

Multicomputers and Image Processing

Algorithms and Programs

Edited by

Kendall Preston, Jr.

Leonard Uhr

Multicomputers and Image Processing

Algorithms and Programs

Edited by

Kendall Preston, Jr.

Department of Electrical Engineering
Carnegie-Mellon University
and Department of Radiation Health
University of Pittsburgh
Pittsburgh, Pennsylvania

Leonard Uhr

Department of Computer Sciences
University of Wisconsin, Madison
Madison, Wisconsin

1982



ACADEMIC PRESS

A Subsidiary of Harcourt Brace Jovanovich, Publishers

New York London

Paris San Diego San Francisco São Paulo Sydney Tokyo Toronto

Academic Press Rapid Manuscript Reproduction

COPYRIGHT © 1982, BY ACADEMIC PRESS, INC.
ALL RIGHTS RESERVED.

NO PART OF THIS PUBLICATION MAY BE REPRODUCED OR
TRANSMITTED IN ANY FORM OR BY ANY MEANS, ELECTRONIC
OR MECHANICAL, INCLUDING PHOTOCOPY, RECORDING, OR ANY
INFORMATION STORAGE AND RETRIEVAL SYSTEM, WITHOUT
PERMISSION IN WRITING FROM THE PUBLISHER.

ACADEMIC PRESS, INC.
111 Fifth Avenue, New York, New York 10003

United Kingdom Edition published by
ACADEMIC PRESS, INC. (LONDON) LTD.
24/28 Oval Road, London NW1 7DX

Library of Congress Cataloging in Publication Data
Main entry under title:

Multicomputers and image processing.

(Notes and reports in computer science and applied
mathematics)

Papers from a workshop held in Madison, Wis., May 27-
30, 1981.

Includes index.

1. Image processing--Digital techniques--Congresses.

2. Parallel processing (Electronic computers)--

Congresses. 3. Computer networks--Congresses.

I. Preston, Kendall. II. Uhr, Leonard, Merrick.

III. Series.

TA1632.M84 621.36'7 82-1623
ISBN 0-12-564480-9 AACR2

PRINTED IN THE UNITED STATES OF AMERICA

82 83 84 85 9 8 7 6 5 4 3 2 1

Contributors

Numbers in parentheses indicate the pages on which the authors' contributions begin.

- Dan Antonsson (31), *Department of Electrical Engineering, Linköping University, S-581 83 Linköping, Sweden*
- Jean-Luc Basille (99), *Laboratoire C. E. R. F. I. A., Université Paul Sabatier, Toulouse, France*
- Faye A. Briggs (319), *School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Serge Castan (99), *Laboratoire C. E. R. F. I. A., Université Paul Sabatier, Toulouse, France*
- Ronald S. Curtis (307), *Department of Computer Science, State University of New York at Buffalo, Amherst, New York 14226*
- Renato De Mori (193), *Istituto di Scienze dell' Informazione, Università di Torino, Torino, Italy*
- Bernard Delres (99), *Laboratoire C. E. R. F. I. A., Université Paul Sabatier, Toulouse, France*
- Robert J. Douglass (207), *Department of Applied Mathematics and Computer Science, University of Virginia, Charlottesville, Virginia*
- M. J. B. Duff (261), *Department of Physics and Astronomy, University College London, England*
- Charles R. Dyer (409, 453), *Department of Information Engineering, University of Illinois at Chicago Circle, Chicago, Illinois 60607*
- Glenn S. Fowier (431), *Department of Computer Science and Department of Electrical Engineering, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24060*

