiplicities have with relevant graph invariants and topological and communication properties such as diameter, bisection width, distances, connectivity, expansion, partitions, edge-loading distribution etc. Here we determine completely the spectra of several families of deterministic small-world graphs based on the technique described before.

Joint work with: Silvia Gago

Full version available at: http://www-mat.upc.es/~comellas/papers/CoGa03-IWIN.djvu

Traffic Grooming in Unidirectional WDM Ring Networks

By: David Coudert

Affiliation: Mascotte, cnrs-i3s-inria, Sophia Antipolis, France.

Contact: http://www-sop.inria.fr/mascotte/David.Coudert
email: David.Coudert@sophia.inria.fr

Abstract. We address the problem of traffic grooming in WDM rings with all-to-all uniform unitary traffic. The goal is to minimize the total number of SONET add-drop multiplexers (ADMs) required. We have shown that this problem corresponds to a partition of the edges of the complete graph into subgraphs, where each subgraph has at most C edges (where C is the grooming ratio) and where the total number of vertices has to be minimized. Using tools of graph and design theory, we optimally solve the problem for practical values and infinite congruence classes of values for a given C. Among others, we give optimal constructions when $C = 3$, $C = 4$, $C = 5$, $C \geq N(N-1)/6$ and results when $C = 12$.

Joint work with:

- Jean-Claude Bermond, Mascotte, cnrs-i3s-inria, France
  email: Jean-Claude.Bermond@sophia.inria.fr
- Xavier Muñoz, DMA4–UPC, Barcelona, Catalonia, Spain
  email: xml@mat.upc.es

Full version available at:

J-C. Bermond, D. Coudert, and X. Muñoz. Traffic grooming in unidirectional WDM ring networks: The all-to-all unitary case. In *The 7th IFIP Working Conference on Optical Network Design & Modelling – ONDM*, pages 1135–1153, 2003.

Configurations as Interconnection Networks

By: Harald Gropp

Affiliation: Universitt Heidelberg

Contact: d12@ix.urz.uni-heidelberg.de

Abstract. Configurations (v_r, b_k) are finite incidence structures of v points and b lines such that there are k points on each line and r lines through each point and such that through 2 different points there is at most one common line.

If $v = b$ and $r = k$ the configuration is called symmetric and denoted by v_k. Particular subclasses of configurations are projective planes, Steiner systems, partial geometries or generalized quadrangles. However, configurations were already defined in 1876, i.e. 125 years ago.

The Levi graph of a configuration is the bipartite point-line incidence graph with $v + b$ vertices and degrees k and r of girth at least 6. Configurations are just linear uniform regular hypergraphs.

In my talk the role of configurations in the business of constructing interconnection networks will be discussed.

Joint work with:

Full version available at:

On Clustering Algorithms in Ad Hoc Networks

By: Tomas Johansson

Affiliation: Luleå University of Technology

Contact: http://www.sm.luth.se/∼tomasjo
email:tomasjo@sm.luth.se

Abstract. A wireless ad hoc network consists of nodes that move freely and communicate with each other using wireless links. Ad-hoc networks do not use specialized routers for path discovery and traffic routing. One way to support efficient communication between nodes is to develop a wireless backbone architecture; this means that certain nodes must be selected to form the backbone.

Over time, the mobility of the individual nodes will cause links to activate and fail constantly. The structure of the backbone must change to reflect the changes in the network topology as nodes move around. The algorithm that selects the members of the backbone should be fast, and require as little communication between nodes as possible, since mobile nodes are often powered by batteries. One way to solve this problem is to group nodes into clusters, where one node in each cluster functions as clusterhead, responsible for routing.

Over time, a number of different clustering algorithms with different priorities have been proposed. Some seek to produce as few clusters as possible, while others focus on keeping the cluster size, the number of nodes in each cluster, between two predetermined values. One objective that all clustering algorithms share is to keep the amount of overhead communication as low as possible.

In our current work, we are developing a clustering algorithm that will allow us to define a limit for the cluster size as well as for the maximum number of hops between a node and its clusterhead. Instead of periodically constructing an entirely new clustering structure, the structure is only being constructed from scratch when the clustering algorithm is initiated. After that, the existing clustering structure is constantly modified in order to adapt to the changes in the network. The algorithm uses the assumption that nodes that have remained in the cluster for a long time are more likely to stay in the cluster in the future.

