SPAA '92

4th Annual ACM Symposium on Parallel Algorithms and Architectures

June 29 – July 1, 1992
San Diego, California

Sponsored by
ACM SIGACT
ACM SIGARCH

In Cooperation with
EATCS

SPAA '92

4th Annual ACM Symposium on Parallel Algorithms and Architectures

June 29 – July 1, 1992
San Diego, California

Sponsored by
ACM SIGACT
ACM SIGARCH

In Cooperation with
EATCS

The Association for Computing Machinery
1515 Broadway
New York, New York 10036

Soft Cover ISBN: 0-89791-483-X
Series Hard Cover ISBN: 0-89791-484-8

Additional copies may be ordered prepaid from:

ACM Order Department
P.O. Box 64145
Baltimôre, MD 21264

1-800-342-6626
1-410-528-4261 (Outside U.S., MD, and AK)

ACM Order Number : 417920

Foreword

The papers in this volume were presented at the 4th Annual ACM Symposium on Parallel Algorithms and Architectures, held on June 29 - July 1, 1992 at San Diego California. The symposium was sponsored by the Special Interest Group on Automata and Computability Theory (SIGACT) and the Special Interest Group on Computer Architecture (SIGARCH) of the ACM, in cooperation with European Association for Theoretical Computer Science (EATCS).

The 45 papers were chosen by the program committee from the 120 extended abstracts submitted. The Committee's decisions were based on the perceived originality, quality, and appropriateness to the theme of the Symposium. The submissions were not refereed; it is expected that the authors will publish polished and complete versions in scientific journals.

The Program Committee would like to express its thanks to all who submitted papers for consideration.

Program Committee

Allan Fisher, CMU
John Gurd, Univ. of Manchester
John Hennessy, Stanford
H. T. Kung, CMU, Program Chair
Monica Lam, Stanford
Tom Leighton, MIT
Wolfgang Paul, Univ. of Saarland
Larry Rudolph, Hebrew Univ.
Marc Snir, IBM Watson Research
Les Valiant, Harvard
Jean Vuillemin, DEC Paris Research

Conference Organization

Conference Chair:	Lawrence Snyder University of Washington
Local Arrangements Chair:	Lawrence L. Larmore University of California, Riverside
Program Chair:	H.T. Kung Carnegie Mellon University

SPAA Corporate Affiliates

Computer Technology Institute	Patras, Greece
IBM Almaden Research Center	San Jose, CA
IBM Argentina	Buenos Aires, Argentina
Jones and Bartlett Publishers	Boston, MA
Massachusetts Institute of Technology	Cambridge, MA
Morgan Kaufmann Publishers	San Mateo, CA
NEC Research Institute	Princeton, NJ

.A'92: 4th Annual ACM Symposium on Parallel Algorithms and Architectures

following two pages were omitted from the SPAA'92 proceedings, 4th Annual ACM posium on Parallel Algorithms and Architectures. The figures are referred to in the article .ASH 2, " pages 316 - 322, by Jeffrey M. Arnold, Duncan A. Buell, and Elaine G. Davis, of Supercomputing Research Center in Bowie, Maryland.