- Ariel J. Frank (307), *Department of Computer Science, State University of New York at Buffalo, Amherst, New York 14226*
- K. S. Fu (319), *School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Peter Gemmar (87), *Research Institute for Information Processing and Pattern Recognition, Karlsruhe, Federal Republic of Germany*
- Barry K. Gilbert (385), *Biodynamics Research Unit, Department of Physiology and Biophysics, Mayo Foundation, Rochester, Minnesota*
- Attilio Giordana (193), *Istituto di Scienze dell' Informazione, Universita di Torino, Torino, Italy*
- Goesta H. Granlund (19), *Picture Processing Laboratory, Linköping University, S-581 83 Linköping, Sweden*
- F. Gail Gray (431), *Department of Computer Science and Department of Electrical Engineering, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24060*
- Björn Gudmundsson (31, 231), *Department of Electrical Engineering, Linköping University, S-581 83 Linköping, Sweden*
- Concettina Guerra (221), *Istituto di Automatica, University of Rome, Italy*
- Shin-ichi Hanaki (343), *C&C Systems Research Laboratories, Nippon Electric Co., Ltd., Kawasaki-city, Japan*
- Pat Hanrahan (179), *Department of Computer Sciences, University of Wisconsin, Madison, Wisconsin 53706*
- Robert M. Haralick (431), *Department of Computer Science and Department of Electrical Engineering, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24060*
- Kai Hwang (319), *School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Taizo Iijima (361), *Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan*
- Yasushi Inamoto (353), *Fujitsu Laboratories Ltd., Kawasaki, Japan*
- Mitsuo Ishii (353), *Fujitsu Laboratories Ltd., Kawasaki, Japan*
- Ernest Kent (453), *Integrated Systems Laboratory, University of Illinois at Chicago Circle, Chicago, Illinois 60607*
- Thomas M. Kinter (385), *Biodynamics Research Unit, Department of Physiology and Biophysics, Mayo Foundation, Rochester, Minnesota 55901*
- Loren M. Krueger (385), *Biodynamics Research Unit, Department of Physiology and Biophysics, Mayo Foundation, Rochester, Minnesota 55901*
- Björn Kruse (31, 125), *Department of Electrical Engineering, Linköping University, S-581 83 Linköping, Sweden*

- Ichiro Kubota (125), *Institute of Industrial Science, University of Tokyo, Roppongi, Tokyo, Japan*
- H. T. Kung (373), *Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213*
- Pietro Laface (193), *C. E. N. S. — Istituto di Elettrotecnica Generale, Politecnico di Torino, Torino, Italy*
- Christian Lantuéjoul (111), *Centre de Géostatistique et de Morphologie Mathématique, Ecole des Mines de Paris, Fontainebleau, France*
- Jean-Yves Latil (99), *Laboratoire C. E. R. F. I. A., Université Paul Sabatier, Toulouse, France*
- Stefano Levioldi (221), *Istituto Scienze dell'Informazione, University of Bari, Bari, Italy*
- Martin D. Levine (149), *Department of Electrical Engineering, McGill University, Montreal, Quebec, Canada*
- Kenneth Lundgren (19), *Picture Processing Laboratory, Linköping University, S-581 83 Linköping, Sweden*
- Hiroyuki Matsuura (361), *Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan*
- W. M. McCormack (431), *Department of Computer Science and Department of Electrical Engineering, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24060*
- Bruce H. McCormick (453), *Integrated Systems Laboratory, University of Illinois at Chicago Circle, Chicago, Illinois 60607*
- Ahmed Nazif (149), *Department of Electrical Engineering, McGill University, Montreal, Quebec, Canada*
- Hidemitsu Ogawa (361), *Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan*
- Morio Onoe (125), *Institute of Industrial Science, University of Tokyo, Roppongi, Tokyo, Japan*
- James Piper (161), *Medical Research Council, Western General Hospital, Edinburgh, Scotland*
- J. L. Potter (275), *Digital Technology, Goodyear Aerospace Corporation, Akron, Ohio*
- Kendall Preston, Jr. (135), *Department of Electrical Engineering, Carnegie-Mellon University, and Department of Radiation Health, University of Pittsburgh, Pittsburgh, Pennsylvania 15213*
- Anthony P. Reeves (7), *School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Azriel Rosenfeld (253), *Computer Vision Laboratory, Computer Science Center, University of Maryland, College Park, Maryland 20740*
- Denis Rutovitz (161), *Medical Research Council Science, Western General Hospital, Edinburgh, Scotland*