We are working on comparing our new algorithm with existing clustering algorithms in the ns-2 simulator. The simulations will be run for different topologies, sizes of network, and various limitations of the cluster structure properties. Future work includes analysis of routing algorithms for ad hoc networks and a design of a logical (cluster) structure that is optimal for the particular routing. Matching routing protocols and optimal topologies under different mobility scenarios for ad hoc networks would be the next step.

Joint work with: Lenka Carr-Motyčková

Full version available at:

The Potential of Relaying in Cellular Wireless Computer Networks

By: Arie M.C.A. Koster

Affiliation: Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Takustrasse 7, D-14195 Berlin, Germany.

Contact: `http://www.zib.de/koster/`
`email: koster@zib.de`

Abstract. Bandwidth is a scarce resource in wireless computer networks. *Relaying* is a technique to use existing technology in cellular data networks more efficiently and increase data throughput. Traditional cellular networks allow radio data transmission only between a mobile station and the cell's base station. Relaying enables mobile nodes to act as intermediaries, as *relaying nodes*, on the way to the base station. Since data can be sent faster over shorter distances, the total data throughput of a cell can be increased. Another potential benefit is the reduction of transmission power and hence, interference, improving the transmission situation.

Karl and Mengesha (2001) have shown that relaying increases data throughput significantly for special settings if data transmissions are cleverly scheduled in time. In this talk we discuss the potential of relaying in more general settings. We approach this question by defining a formal mathematical optimization problem and developing special mathematical models and algorithmic tools to determine transmission schedules that maximize the minimum data rate that any active mobile station achieves. In this way, we are able to quantitatively compare the best way of organizing data transmission with and without relaying for concrete networks consisting of multiple cells. We limit ourselves to the uplink case and a single frequency channel.

We apply our methodology to example networks and assess the average potential of relaying under the assumption of uniform placement of mobile terminals. This yields statistical evidence that relaying increases the *average* data throughput in an example network class by 30%. In addition, the method presented here can serve as an engineering benchmark for any concrete scheduling/ routing algorithm for relaying.

Joint work with: Hans-Florian Geerdes (ZIB)

Full version available at: Master thesis of Hans-Florian Geerdes available on `http://www.zib.de/groetschel/students/thesis_relaying.ps.zip`.

References:

H. Karl and S. Mengesha Analysing capacity improvements in wireless networks by relaying. In *Proc. of IEEE Intl. Conf. on Wireless LANs and Home Networks*, pages 339–348, Singapore, December 2001.

Performing Tasks in Asynchronous Environment

By: Dariusz Kowalski

Affiliation: University of Connecticut, Storrs, USA, and
Uniwersytet Warszawski, Poland

Contact: http://www.cse.uconn.edu/ kowalski
email: kowalski@cse.uconn.edu

Abstract. The problem of performing n tasks on p asynchronous or undependable processors is a basic problem in distributed computing. This problem, called *Do-All* (and also the variation of the *Do-All* problem, called *Write-All*), has been extensively studied in various synchronous models and in the asynchronous shared-memory system. However, there is a lack of efficient algorithms for both asynchronous message-passing and shared-memory systems. The efficiency of algorithms is measured in terms of *work* that accounts for all local steps performed by the processors in solving the problem.

In the message-passing model, where processors communicate by sending messages, *Do-All* can be trivially solved without any communication by an algorithm where each processor performs all tasks. This requires work $\Theta(p \cdot n)$. Thus it is important to develop *subquadratic* solutions (when p and n are comparable) by trading computation for communication. Following the observation that it is not possible to obtain subquadratic work when the message delay d is substantial, e.g., $d = \Theta(n)$, this work pursues a *message-delay-sensitive* approach. Here the bounds on work and communication are given as functions of p, n, and d – the upper bound on message delays, however algorithms have no knowledge of d and they cannot rely on the existence of an upper bound on d. We show algorithms which perform work $o(p \cdot n)$ as long as $d = o(n)$.

One of the variation of the *Do-All* problem is the *Write-All* problem—*using p processors write 1's into all locations of an array of size n.* It has been studied in various shared-memory models, e.g., in the asynchronous shared-memory system. In this problem writing 1 abstracts the notion of performing a simple task. Despite extensive research, there is a lack of efficient deterministic asynchronous algorithms for *Write-All*. An optimal algorithm would have work $\Theta(n)$, however it is known that optimality cannot be achieved when $p = \Omega(n)$. The quest then is to obtain work-optimal solutions for this problem using a non-trivial (compared to n) number p of processors. Recently it was shown that optimality can be achieved using a non-trivial number M of processors, where $M = \sqrt[4]{n/\log n}$. Our result shows that optimality can be achieved using close to M^2 processors; more precisely, using $M^{2-\varepsilon}$ processors, for any $\varepsilon > 0$.

