Workshop 13+14

Architectures and Networks

Kieran Herley and David Snelling

Co-chairmen

This workshop embraces two broad themes: routing and communication networks and parallel computer architecture. The first theme is devoted to communication in parallel computers and covers all aspects of the theory and practice of the design, analysis and evaluation of interconnection networks and their routing algorithms. The continued vitality of this area is reflected by the strength and quality of the papers presented in Sessions 1 and 3 (Routing Networks and Algorithms I and II) of this workshop. Although recent years have seen computer architecture undergoing a period of contraction and a tendency toward uniformity in the form of shared memory systems composed of super scalar processors, there is nonetheless no lack of innovation in the field. The computer architecture theme, represented by the papers of Sessions 2 and 4 (Computer Architecture I and II), highlights not only innovation within the converging area of super scalar systems, but the continued presence of novel ideas as well.

A total of 32 submissions were received (19 on the theme of routing and communication networks and 13 on the computer architecture theme) of which a total of 15 were accepted by the workshop committee for presentation at the conference.

Session 1 is devoted to the issue of routing algorithms for unstructured or loosely structured communication on various networks. The paper by S. R. Donaldson, J. M. D. Hill and D. B. Skillicorn entitled "Predictable Communication on Unpredictable Networks: Implementing BSP over TCP/IP" addresses the problem of supporting the well-known BSP model on networks of workstations running TCP/IP. Techniques that offer high throughput with low variance are presented and evaluated experimentally with encouraging results. N. Agrawal and C. P. Ravikumar consider novel techniques for efficient wormhole routing where deadlock is handled not by avoidance but by detection and recovery in their paper "Adaptive Routing Based on Deadlock Recovery". An adaptive routing technique is presented that ensures that deadlocks are infrequent and two schemes are outlined for deadlock recovery are described and evaluated. M. Ould-Khaoua's paper "On the Optimal Topology for Multicomputers with Fully-Adaptive Wormhole Routing: Torus or Hypercube?" revisits the hypercube versus torus comparison for wormhole routing networks. Whereas previous studies, generally based on nonadaptive routing, indicated the superiority of torus topologies, the paper suggests that under certain conditions the hypercube may offer superior performance when adaptive techniques are employed. The fundamental h -relation routing problem is the focus of the paper "Constant Thinning Protocol for Routing in Complete Networks" by A. Kautonen, V. Leppänen

and M. Penttonen. They present and evaluate a randomized algorithm for this problem on the OCP model with encouraging results.

T. Grün and M. A. Hillebrand's paper, "NAS Integer Sort on Multi-Threaded Shared Memory Machines", proposes novel techniques for improving performance in the already novel environment of multi-threaded computer architecture. They use both the commercial Tera MTA and the experimental SB-PRAM machines as platforms for these techniques. In "Analysing a Multistreamed Superscalar Speculative Instruction Fetch Mechanism", R. R. dos Santos and P. O. A. Navaux address the possibility of a more conservative step toward multi-threading within the more traditional super scalar environment. Their primary target is the performance loss due to instruction fetch latency resulting from miss predicted branches. Not surprisingly we find detailed consideration of implementation costs critical in the design of special purpose processor arrays. In "Design of Processor Arrays for Real-time Applications", D. Fimmel and R. Merker study various algorithms for optimizing the layout of uniform processes on processor arrays.

Session 3 is largely devoted to routing algorithms for relatively structured communication such as gossiping or all-to-all scattering. The paper "Interval Routing and Layered Cross Product: Compact Routing Schemes for Butterflies, Mesh of Trees and Fat Trees" by T. Calamoneri and M. Di Ianni show how shortest-path routing information may be maintained in a compact (space-efficient) way for a class of interconnections that includes the butterfly, the mesh of trees and certain fat-trees. The results rely on an elegant exploitation of interval routing techniques. The gossiping problem involves simultaneously broadcasting a set of packets where each node is the origin of a distinct packet and must receive a copy of the packet that originates at every other node. This problem is the subject of both the paper "Gossiping Large Packets on Full-Port Tori" by U. Meyer and J. F. Sibeyn and "Time-Optimal Gossip in Noncombining 2-D Tori with Constant Buffers" by M. Šoch and P. Trvdík. The former presents a near-optimal algorithm for two- and higher-dimensional tori that are time-independent (each node repeats the same simple routing action repeatedly for the duration of the algorithm's execution), whereas the latter presents a time-optimal algorithm for two-dimensional tori that requires only a constant amount of buffer-space per node. The paper "Divide-and-Conquer Algorithms on Two-Dimensional Meshes" by M. Valero-García, A. González, L. Díaz de Cerio and D. Royo explores techniques for supporting divide-and-conquer computations on the two-dimensional mesh. The approach provides efficient solutions to this problem under both the wormhole and store-and-forward routing regime. The all-to-all-scatter problem involves routing a set of distinct packets, one for each source-destination node-pair, where each packet is to be routed from its source to its destination. The paper "Average Distance and All-to-All Scatter in Kautz Networks" by P. Salinger and P. Trvdík describes an asymptotically optimal algorithm for this problem for the Kautz network.

Modern ccNUMA systems represent the merger of easy to build distributed memory systems and easy to program SMP systems. In their paper, "Reac-

tive Proxies: a Flexible Protocol Extension to Reduce ccNUMA Node Controller Contention”, S. A. M. Talbot and P. H. J. Kelly propose and evaluate a technique for distributing the contention that arises naturally in ccNUMA systems when processes share a data structure. The approach is novel in that the degree of distribution is driven by the degree of contention and that well behaved applications do not suffer because of the protocol. The flip side of performance is always reliability. T. Skeie presents techniques for “Handling Multiple Faults in Wormhole Mesh Networks”. The infrastructure for large distributed systems needs to provide both performance and reliability. The presence of multiple paths between nodes on a mesh provides both. In this case, fault tolerance is achieved without dramatic losses in performance. In another instance of novel approaches in modern computer architecture, “Shared Control Supporting Control Parallelism using a SIMD-like Architecture”, N. Abu-Ghazaleh and P. Wilsey provide mechanisms for supporting control parallelism in SIMD systems. The aim is to provide flexibility in conditional and similar computation that does not require global redundant computation necessary in traditional SIMD systems.

The members of the workshop committees wish to extend their sincere thanks to those who acted as reviewers during the selection process for their invaluable help during that period and to all of those who by their submissions have contributed to the success of this workshop.