- Makoto Sato (361), *Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan*
- Larry Schmitt (179), *Department of Computer Sciences, University of Wisconsin, Madison, Wisconsin 53706*
- Howard Jay Siegel (241, 331), *Laboratory for Applications of Remote Sensing, and School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Leah J. Siegel (241), *Laboratory for Applications of Remote Sensing, and School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Bradley W. Smith (331), *Laboratory for Applications of Remote Sensing, and School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- S. W. Song (373), *Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213*
- Stanley R. Sternberg (291), *Environmental Research Institute of Michigan, and University of Michigan, Ann Arbor, Michigan 48106*
- James P. Strong (47), *Goddard Space Flight Center, Greenbelt, Maryland 20770*
- Philip H. Swain (241, 331), *Laboratory for Applications of Remote Sensing, and School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907*
- Steven L. Tanimoto (421), *Department of Computer Science, University of Washington, Seattle, Washington 98105*
- Tsutomu Temma (343), *C&C Systems Research Laboratories, Nippon Electric Co., Ltd., Kawasaki-city, Japan*
- Joseph G. Tront (431), *Department of Computer Science and Department of Electrical Engineering, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24060*
- Leonard Uhr (1, 179), *Department of Computer Sciences, University of Wisconsin, Madison, Wisconsin 53706*
- Larry D. Wittie (307), *Department of Computer Science, State University of New York at Buffalo, Amherst, New York*

Preface

Introduction: Fitting Algorithms to Multicomputer Architectures

This book is the second of a set presenting papers from a series of workshops exploring the large and powerful new multicomputer arrays and networks that are today just beginning to be built.

Three major topics are being examined in these workshops: (a) the architectures of multicomputer arrays and networks, with emphasis on those used for image processing; (b) higher level programming languages that aid and encourage the programmer in formulating parallel programs; (c) algorithms and programs that exploit the potentially enormous increases in speed that arrays and networks can offer.

The workshops were deliberately kept small and quite informal, in order to establish a congenial environment for lively and thoughtful discussions and examinations of problems of mutual interest. Many of the workers who are building multicomputer arrays and networks and designing higher level languages, and whose primary focus is on image processing, have participated in these workshops. Many of the leading researchers in image processing, pattern recognition, and perception systems, chosen because they have been taking parallel approaches, have also participated. The great interest and enthusiasm of the participants suggests that they have found these workshops worthwhile. The quality of the formal papers and of the research they describe, as attested to by these first two volumes, suggests that these meetings are beginning to form an important focal point for the researchers developing these new kinds of multicomputers.

The first two meetings, held in Great Windsor Park, near London, in 1979, and on Ischia, in the Bay of Naples, in 1980, focused on first languages and second architectures. The first volume in this set, "Languages and Architectures for Image Processing," edited by Michael Duff and Stefano Levialdi, and published in 1981

by Academic Press, London, contains papers from these two meetings. The papers collected in the present volume are from the third meeting, held in Madison, Wisconsin, May 27-30, 1981.

At the Madison meeting algorithms and programs were emphasized. But each workshop has explored all of the closely related aspects of developing algorithms and programs for new multicomputer architectures, and each of the two volumes contains chapters on each of the major topics. Indeed, a number of papers address several or all of these issues.

This reflects the need to explore this problem in an interrelated way. Algorithms suggest architectures that can most efficiently execute them, and languages in which they can most fluently be expressed, and programmed. Architectures and languages can encourage, or discourage, the efficient programming of algorithms.

The Organization and Contents of This Book

The papers in the present volume are organized to reflect the three major aspects of the problem: user algorithms and programs, higher level languages, and multicomputer architectures (see the Introduction).

User Algorithms and Programs for Arrays and Networks

Architectures provide us with the tool. Languages allow us to express ourselves and use the tool. But our ultimate purpose is to solve problems, in our case image processing and perceptual problems. It is the algorithms, and the programs that embody and execute these algorithms, that finally make the problem-oriented user happy, and justify all the hard work and frequent pain inherent in developing complex new multicomputer systems.