Joint work with: Alex Shvartsman.

On Disputes in Policy Based Path-Vector Routing

By: Johan Nykvist

Affiliation: Luleå University of Technology

Contact: `http://www.sm.luth.se/~johan`
`email:johan@sm.luth.se`

Abstract. The standard inter-domain routing protocol used in the Internet today is the Border Gateway Protocol version 4 (BGP). BGP is a policy based path-vector routing protocol that distributes network reachability information between Autonomous Systems (AS:s). By default, BGP chooses the shortest path to a destination, in terms of the number of AS hops. However, the BGP protocol admits routing policies to control the path selection.

Varadhan et al have shown that certain configurations of routing policies in BGP routed networks can cause routing divergence, where no routing solution is found. The symptom of routing divergence is a never-ending routing information exchange, i.e., persistent route oscillation.

To remedy routing divergence due to policy disputes, Griffin and Wilfong have suggested an extension of the path-vector routing algorithm used in BGP, such that route histories are maintained in disseminated routes. Each router that mediates a route adds information to the route's history attribute.

If routing divergence arises due to policy disputes, the history will eventually contain cyclic patterns of entries, which can be utilized to detect persistent route oscillations by the routers. However, this merely detects oscillations — a suitable method to resolve such oscillations is yet to be found.

The presentation covers our ongoing work in which we study the nature of policy based routing and seek algorithmic solutions to problems related to policy based routing. Specifically, the problem of policy disputes is being addressed. The effect of the route history mechanism is under evaluation, with the objective to improve the mechanism. For this, we have created a simulator suited for experiments on policy disputes, which incorporates a simulation model of the route history mechanism. Using the simulator, we study policy disputes in different topologies of various sizes, using different policy configurations. We have also performed initial experiments in a testbench that consists of a real implementation of BGP. The results from the testbench corroborates the behavior of routing exchange at policy disputes in the simulator. In the next phase, we are going to address the problem of how policy disputes should be resolved, once they have been detected.

Joint work with: Lenka Carr-Motyčková

All-to-all, 2-hop Virtual Path Layout in Meshes and Chordal Rings

By: J. Opatrny

Affiliation: Dept. of Computer Science, Concordia University, Canada

Contact: `http://www.cs.concordia.ca/ faculty/opatrny/`
`email: opatrny@cs.concordia.ca`

Abstract. We show how projective planes can be used to obtain a solution of All-to-all, 2-hop Virtual Path Layout in Meshes and Chordal Rings. In many instances, our solutions are asymptotically optimal, and in specific cases it can lead to optimal solutions. Extensions of these results discussed.

Joint work with: S. Choplin, INRIA, Sophia Antipolis, and L. Narayanan, Concordia University.

Full version available at:

Exploration and rendezvous problems in networks

By: Andrzej Pelc

Affiliation: Université du Québec en Outaouais

Contact: `http://w3.uqah.uquebec.ca/pelc/main.html`
email: `pelc@uqo.ca`

Abstract. One or more robots (mobile agents) have to explore a connected unknown or partially known graph. Every edge must be traversed by at least one robot and exploration has to be completed as fast as possible. We compare the performance of deterministic exploration algorithms that do not know the graph (or know it only partially) to that of an optimal algorithm having full knowledge of the graph. We seek exploration algorithms for which this worst-case ratio is the smallest. In the case of one robot, we consider scenarios where the robot has no knowledge of the graph and when it has an unoriented map (isomorphic copy) of it. In the case of many robots we restrict attention to trees and assume that they are totally unknown to exploring agents. For this scenario we investigate distributed collective exploration and show the influence of communication between agents on efficiency of exploration.

A problem related to graph exploration is rendezvous in graphs. Two mobile agents having distinct identifiers and located in nodes of an unknown anonymous connected graph, have to meet at some node of the graph. We seek fast deterministic algorithms for this rendezvous problem, under two scenarios: simultaneous startup, when both agents start executing the algorithm at the same time, and arbitrary startup, when starting times of the agents are arbitrarily decided by the adversary.

Joint work with: A. Dessmark, P. Fraigniaud, L. Gasieniec, D. Kowalski

Full version available at:

1. A. Dessmark, A. Pelc, Optimal graph exploration without good maps, Proc. 10th European Symposium on Algorithms (ESA 2002), September 2002, Rome, Italy, LNCS 2461, 374-386.

