Workshop 02

Routing and Communication in Networks

Workshop 02: Routing and Communication in Interconnection Networks

Ernst Mayr and Chita Das

All aspects of communication, including routing and communication algorithms, the design and packaging of interconnection networks, and the communication costs of parallel algorithms, are within the scope of this workshop.

Particular, but not exclusive topics of interest include:

- routing algorithms,
- communication costs of parallel algorithms,
- interconnection networks,
- fault-tolerant communication,
- synchronization in parallel computers,
- deadlock-free routing,
- algorithms for collective communication,
- graph embedding.

This year, we received a total of 31 submissions for this workshop (which, thus, turned out to be one of the most "attractive" of the conference), and the programme committee selected 14 papers. which very much reflect the breadth, richness and high quality of research ongoing in the routing and communication area.

The first session of presentations in this workshop is entitled *Fat-Trees and Wormhole Routing* and has five papers. This session deals with two important technologies for modern routing and communication networks. The first is the concept of wormhole routing, relying on fast communication circuitry and avoiding the big and expensive overhead for queueing packets. The first paper of this session presents a simple randomized algorithm for the total exchange of messages on a wormhole-routed torus. It also presents impressive simulation results supporting the claim that this algorithm can provide high performance for many parallel applications. The second paper, on "Deflection-Based Wormhole Routing with Virtual Circuits" combines wormhole routing and deflection routing to provide efficient deadlock-free adaptive routing. It takes advantage of the virtual channels to avoid deadlocks where deflection routing provides alternate paths. The third paper deals with deadlock and deadlock prediction in wormhole routed networks. It is shown that wormhole deadlock prediction is a hard problem and so is an optimal deadlock avoidance mechanism. The second main topic of this first session are fat trees, a very versatile and powerful interconnect structure. The fourth paper of the session analyzes the cost of performing broadcasting, product and prefix computations on a generic model of fat-trees. Algorithms are developed and lower bounds are derived considering the channel capacity at different levels of the tree. The last paper of the session then deals with fault

tolerance properties of fat-trees. It examines the reliability of fat-trees in the presence of faults and analyzes their robustness.

The second session of the workshop is on *Meshes and Tori.* It contains five papers dealing with routing and embedding in mesh or torus-like architectures. The first paper proposes a nice and elegant shortest path routing algorithm for a triangular grid topology. Starting with an infinite grid, the paper discusses minimal routing in a family of Cayley graphs. The second paper develops an algorithm for optimal embedding of k-ary complete trees into 2-dimensional square meshes. The embeddings have load 1, optimal dilation and expansion 2. The next paper presents an algorithm for optimal all-to-all broadcast in store-and-forward noncombining 2-dimensional tori. The algorithm is based on a minimal height construction of a sufficient number of time-arc disjoint spanning trees. The fourth paper studies the cutwidth of the mesh of d-ary trees. The cutwidth of a graph is an important parameter e.g., when embedding the graph in the plane, with all nodes on a line. The paper improves on previously known upper and lower bounds. The last paper of the session studies the emulation of common graphs such as meshes, fat-trees and hypercubes on d-dimensional static multichannel meshes of optical buses. This model uses wavelength division multiplexing, and the goal of the paper is minimizing the number of channels used per bus and the physical distance between nodes on a bus.

The third and final session then is on complexity bounds and some models for parallel processor networks. It contains four papers. The first paper outlines several approaches and models for routing in asynchronous networks, and for the analysis of the corresponding algorithms. The following paper presents new techniques for deriving upper and lower bounds for the compacity of dilation bounded interval routing schemes, i.e., the space required to write down the interval routing information. The third paper considers a mesh architecture with exclusive-write row and column buses. It is shown, in a somewhat more general framework, that in this model solving a variant of the element distinctness problem is hard. The fourth and final paper studies yet another model for parallel computation, the so-called Parallel Alternating-Direction Access Machine has processors communicate via buses connecting memory modules via row and column buses to the processors.

We would like to express our sincere appreciation to the two other members of the programme committee, Prof. Abhiram Ranade and Prof. Pierre Fraigniaud, for their invaluable help in the entire selection process. We would also like to thank the reviewers for their time and effort.