The development of algorithms with the extremely high degree of parallelism that the new highly parallel arrays and other kinds of networks can handle is an unusually difficult, at times almost mysterious, task. Almost all of our previous experience has been with serial computers, that is with computers that do only one thing at a time. Almost all of the algorithms developed to date are serial; and our programs are entirely serial in that they map algorithms onto a serial computer. In sharp contrast, we know that many problems have great degrees of parallelness, even though it is rarely known precisely how great. We are just beginning to attack what will become an increasingly important research problem, the development of parallel methods for solving problems with parallel multicomputers.

Several of the papers in this volume present relatively specific algorithms for specific applications; others present and examine larger structures of algorithms; others examine whole programs.

Examinations of a Variety of Image Processing Algorithms

Anthony Reeves examines basic algorithms for binary array processors (and gives a good succinct description of such systems). The algorithms described handle a wide variety of processes, from image registration and image enhancement to motion detection and object description.

Kenneth Lundgren and Goesta Granlund describe a very general operator that is the basis of and is embodied in Granlund's GOP computer (described briefly in this paper, and more extensively in the first volume, edited by Duff and Levialdi). They also show how this operator is used for contour detection.

Björn Kruse, Björn Gudmundsson, and Dan Antonsson describe the parallel filter processor (FIP) that is the local-neighborhood processing subsystem of PICAP II.

James Strong gives an overview of the very large and fast MPP array that is now being built by Goodyear-Aerospace for NASA-GODDARD, and examines a wide range of basic algorithms.

Peter Gemmar, in addition to describing the FLIP architecture, shows how it can be used for image correlation and object tracking algorithms.

Jean-Luc Basille, Serge Castan, Bernard Delres, and Jean-Yves Latil present a propagation algorithm and examine its implementation on the projected SYMPATI system, a two-level multicomputer with the lower level SIMD (single instruction-multiple data stream) and the upper level MIMD (multiple instruction-multiple data stream) processors.

Christian Lantuejoul describes geodesic segmentation techniques for transforming complex images, and examines their application to several interesting practical problems.

Mario Onoe and Ichiro Kubota describe a new algorithm for implementing the two-dimensional Fourier transform that eliminates the need for matrix transposition. This algorithm is of great significance in reducing synthetic aperture radar holograms to images.

Kendall Preston analyzes the maxmin propagation transform. He further shows how it may be carried out on a two-dimensional cellular logic machine such as the PHP now running at the joint Carnegie-Mellon University-University of Pittsburgh Image Processing Laboratory.

H. J. Siegel, Philip Swain, and Bradley Smith compare their proposed PASM multicomputer with the CDC Flexible Processor on a variety of basic image processing algorithms.

Perception Programs That Embody Complex Sets of Algorithms

The specific image processing algorithms must be combined into a very large and complex total system in order to actually recognize objects and describe whole

scenes. This is an extremely difficult problem, when highly variable and unconstrained real-world images are used, and researchers are only beginning to develop such systems.

Martin Levine and Ahmed Nazif examine parallel and serial techniques for model-based segmentation in multilevel vision systems.

Denis Rutovitz and James Piper examine the problem of chromosome analysis, and how it might be attacked with systems that use a parallel array like CLIP and/or a serial computer.

Leonard Uhr, Larry Schmitt, and Pat Hanrahan describe a complex type of program that embodies a set of lower level and higher level algorithms for effecting visual perception taking a cone/pyramid multilayered converging approach. They then compare different array, network, and pipeline structures with respect to such a structure of algorithms.

Renato De Mori and Attilio Giordana describe a set of complex high-level algorithms for the closely related and comparably difficult problem of recognition of continuous spoken speech. These algorithms form a complete program for this task, one that would most suitably be executed on a network of the sort described in their paper.

General Issues and Problems of Algorithm Development

Robert Douglass examines the perception of occluded objects from the psychological point of view. He also proposes parallel algorithms and appropriate network architectures for handling this difficult and important problem.

