

Capsule Reviews

The Capsule Reviews are intended to provide a short succinct review of each paper in the issue, in order to bring the content to a wider readership. This issue's Capsule Reviews were compiled by Fairouz Kamarreddine. Professor Kamarreddine is an Associate Editor of *The Computer Journal* and is based in the Department of Computing and Electrical Engineering at Heriot-Watt University, Edinburgh, UK.

Support for reasoning about interactive systems through human-computer interaction designers' representations. F. PATERNÒ AND C. SANTORO

Usability is a fundamental aspect of interactive systems and human-computer interaction (HCI) techniques are thriving. However, although formal verification techniques have brought many advantages to fields such as hardware verification, these techniques are more difficult to apply to the more demanding field of HCI. Nonetheless, different formal approaches such as model checkers and theorem provers have been used to verify properties of interactive systems and to improve reliability of the resulting software. However, given the complexity of interactive systems, adopting a unified formal specification will result in a model which is too complex for verification and also which is too difficult for novices of formal techniques to use or understand. Armed with this fact, this paper proposes the use of a number of partial property-oriented specifications (phases), instead of the fully formalized specification. These partial specifications enable the selection of the most appropriate technique for each specific case. This approach additionally enables a good integration of formal model checking with tools for task modelling. The main result is that designers can take advantage of model checking tools, while still working with representations of the relevant models familiar to them, even allowing people with little background in formal methods to use it.

A hierarchical technique for constructing efficient declustering schemes for range queries. R. BHATIA, R. K. SINHA AND C.-M. CHEN

Consider multi-dimensional examples such as remote-sensing data collected by earth-orbiting satellites. Data for such examples can be organized as k -dimensional grids where each grid tile represents a data block that is placed on one of the disks. This multi-disk usage enables parallel data retrieval and avoids idle CPU cycles. This paper deals with the problem of allocating uniform, multi-dimensional data among parallel disks to minimize disk access time for range queries, by devising a good declustering scheme which assigns each tile in the data set to one of the disks so that the data blocks requested by a query are distributed as evenly as possible among the disks. Designing good declustering schemes for large numbers of disks is vital in many scientific applications where thousands of disks may be used to deal with the ever-increasing data size. As the problem studied in this paper is combinatorial, it is not

computationally feasible to perform an exhaustive search for the best scheme for large numbers of disks. Thus, the main proposal of the paper consists of a hierarchical technique to efficiently construct declustering schemes for large numbers of disks based on declustering techniques for small numbers of disks, which are called the base schemes. The paper shows that the performance of the resulting declustering scheme is proportional to the performance of the base schemes. Simulation results are also given to both support the main claim of the paper regarding the proportionality of performance of the hierarchical declustering scheme with respect to that of the base schemes, and to compare the performance of the hierarchical declustering method with already existing declustering schemes.

Placement of Web-server proxies with consideration of read and update operations on the Internet. X. JIA, D. LI, X. HU, W. WU AND D. DU

The use of the Web is ever-increasing and this amount of Web traffic results in a slow down of access. In order to ease traffic congestion, caching is usually adopted. Caching can either be client-based or server-based and can hence be done at either client or server proxies. Server proxies aim to improve client response times, distributing the workload of the Web server and reducing the network traffic. A server proxy stores partially replicated data of a Web site so that when a client requests access to that Web site, the server proxy provides it. This brings about the costs of keeping the proxy's replicated data up to date with its replica data. This paper deals with this problem and investigates the optimal placement of server proxies of a Web server on the Internet with the consideration of both read and update operations of the data on the Web server. In particular, two sub-problems of placement of server proxies on the Internet are studied:

1. The problem of optimal placement of k server proxies in a system to minimize the total access cost to the Web server.
2. The problem of finding for an unconstrained number of server proxies, the optimal number of server proxies required so that the overall cost of access is minimized.

The paper proposes optimal solutions to these problems by formulating them in a dynamic programming environment. The paper uses various simulations to evaluate the performance of the proposed solutions and to assess the factors that affect the performance of these solutions.

Coarse-grained parallel transitive closure algorithm: path decomposition technique.

A. GIBBONS, A. PAGOURTZIS, I. POTAPOV AND W. RYTTER

The generic transitive closure problem (TC) can be defined using the semiring notion with generalized sum and product. The value of each path from a vertex i to a vertex j is the product of the values of its edges and the TC problem involves calculating the sum of the values of the simple paths from a given vertex i to a given vertex j (simple in the sense that no vertex appears twice except possibly the first and last). The TC problem has many useful applications in graph theory and in dynamic programming. Matrix multiplication which is well suited for parallel implementation plays a helpful role in achieving efficient solutions to the TC problem. This paper studies the TC problem restricted to the so-called simple semirings and shows that in order to solve this restricted TC problem, it is sufficient to add the values of all simple paths at least once. This result is important because adding non-simple paths is not avoided in any of the known algorithms for TC. In its study of the TC problem, this paper investigates the relation between fine-grained and coarse-grained distributed computations of a class of problems related to it. It implements two different coarse-grained algorithms (a coarse-grained version of the known Warshall algorithm, and the new 2-pass algorithm designed for distributed systems with slow communication) and shows that the latter performs better.