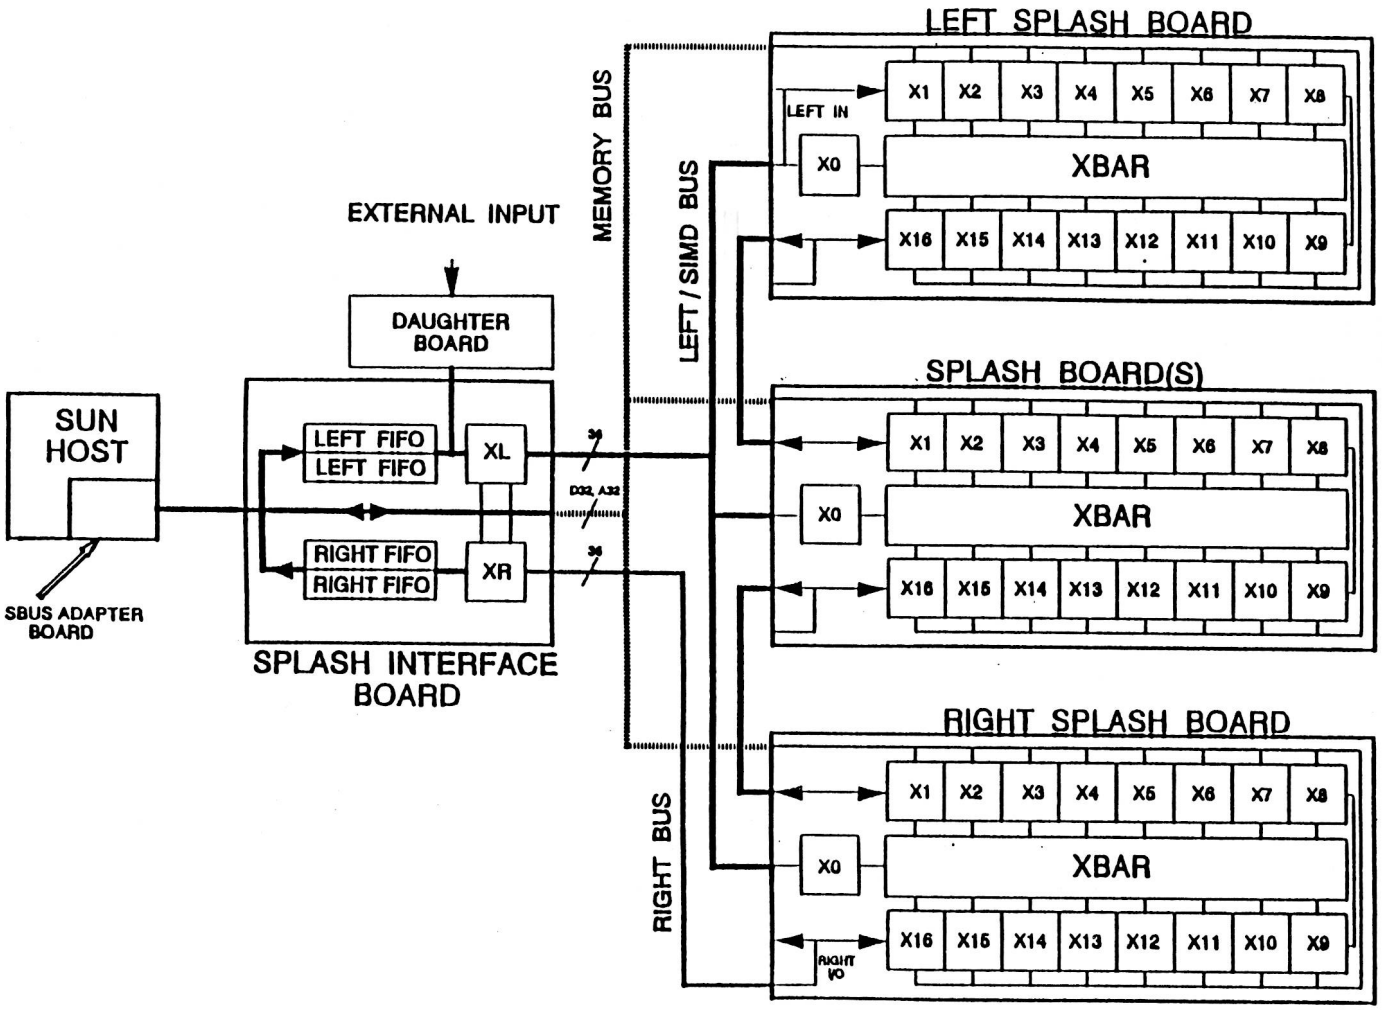

Figure 1: The SPLASH II System

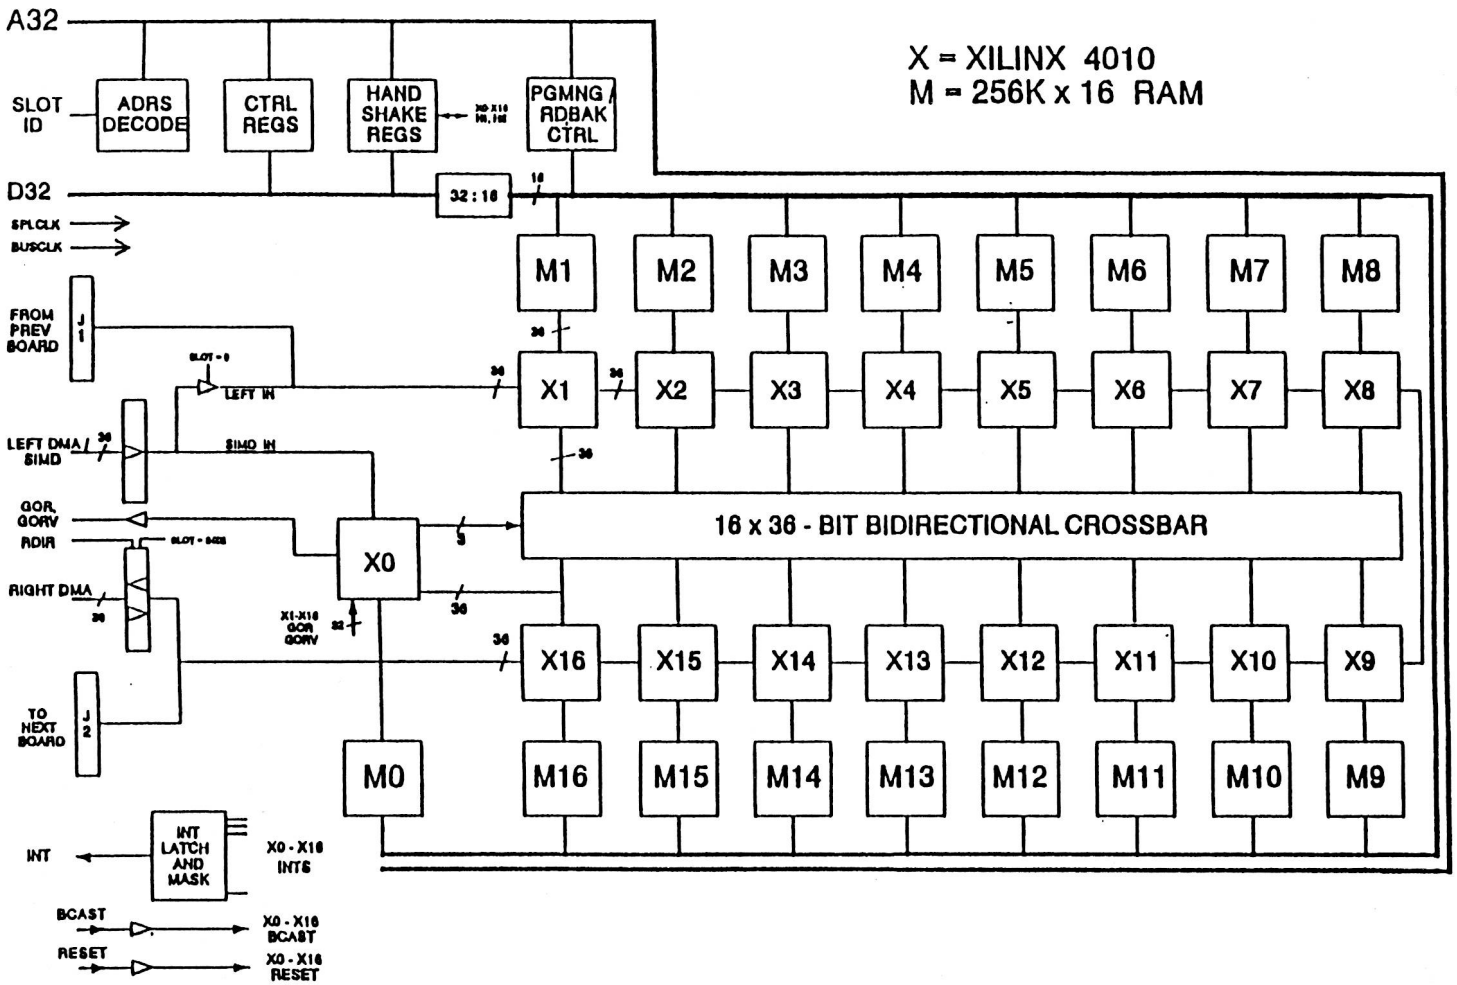

Figure 2: The SPLASH II Board

Keynote Speech

Charles L. Seitz
Caltech

"The Caltech Mosaic C: an Experimental, Fine-Grain Multicomputer"