2. A. Dessmark, P. Fraigniaud, A. Pelc, Deterministic rendezvous in graphs, available upon request.

3. P. Fraigniaud, L. Gasieniec, D. Kowalski, A. Pelc, Collective tree exploration, available upon request.

Deterministic Truthful Approximation Schemes for Scheduling Related Machines

By: Pino Persiano

Affiliation: Dipartimento di Informatica ed Appl. Università di Salerno – Italy

Contact: `http://www.dia.unisa.it/~giuper`
`email: giuper@dia.unisa.it`

Abstract. We consider the problem of scheduling on parallel related machines owned by selfish agents. We provide a deterministic polynomial-time approximation scheme and suitable payment functions that yield a $(1+\epsilon)$-truthful mechanism for several NP-hard restrictions of this problem. Up to our knowledge, our result is the first deterministic truthful mechanism for an NP-hard problem that guarantees $(1+\epsilon)$-approximate solutions, for any $\epsilon > 0$.

Our result also yields a family of deterministic truthful $(2+\epsilon)$-approximation mechanisms for any fixed number of machines. This improves over the randomized 3-approximation mechanism by Archer and Tardos [2001] which is truthful under a weaker notion of truthfulness.

To obtain our results we introduce a technique to transform the PTAS by Graham into a deterministic truthful mechanism achieving the same approximation. As a byproduct we obtain a sufficient condition for the existence of a truthful $(1+\epsilon)$-approximation mechanism for any fixed number of machines.

Joint work with: E. Auletta, R. De Prisco, P. Penna

Full version available at: htpp://ec.tcfs.it/Papers/related-MDfull.ps.gz

The Black Hole Search Problem

Facing the harmful host threats in distributed mobile computing environments

By: Giuseppe Prencipe

Affiliation: Università di Pisa

Contact: `http://sbrinz.di.unipi.it/~peppe`
`email: prencipe@di.unipi.it`

Abstract. The use of mobile agents is becoming increasingly popular when computing in networked environments, ranging from Internet to the Data Grid. Computing in such environments is termed *distributed mobile computing*[1] (DMC). In DMC terminology, a network site is a *host*; local processes are *stationary agents*; mobile agents *navigate asynchronously*, moving from host to neighboring host, and perform computations at each host, according to a predefined set of behavioral rules called *protocol*, the same for all agents. The hosts provide a storage area called *whiteboard* for incoming agents to communicate and compute, whose access is held in fair mutual exclusion.

The major practical concern in these systems is definitely *security*. Among the severe security threats faced in distributed mobile computing environments, two are particularly troublesome: *harmful agent* (that is, the presence of malicious mobile processes), and *harmful host* (that is, the presence at a network site of harmful stationary processes).

We focus on the presence in the network of a harmful host, called *black hole*: a host which *disposes* of visiting agents upon their arrival, leaving *no observable trace* of such a destruction. The task is to unambiguously determine and report the location of the black hole. This problem has been intensively studied from a programming point of view, and recently also from an algorithmic perspective. The research concern is to determine under what conditions and at what cost mobile agents can successfully accomplish this task. In this presentation, we will describe bounds and algorithmic techniques that allow the agents to locate the black hole for most of the frequently used network topologies: *rings*, *hypercubes*, *cube-related networks*, *chordal rings*, and multidimensional *tori* and *meshes*.

Joint work with: Stephan Dobrev, Paola Flocchini, R. Kralovic, P. Ruzieka and Nicola Santoro.

Full version available at: http://sbrinz.di.unipi.it/~peppe/Publications.html

[1] The term "mobile computing" alone is traditionally used in reference to wireless (e.g., "ad-hoc") networks.