An algorithm for validating ASN.1 (X.680) specifications using set constraints.

C. RINDERKNECHT

This paper rightly disagrees with the common belief among practitioners in the networking field that theoretical studies cannot help the industrial audience, and states that many tools used by engineers can benefit from theoretical studies. The paper presents the first complete and formal semantics of X.680 (the main part of ASN.1) using set constraints. ASN.1 (Abstract Syntax Notation One) is a standard language for defining data types whose values may be exchanged across a network between two communicating applications, allowing the protocol designer to capture numerous networking concepts without worrying about the possible heterogeneity of the peers. The peers share a set of ASN.1 modules and agree upon encoding rules which encode values into series of bits. An ASN.1 compiler accepts a set of ASN.1 modules and produces (according to the encoding rules and a peer-specific target programming language), a set of data type definitions in that language together with codes for the values to be exchanged. These pieces of source code are compiled and linked separately against the communicating application. However, due to the expressiveness of ASN.1 (in order to fulfil different user needs), it remains a challenge to write an ASN.1 compiler. Hence, some compilers may reject valid specifications or may accept invalid ones. This paper fully validates the specifications written in the main part of ASN.1 using an algorithm based on set constraint theory. This algorithm is two-fold (a collecting algorithm that

extracts some constraints and a solving procedure) and can be implemented and integrated in the front-end analyser of an ASN.1 compiler.

Determining the asynchronous threshold values of the timed token medium access protocol.

A. GENÇATA AND F. BUZLUCA

Synchronous messages are periodic: they arrive at regular intervals and have delivery constraints. Asynchronous messages are non-periodic: they may arrive in a random way and have no time constraints. The timed-token medium access protocol provides support for both synchronous and asynchronous services in the following sense: each node in the network is guaranteed a certain bandwidth for its synchronous traffic and the remaining bandwidth is dynamically shared by all nodes for asynchronous traffic. If the protocol parameters are chosen carefully, synchronous services can be used to guarantee the deadlines of real-time messages and the asynchronous service can be used to transmit non-real-time messages. Many studies of protocol parameters have taken place. In this paper two methods are developed to calculate suitable asynchronous threshold values for different purposes. Both methods can easily be applied to existing fibre distributed data interface networks to improve their functionality.

Solving the converter placement problem in WDM ring networks using genetic algorithms.

T.-M. CHAN, S. KWONG AND K. F. MAN

This article starts from the observation that because of the tremendous growth of network traffic and services and because of the large number of users, optical networks using the technology of wavelength division multiplexing (WDM) have great potential for becoming the next generation of large bandwidth networks. In a WDM network, both static and dynamic traffic can be considered. In each of these traffic patterns, one needs to select the lightpaths and assign them wavelengths so that no two paths using the same wavelength pass through the same link to prevent collision. This problem is known as the routing and wavelength assignment (RWA) problem. This paper studies the RWA problem in ring networks that have the advantages of providing simple routing policies and simple control and management of network resources. As there are two strategies to transmit optical signals in WDM networks (one with wavelength conversion and the other without), the paper chooses the strategy that uses wavelength conversion defending its choice by the fact that this increases efficiency. However, wavelength conversion is neither economical nor practical and so the paper considers sparse wavelength conversion which is a small fraction of routing nodes equipped with full wavelength converters to strike the balance between the total converter cost and performance improvement. The paper studies the converter placement problem in a WDM ring network and proposes three heuristic and search techniques to solve this combinatorial optimization problem.

Design of reconfigurable VLSI architecture for hybrid arithmetic in $GF(2^m)$. C. N. ZHANG AND H. LI

The finite fields $GF(2^m)$ of characteristic 2 have important uses. However, although addition is a fast operation in these fields, multiplication is costly and slow and hence various classifications of finite field architectures have been given. On the one hand, multipliers of these fields can be classified into fast bit-parallel architectures and low-complexity bit-serial multipliers. On the other hand, finite fields architecture can be classified with respect to the representation of the field elements. These basis representations can be canonical, dual or normal. Each of

these representations has its advantages and disadvantages. This paper defines a redundant canonical basis with the reducible all one polynomial (AOP) in order to reduce the complexity of the arithmetic operations in $GF(2^m)$. There are many advantages of such a proposal which are highlighted in the paper. The authors give the fundamental theorems of the arithmetic operations in the redundant basis and propose three new multipliers comparing them to previous related work. A new reconfigurable hybrid arithmetic VLSI architecture is then derived which satisfies nice properties. Simulation results bring out the advantages of the framework.