SESSION 1: Routing

TABLE OF CONTENTS

Monday, June 29, 1992

Session 1: Routing

Session 2: Sorting and Interconnection Networks

Session 3: Parallel Algorithms on Networks

Tuesday, June 30, 1992

Session 4: Parallel Algorithm Implementations

Session 5: Asynchronous Parallel Computing

Session 6: Network Fault-Tolerance and Embedding

Wednesday, July 1, 1992

Session 7: Parallel Architectures

Session 8: PRAM Algorithms

Keynote Speech

Charles L. Seitz
Caltech

"The Caltech Mosaic C: an Experimental, Fine-Grain Multicomputer"

SESSION 1: Routing

Adaptive Deadlock- and Livelock-Free Routing With All Minimal Paths in Torus Networks

Pablo E. Berman *†

Luis Gravano †‡
e-mail: gravano@buevm2.vnet.ibm.com

Gustavo D. Pifarré †‡§
e-mail: pifarre@buevm2.vnet.ibm.com

Jorge L. C. Sanz ‡†
e-mail: sanz@ibm.com

Abstract

This paper consists of two parts. In the first part, a new algorithm for deadlock- and livelock-free routing for the n-dimensional torus network is presented. This algorithm, called **-Channels*, is *fully-adaptive minimal*, i.e. all paths with a minimal number of hops from source to destination are available for routing. **-Channels* works for messages of unknown size, thus yielding new routing techniques for both packet-switched and *worm-hole* models. **-Channels* differs radically from the packet-switched fully-adaptive minimal methods presented in SPAA '91 by Pifarré, Gravano, Felperin, and Sanz [PGFS91]. In particular, the packet-based techniques in [PGFS91] do not work for worm-hole routing as deadlock situations can be constructed.

**-Channels* requires only five virtual channels per bidirectional link of the n-dimensional torus. In fact, only three virtual channels are necessary for the links in one of the dimensions, thus yielding a total of $10(n-1)+6$ buffers per node.

This bound gives the smallest number of channels known in the literature for fully-adaptive minimal *worm-hole* routing. Previous algorithms for fully-adaptive worm-hole routing in n-dimensional tori require more than $O(2^{n-1})$ virtual channels per bidirectional link (see [LH91]). In addition, these results also yield the smallest number of buffers known in the literature for packet-switched fully-adaptive minimal routing.

In the second part of this paper, a comparison of four worm-hole techniques in the 2-dimensional torus is shown in terms of activity in the routing nodes and experimental performance evaluation. Simulation results on the performance of the four worm-hole algorithms are shown for dynamic injection models. Meaningful comparisons required the equalization of the number of virtual channels for all of the techniques, and this is accomplished by resorting to the concept of lane-channels introduced in [Dal90]. The performance of these schemes is measured for different traffic models: random and bit-reversal. Two worm lengths are tried.

1 Introduction.

Message routing in large interconnection networks has attracted a great deal of interest in recent years. Different underlying machine models and hardware architectures have been used and proposed [DS87], [RBJ88, Ran85], [Upf89, LM89], [Val88], [KS90], [NS87], [Hil85], [CEDK91], [BCC+88], [LLK+91].

A desirable feature of routing algorithms is adaptivity, i.e., the ability of messages to use alternative paths toward their destinations according to traffic congestion in the nodes of the network. The amount of hardware resources grow with the degree of adaptivity desired, and resources may become critical if deadlock and live-

* ESLAI, Escuela Superior Latino Americana de Informática, CC 3193,(1000) Buenos Aires, Argentina.

† Computer Research and Advanced Applications Group, IBM Argentina, Ing. E. Butty 275, (1300) Buenos Aires, Argentina.

‡ Computer Science Dept., IBM Almaden Research Center, San José, California.

§ Departamento de Computación, Fac. de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Argentina.

SPAA '92 - 6/92/CA
©1992 ACM 0-89791-484-8/92/0006/0003 ..... $1.50