Concettina Guerra and Stefano Levialdi survey major models for effecting local operations at a large number of locations in parallel, and explore techniques for describing these models and for expressing algorithms for computer systems built to embody these models.

Björn Kruse and Björn Gudmundsson briefly describe the PICAP system of specialized processors and pipelines, and examine the variety of types of local and global parallelism that it embodies.

Leah Siegel, H. J. Siegel, and Philip Swain examine and compare a variety of different criteria for evaluating the performance of parallel algorithms.

Azriel Rosenfeld reviews work on cellular bounded automata, both one-dimensional and two-dimensional. He also examines extensions to these near-neighbor systems that give them a more general interconnection topology, regular connections at increasing distances, and pyramid connections through successive converging layers.

Michael Duff examines the set of interrelated issues of computer architecture, programming language, and algorithm structure. He works through several interesting examples of algorithms to explore the need for and value of having the programmer be aware of and work with the architecture's structure.

Higher Level Programming Languages for Arrays and for Networks

We have been especially fortunate that virtually every researcher of whom we are aware working on higher level languages for image processing and perceptual tasks using arrays and networks has attended at least one in this series of workshops.

Most of these languages are described in the first volume of this series (edited by Duff and Levialdi). These include the ALGOL-based PIXAL developed by Stefano Levialdi and his associates, Zenon Kulpa's PICASSO system, Björn Gudmundsson's higher level language for PICAP, the language L being developed by Adolfo Guzman and his associates, Leonard Uhr's parallel extensions to PASCAL, and Robert Douglass's proposal for MAC, the first language of which we are aware that attempts to handle both SIMD and MIMD processes.

The present volume extends this list with papers on two important new languages.

Jerry Potter describes the new Fortran-based language being developed for the very large MPP. This language, because of the great potential power of the MPP, and the preponderance of Fortran programmers in the real world, is likely to have a major practical impact.

Stanley Sternberg describes his unusually elegant set-theory-based language for expressing the structural image processing and perceptual operations effected by his Cytocomputer and other cellular-array-motivated architectures.

The Architectural Structure of Arrays, Networks, and Pipelines

Several multicomputer architectures are described in this book. Often the description is combined with an examination of languages or algorithms for that computer (in which case we will refer to that chapter).

The major computer architectures that have been developed can be viewed and categorized from several points of view. Are they serial or parallel, general-purpose, specialized (but still general-purpose) or special-purpose? Does each processor work independently in MIMD mode; or do all processors execute the same instruction, in SIMD mode; or is there a pipeline, each processor continuing to execute the same instruction, but passing results on to the next processor in MISD assembly-line fashion? What kind of architecture does the individual processor have (e. g., 1-bit or 32-bit), and what interconnection topology exists between processors (e.g., bus, ring, cross-point switch, star, n -cube, array, tree, pyramid)? The following set of papers examine several major types of new architectures that are being developed for large parallel-serial multicomputer systems.

Very Large Parallel SIMD Arrays of 1-Bit Computers

Several very large arrays, typically organized in two-dimensional grids with each computer linked directly to its four, eight, or (when a hexagonal design is used) six nearest-neighbors have recently been built, or designed.

In order to build such large arrays, each computer has been made as simple as possible, usually executing 1-bit processes on 1-bit words. Longer strings and larger numbers are therefore processed in bit-serial fashion, as the price of achieving the highly parallel processing over the large array.

These arrays are also given a single controller, so that all computers are executing the same instruction at each moment in time. This makes them "SIMD" (single instruction-multiple data stream) multicomputers.

They are often called "special-purpose;" but since each of their thousands of computers is itself general-purpose (albeit small) that is not really the case. They might more aptly be called "specialized" (and one can argue that any computer is more or less specialized).

These arrays include the 96 by 96 CLIP4 built by Michael Duff and his associates and running at University College London, the 64 by 64 DAP built by Stewart Reddaway and his associates and sold by ICL England, and the planned 128 by 128 MPP, developed by James Strong and others at NASA-Goddard and being designed and built by Kenneth Batchner, Jerry Potter, and others at Goodyear-Aerospace. CLIP4 and DAP are described in the first volume of this series, while the MPP is described in this volume, by Jerry Potter and by James Strong.

