Chicago Journal of Theoretical Computer Science

The MIT Press

Volume 1996, Abstracts
31 December 1996

ISSN 1073–0486. MIT Press Journals, 55 Hayward St., Cambridge, MA 02142; (617)253-2889; *journals-orders@mit.edu, journals-info@mit.edu*. Published one article at a time in LaTeX source form on the Internet. Pagination varies from copy to copy. For more information and other articles see:

- *http://www-mitpress.mit.edu/jrnls-catalog/chicago.html/*
- *http://www.cs.uchicago.edu/publications/cjtcs/*
- *ftp://mitpress.mit.edu/pub/CJTCS*
- *ftp://cs.uchicago.edu/pub/publications/cjtcs*

The *Chicago Journal of Theoretical Computer Science* is abstracted or indexed in *Research Alert,®* *SciSearch,®* *Current Contents® /Engineering Computing & Technology*, and *CompuMath Citation Index.®*

The *Chicago Journal of Theoretical Computer Science* is a peer-reviewed scholarly journal in theoretical computer science. The journal is committed to providing a forum for significant results on theoretical aspects of all topics in Computer Science.

Abstracts from Volume 1996

31 December, 1996

1. Rank Predicates vs. Progress Measures in Concurrent-Program Verification

Moshe Y. Vardi

9 February 1996

This note describes a direct relationship between rank predicates and progress measures in concurrent-program verification.

2. Sparse Hard Sets for P Yield Space-Efficient Algorithms

Mitsunori Ogihara

27 March 1996

In 1978, Hartmanis conjectured that there exist no sparse complete sets for P under logspace many-one reductions. In this paper, in support of the conjecture, it is shown that if P has sparse hard sets under logspace many-one reductions, then $\mathrm{P} \subseteq \mathrm{DSPACE}[\log^2 n]$.

3. Optimal Virtual Path Layout in ATM Networks with Shared Routing Table Switches

Ornan Gerstel, Israel Cidon, Shmuel Zaks

31 October 1996

In this paper we present a new model for routing that occurs in high-speed ATM networks. Within this model we define a general routing problem, called a virtual path layout. Given a network of nodes (switches) and links, one must find a set of paths in the network, termed the virtual path layout, whereby each pair of nodes may connect using a route that is a concatenation of a small number of virtual paths and is also short in terms of the network links it traverses. Each such layout implies a utilization of the routing tables in the network's nodes. Our goal is to find a layout that minimizes this utilization, assuming that each such node has a central routing table. We prove that this problem is NP-complete, and consequently focus on a simpler problem, in which it is required to connect all nodes to a single switch. Next, we prove that even this problem is NP-complete, and restrict some of the assumptions to yield a practical subproblem, for which we present a polynomial-time greedy algorithm that produces an optimal solution. Finally, we use this restricted problem as a building block in finding a suboptimal solution to the original problem. The results exhibit a tradeoff between the performance of a routing scheme and its resource utilization.

4. Weakly Growing Context-Sensitive Grammars

Gerhard Buntrock, Gundula Niemann

13 November 1996

This paper introduces *weakly growing context-sensitive grammars*. Such grammars generalize the class of growing context-sensitive grammars (studied by several authors), in that these grammars have rules that "grow" according to a position valuation.

If a position valuation coincides with the initial part of an exponential function, it is called a *steady* position valuation. All others are called *unsteady*. The complexity of the language generated by a grammar depends crucially on whether the position valuation is steady or not. More precisely, for every unsteady position valuation, the class of languages generated by *WGCSG*s with this valuation coincides with the class *CSL* of context-sensitive languages. On the other hand, for every steady position valuation, the class of languages generated corresponds to a level of the hierarchy of exponential time-bounded languages in *CSL*. We show that the following three conditions are equivalent:

- The hierarchy of exponential time-bounded languages in *CSL* collapses.
- There exists a class defined by an unsteady position valuation such that there is also a normal form of order 2 (e.g., Cremers or Kuroda normal form) for that class.

- There exists a class defined by a steady position valuation that is closed under inverse homomorphisms.

5. Uniform Self-Stabilizing Orientation of Unicyclic Networks under Read/Write Atomicity

H. James Hoover

5 December 1996

A distributed system is *self-stabilizing* if its behavior is correct regardless of its initial state. A *ring* is a distributed system in which all processors are connected in a cycle and each processor communicates directly with only its two neighbors. A ring is *oriented* when all processors have a consistent notion of their left and right neighbors. A ring is *uniform* when all processors run the same program and have no distinguishing attributes, such as processor IDs. A well-known self-stabilizing uniform protocol for ring orientation is that of Greenberg, Jájá, and Krishnamurthy. For a ring of size n, this protocol will stabilize in expected $O(n^2)$ processor activations. This assumes that processors are scheduled by a *distributed demon*—one in which the communication registers between processors can be atomically updated (a read followed by a write), and the processors have the ability to make random choices.

This paper generalizes the notion of orienting a ring to that of orienting a *unicyclic* network, that is, a ring with trees attached. We present a very simple protocol for the self-stabilizing orientation of such unicyclic networks. It has the same expected $O(n^2)$ processor activation performance as the Israeli and Jalfon protocol for rings. But ours works under the more general scheduling assumption of a *read/write demon*—one in which a read or write of a communication register is atomic, but an update (a read followed by a write) is not. We similarly assume the ability to make random choices. A subresult of our protocol is a uniform self-stabilizing algorithm for the two processor token-passing problem under the read/write demon.

Our protocol is uniform in the sense that all processors of the same degree are identical. In addition, observations of the behavior of the protocol on an edge yield no information about the degree of the processors at the ends of the edge.

6. Manhattan Channel Routing is NP-Complete under Truly Restricted Settings

Martin Middendorf

30 December 1996

Settling an open problem that is over ten years old, we show that Manhattan channel routing—with doglegs allowed—is NP-complete when all nets have two terminals. This result fills the gap left by Szymanski, who showed the NP-completeness for nets with four terminals. Answering a question posed by Schmalenbach and Greenberg, Jájá, and Krishnamurty, we prove that the problem remains NP-complete if in addition the nets are single-sided and the density of the bottom nets is at most one. Moreover, we show that Manhattan channel routing is NP-complete if the bottom boundary is irregular and there are only 2-terminal top nets. All of our results also hold for the restricted Manhattan model where doglegs are not allowed.