# 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.**

**A power-saving network architecture for accessing the Internet from mobile computers: design, implementation and measurement.** G. ANASTASI, M. CONTI AND W. LAPENNA

A big limitation of a portable computer is the short life of its battery power. There does not seem to be a quick solution to the battery technology. For this reason, ways need to be found to utilize power efficiently. This paper proposes a power-saving strategy for accessing the Internet from a mobile computer. Experimental results show that power consumption related to networking rises to over 50% of the overall power consumption in a laptop computer. The paper defines an index for measuring the impact of mobile–host data transfers on the consumption of battery power and to estimate the power consumption of the TCP/IP protocol stack (which is the *de facto* standard architecture at the inter-network level). The paper identifies idle states of the network interface and defines a power-saving network architecture (PSNA) for accessing the Internet from a mobile host in a manner that eliminates the consumption of the battery power caused by the idle states. Experimental results show that PSNA helps to increase the battery lifetime of the processing capacity of a mobile host. The paper provides an implementation of this architecture and evaluates its performance.

**Byzantine fault detectors for solving consensus.** K. P. KIHLSTROM, L. E. MOSER AND P. M. MELLIAR-SMITH

This paper investigates the use of fault detectors for solving the consensus problem in asynchronous distributed systems that are subject to Byzantine faults. The consensus problem is fundamental in distributed computing. It had previously been established, on one hand, that it is impossible to achieve consensus in an asynchronous distributed system that is subject to even one crash fault, and on the other hand, that consensus is possible in an asynchronous distributed system that is subject to crash faults if the asynchronous model is augmented with an unreliable fault detector. This paper concentrates on this latter result and considers unreliable fault detectors for solving consensus in an asynchronous distributed system that is subject to Byzantine faults, which embody malicious designs and attempt to bring the system down. The paper provides a more accurate characterization of Byzantine faults. Completeness and accuracy are two critical properties of unreliable fault detectors. The paper subsequently provides eventual weak and strong completeness properties (used to solve the consensus problem) which it uses to introduce four new classes of unreliable Byzantine fault detectors. Two algorithms are then given: one to solve the consensus problem in a particular asynchronous distributed system where the number of Byzantine faults is always bounded in relation to the number of processes, the other to implement a Byzantine fault detector in a model of partial synchrony. This paper attempts to capture as much information about Byzantine faults as possible allowing faulty processes to be detected and excluded from the system.

**Viable architectures for high-performance computing.** S. G. ZIAVRAS, Q. WANG AND P. PAPATHANASIOU

The demands for good performance of computational systems led to the development of computers with thousands of processors. But high performance is difficult to achieve because of the difficulty in developing low-complexity bandwidth and low-latency networks to interconnect the thousands of processors. Interprocessor connection networks have poor topological properties which lead to large memory latencies, and scalable networks with good topological properties are hard to build because of their high wiring complexities. The hypercube is such a network with good topological properties and has dominated the high-performance computing field, but its high wiring complexity is its greatest disadvantage. The versatility of the hypercube in emulating important topological properties efficiently is an incentive for introducing hypercube-like interconnection networks of lower complexity and better scalability, without sacrificing the nice topological properties. Hence, the paper proposes the highly overlapping windows (HOWs) networks which have nice topological properties (extremely small diameter, average internode distance, and immense bisection width), yet also enjoy low complexity and good scalability. HOWs can be obtained from hypercubes by removing some of their processor interconnections in order to reduce their wiring complexity. The architecture of HOWs is presented together with a fault-tolerant routing algorithm to find shortest paths. The cost analysis of HOWs is discussed as is the embedding of some interconnection networks into two-dimensional HOWs systems.

**Fair exchange.** H. PAGNIA, H. VOGT AND F. C. GÄRTNER

Nowadays, the network is being used for business transactions and sales of goods and services (especially digital services). A basic property of digital services

is that they cannot be revoked. For example, once the service is granted, the service provider can no longer force the recipient to return it without keeping a private copy. This makes fair exchange a central problem where it is important to exchange two electronic items in a fair manner. Several models have been adopted for this purpose. However, the fair exchange protocols which have been presented in the literature are diverse and largely incomparable. This paper reviews the research area of fair exchange and several formalizations of fairness and of the fair exchange protocols. It shows that these protocols are very diverse and incomparable but that they can be understood as a composition of different protocol modules. The paper then gives a general framework for modelling fair exchange protocols. This generalizing framework accommodates several different fair exchange protocols where customers and vendors can select the protocol that suits a particular need. Furthermore, as the approach of the paper is compositional, one can dynamically select exchange protocols for a given level of fairness.