Multicomputer MIMD Networks, Fixed and Reconfigurable

A number of networks have been designed, where, typically, more powerful computers are used, interconnected in a wide variety of ways, and not only in a near-neighbor array. These networks are typically designed with a much smaller number of each much more powerful 8-bit, 16-bit, 32-bit, or 64-bit computers, each with its own controller.

Today only two or three such networks with more than 50 computers actually appear to be running. These include Cm* (built at Carnegie-Mellon University) and the system described by Manara and Stringa in the first volume of this series, which is being used in about a dozen Italian and French Post Offices to recognize portions of addresses on letters.

Most of these networks are being designed with the hope they will handle a general mix of programs, rather than with any particular classes of applications in mind. This is turning out to be a very difficult task, and very few of these systems have yet reached the point where the design is ready for actual building.

Among the most promising and interesting of these is Larry Wittie's MICRONET (described with co-authors Ronald Curtis and Ariel Frank). It has 16

LSI-11s in its first prototype version, but it is being designed looking toward future networks of thousands of computers.

Banks of switches can be added, between processors and processors or between processors and memories, so that a network can be reconfigured under program control.

Peter Gemmar has designed and built a reconfigurable "flexible image processing system" (FLIP) with 16 computers that is optimized to handle correlations and other operations frequently used in image processing systems (see also the first volume of this series).

Two large reconfigurable systems of much more powerful computers are being developed by Faye Briggs, Kai Hwang, and K. S. Fu, and by H. J. Siegel and his associates. Both of these very ambitious systems are envisioned to use up to 1024 computers. Although they are being designed to handle a wide mix of problems, the focus is on image processing and other perceptual tasks.

Specialized and Special-Purpose Image Processing Architectures

A variety of more or less specialized networks and image processing systems are being developed.

Shin-ichi Hanaki and Tsutomu Temma describe a multicomputer system designed to handle image processing with a data-flow, template/data-driven approach. They show how the "butterfly" perfect shuffle, to compute fast fourier transforms, can be implemented on a ring of computers built to implement their data-flow approach.

Mitsuo Ishii and Yasushi Inamoto describe a new system for image processing and computer aided design graphics that has an unusually large memory.

Makoto Sato, Hiroyuki Matsuura, Hidemitsu Ogawa, and Taizo Iijima describe an architecture for microcomputers that both puts a reasonably large number of them (32) on a ring and also links them all directly to a "father" node.

Two special-purpose systems of unusual interest are described in this volume: H. T. Kung's systolic array for convolutions and Berry Gilbert's very powerful and amazingly fast multicomputer for real-time three-dimensional x-ray CAT scanning.

A systolic array is a very efficient special-purpose system that has been designed to embody a particular algorithm in hardware (meaning off-the-shelf IC chips today, and specially designed VLSI chips tomorrow). Kung and Song describe a VLSI implementation of the first systolic array (for convolution) that has been designed for an image processing task. It is of special interest because an IC version is now being built by TRW and, to the extent a market develops, will probably be produced commercially (this was described at the Madison workshop by Kung and R. L. Picard, but that paper is not included in this volume).

Gilbert's system to handle three-dimensional CAT scans in real time is now being completed. It may well turn out to be one of the fastest convolvers and one of the

most powerful computers ever built. It will push ECL chips to subnanosecond clock times, and is expected to execute more than 3 billion floating-point instructions per second.

Pipelines

Pipelines of computers can be built, and data pumped through them, much as material flows through factory assembly lines.

Björn Kruse's PICAP systems combine this kind of pipeline with a variety of other more or less specialized processors and general-purpose serial computers, into a total MIMD system that has been carefully tuned for efficient image processing and pattern recognition.

Stanley Sternberg's Cytocomputer, with over 100 computers in the pipe, has extended the pipeline to its greatest length to date. There are plans for new Cytocomputers that will be built with each processor on its own VLSI chip. These could have even longer pipes, when desired.