Atomic object services for mobile and dynamic networks

By: Alex A. Shvartsman

Affiliation: University of Connecticut

Contact: `http://www.engr.uconn.edu/~aas`
email: `alex@theory.lcs.mit.edu`

Abstract. Implementing atomic read/write objects in dynamic and mobile network settings is challenging due to the perturbations in the communication medium and due to the dynamic changes in the sets of processors "owning" the objects. Algorithms that assume a static universe of processors or that rely on conventional networks are not adequate for dealing with such environments. Given that replication is the only means for ensuring availability and fault-tolerance, and that atomic consistency must always be maintained, new approaches are needed. These approaches must also deal with the complexities associated with implementing atomic object services directly on top of mobile processors in *ad hoc* networks. The two approaches explored in this work are: *(i)* implementing dynamic reconfiguration of the set of object owners within an unknown universe of processors, and *(ii)* introducing an abstraction that associates virtual processes with geographic regions in an *ad hoc* setting. The new algorithms use reconfigurable quorum systems to ensure atomic consistency in all executions. For *ad hoc* networks the algorithms are structured in two layers, where the top layer uses virtual processes to manage atomicity and the bottom layer uses physical processors to implement the virtual processes. The progress of read and write operations is guaranteed provided the relevant quorum systems are not disabled, and in such cases the operations can terminate even if reconfigurations do not complete. When reconfigurations are infrequent, read and write operations complete in one or two round-trip communication phases.

Joint work with: Shlomi Dolev, Ben-Gurion University
Seth Gilbert, MIT
Nancy Lynch, MIT
Piotr Musial, University of Connecticut
Jennifer Welch, Texas A&M University

Full version available at: Preliminary results dealing with a part of the work abstracted above have been reported by Lynch and Shvartsman under the title "RAMBO: A Reconfigurable Atomic Memory Service for Dynamic Networks" in the *Proceedings of the 16th International Conference on Distributed Computing*, Lecture Notes in Computer Science, Volume 2508, 2002, pages 173–190.

`http://link.springer.de/link/service/series/0558/bibs/2508/25080173.htm`

Online and Offline Routing on Meshes

By: Jop F. Sibeyn

Affiliation: Institut für Informatik
Martin-Luther-Universität Halle-Wittenberg
06099 Halle, Germany

Contact: `http://www.informatik.uni-halle.de/≃jopsi/`
email: `jopsi@informatik.uni-halle.de`

Abstract. We consider permutation routing problems on 2D mesh-connected computers with side length n. The performance metrics are both the time T and queue size Q. The goal is to come as close to $T = 2 \cdot n - 2$ and $Q = 1$ as possible. First we prove a lower bound of $T = 3 \cdot n - 4 \cdot \sqrt{n}$ for any algorithm which achieves $Q = 1$ in a certain natural class of routing algorithms. This shows that there cannot be really simple algorithms achieving optimality. A new elegant offline algorithm achieves $T = 2 \cdot n - 1$ and $Q = 2$. This algorithm can be turned into an online algorithm with $T = 2 \cdot n + O(1)$ and with queue size $Q = 3$. A refinement of the offline algorithm leads to a randomized online algorithm running in $T = 2 \cdot n + O(\log n)$ and $Q = 2$. This algorithm has so much slack (most queues contain at most one packet), that it can be turned into a deterministic online algorithm with $T = 2 \cdot n + o(n)$ and $Q = 2$. Most ideas and algorithms can be generalized for routing on 3D meshes.

Joint work with: Bogdan S. Chlebus, University of Colorado at Denver, USA.

More complete version available at: CDROM of *17th International Parallel and Distributed Processing Symposium*, IEEE, 2003.

Random Walks on Finite Graphs using Local Topological Information

By: M. Yamashita

Affiliation: Dept. Computer Science and Communication Engineering, Kyushu University

Contact: `http://tcslab.csce.kyushu-u.ac.jp/ mak/index.html`
`email: mak@csce.kyushu-u.ac.jp`

Abstract. It is just amazing that the mean hitting time and the cover time of a random walk on a finite graph, in which the vertex visited next is selected from the adjacent vertices at random with the same probability, is bounded by $O(n^3)$ for any graph with order n, despite of the lack of global topological information. Thus a natural guess is that a better transition matrix is designable if more topological information is available. We can indeed reduce the mean hitting time and the cover time to $O(n^2)$ and $O(n^2 \log n)$, respectively, by using the degree information of adjacent vertices. Note that the mean hitting time of a path graph for any transition matrix is $\Omega(n^2)$, which means that the degree information of adjacent vertices is sufficient with respect to the mean hitting time.

Joint work with: S. Ikeda and I. Kubo

Full version available at: S. Ikeda, I. Kubo, N. Okumoto, and M. Yamashita: "Impact of Local Topological Information on Random Walks on Finite Graphs," *ICALP 2003* (to appear).

S. Ikeda, I. Kubo, and M. Yamashita: "Reducing the Hitting and the Cover Times of Random Walks on Finite Graphs by Local Topological Information," *The 2003 International Conference on VLSI* (VLSI'03) June 23-26, 2003, Las Vegas, Nevada, USA (to appear).