**Inverting Dirichlet tessellations.** F. P. SCHOENBERG, T. FERGUSON AND C. LI

If $P_1, \ldots, P_n$ are $n$ points of the plane, then the Dirichlet tessellation corresponding to these points is constructed by dividing up (a portion of) the plane into $n$ distinct cells such that for each point $P_i$, its cell consists of all locations closer to $P_i$ than to any $P_j$ for $j$ different from $i$. Dirichlet tessellations date back to Descartes and have been used in many applications and disciplines. In 1978, an efficient algorithm for generating the boundaries of the cells that constitute the Dirichlet tessellations was established. The present paper attempts to construct the inverse algorithm. That is, given the boundaries of the cells in a Dirichlet tessellation, how can one obtain the points $P_1, \ldots, P_n$ corresponding to these cells? Different algorithms for investigating Dirichlet tessellations are given and their performance properties are discussed. One such algorithm is shown to be very efficient and to only require $O(n)$ computations, giving a high level of accuracy.

**Transient analysis of rewarded continuous time Markov models by regenerative randomization with Laplace transform inversion.** J. A. CARRASCO

For $X$ a continuous time Markov chain (CTMC) model with a finite state space Omega, a rewarded CTMC is obtained by imposing over $X$ a reward rate structure $r_i$, for $i$ in Omega, where $r_i$ gives the rate at which a reward is earned while $X$ is in state $i$. The behaviour of the reward rate can be examined using different measures of which the paper considers two. Computations of the measures require the analysis of the CTMC $X$ and this can be costly and has led to various methods including ordinary differential equation solvers and randomization. However, randomization is expensive when the model is stiff and this led to the development of several variants including the regenerative randomization. However, the truncated transformed model obtained in regenerative randomization is as stiff as the original and its solution may be costly. For this reason, the paper develops a variant: the regenerative randomization with Laplace transform inversion where the truncated transform model is solved using a Laplace transform inversion algorithm. The performance of the proposed variant is compared with those of existing ones and its accuracy is analysed. The paper concludes that for some models, the new variant is significantly faster than regenerative and standard randomizations.

**Construction of symmetrical reversible variable length codes using backtracking.** H.-W. TSENG AND C.-C. CHANG

Variable length codes (VLCs) are usually used in the entropy coding strategies of many image coding standards like JPEG and MPEG-1/-2. Due to the drawbacks of VLCs especially with respect to error propagation, reversible variable length codes (RVLCs) have been adopted. RVLCs avoid continuous errors, can be decoded in both forward and backward directions, and speed up the search of encoded data. RVLCs can be either symmetrical (which share the same code tables in the forward and backward directions) or asymmetrical. Although symmetrical RVLCs are more elegant, asymmetrical ones are more efficient. This paper provides an algorithm for constructing symmetrical RVLCs where the decoding is efficient and the average code length is small. As the title suggests, backtracking techniques are used in this algorithm. The experimental results provided in the paper illustrate that, indeed, the algorithm generates efficient symmetrical RVLCs.

**Loopless generation of Schröder trees.** J. F. KORSCH AND P. S. LAFOLLETTE

In the early seventies, Ehrlich established the interesting result that a number of combinatorial objects can be looplessly generated in the sense that the time taken to generate the next object is at most a constant independent of the size of the object. This result sparked a number of loopless algorithms for generating combinatorial objects. But generating algorithms for combinatorial objects depends on the representation used for the objects. Loopless algorithms for one representation may not be loopless algorithms for another representation. This paper concentrates on particular combinatorial objects known as Schröder trees. Schröder trees and Schröder numbers $s(m)$ (which are the numbers of Schröder trees with $m$ leaf nodes) have a long history dating back over 2200 years. Recently, a one-to-one correspondence has been established between Schröder trees with $m$ leaf nodes and well-weighted binary trees with $m$ leaf nodes. Also recently, a loopless generation algorithm for (a representation of) well-weighted binary trees with $m$ leaf nodes has been given. Via the one-to-one correspondence, this gives an indirect loopless generation algorithm for Schröder trees with $m$ leaf nodes. This paper provides the first loopless algorithms for directly generating Schröder tree representations.