Proposals for Future Three-Dimensional Networks and Arrays

Three-dimensional networks and arrays are just beginning to be seriously examined and designed. Here we have been especially fortunate that four of the most interesting systems were described at the Madison workshop and are included in this volume.

Charles Dyer has been exploring the design of pyramid computers side-by-side with his closely related work on algorithms for cellular arrays and pyramids.

Steven Tanimoto is designing, and hoping to fabricate, a VLSI chip that will be used to build a pyramid that, roughly, starts with a CLIP4-like array at its base. Tanimoto examines the programming of such a system.

Robert Haralick and his co-workers (W. M. McCormack, Gail Gray, Joseph Trout, and Glenn Fowler) examine architectural issues in using multicomputer structures to handle combinatorial problems.

Bruce McCormick was the architect of ILLIAC-III, one of the very first, and still one of the most interesting, of the 1-bit arrays (unfortunately demolished by fire just before it would have been completed). In this book McCormick, with coauthors Ernest Kent and Charles Dyer, gives an intriguing picture of a three-dimensional system.

This volume is probably best read together with the first volume (edited by Duff and Levialdi). For the first time, so far as we are aware, they give the reader access to papers on virtually all of the major new array, network, and special-purpose architectures emphasizing image processing, and to the major higher level languages for these new multicomputers, along with a wide variety of algorithms and programs.

Contents

Contributors	ix
Preface	xiii
 Toward Very Large Multi-Computers <i>Leonard Uhr</i>	 1
Parallel Algorithms for Real-Time Image Processing <i>Anthony P. Reeves</i>	7
Image Processing Applications Enabled by MIMD Processor Structures <i>Kenneth Lundgren and Goesta H. Granlund</i>	19
PICAP and Relational Neighborhood Processing in FIP <i>Björn Kruse, Björn Gudmundsson, Dan Antonsson</i>	31
Basic Image Processing Algorithms on the Massively Parallel Processor <i>James P. Strong</i>	47
Image Correlation: Processing Requirements and Implementation Structures on a Flexible Image Processing System (FLIP) <i>Peter Gemmar</i>	87
A Typical Propagation Algorithm on the Line-Processor SY.MP.A.T.I.: The Región Labelling <i>Jean-Luc Basille, Serge Castan, Bernard Delres, and Jean-Yves Latil</i>	99

Geodesic Segmentation <i>Christian Lantuéjoul</i>	111
A Fast Algorithm for Processing Synthetic Aperture Radar Signals Without Data Transposition <i>Morio Onoe and Ichiro Kubota</i>	125
Cellular Logic Algorithms for Graylevel Image Processing <i>Kendall Preston, Jr.</i>	135
An Experimental Rule-Based System for Testing Low Level Segmentation Strategies <i>Martin D. Levine and Ahmed Nazif</i>	149
The Balance of Special and Conventional Computer Architecture Requirements in an Image Processing Application <i>Denis Rutovitz and James Piper</i>	161
Cone/Pyramid Perception Programs for Arrays and Networks <i>Leonard Uhr, Larry Schmitt, and Pat Hanrahan</i>	180
Parallel Algorithms for Interpreting Speech Patterns <i>Renato De Mori, Attilio Giordana, and Pietro Laface</i>	193
Computing Occlusion with Locally Connected Networks of Parallel Processes <i>Robert J. Douglass</i>	207
Reflections on Local Computations <i>Concettina Guerra, Stefano Levialdi</i>	221
Parallelism in PICAP <i>Björn Kruse and Björn Gudmundsson</i>	231
Parallel Algorithm Performance Measures <i>Leah J. Siegel, Howard Jay Siegel, and Philip H. Swain</i>	241
Cellular Architectures: From Automata to Hardware <i>Azriel Rosenfeld</i>	253
Parallel Algorithms and Their Influence on the Specification of Application Problems <i>M. J. B. Duff</i>	261
MPP Architecture and Programming <i>J. L. Potter</i>	275
Pipeline Architectures for Image Processing <i>Stanley R. Sternberg</i>	291