



PAX COMPUTER

High-Speed Parallel Processing and Scientific Computing

Tsutomu Hoshino
Tsukuba University, Japan

Translated by
Susan Goldman
Courant Institute, New York University
and
Tsutomu Hoshino

Edited by
Harold S. Stone
IBM Watson Research Center and
Courant Institute, New York University



ADDISON-WESLEY PUBLISHING COMPANY

Reading, Massachusetts • Menlo Park, California • New York
Don Mills, Ontario • Wokingham, England • Amsterdam • Bonn
Sydney • Singapore • Tokyo • Madrid • San Juan

This book is in the **Addison-Wesley Series in Electrical and Computer Engineering**

Harold S. Stone, *Consulting Editor*

Library of Congress Cataloging-in-Publication Data

Hoshino, Tsutomu, 1938 –
PAX computer.

1. PAX computer. 2. Parallel processing (Electronic computers)

QA76.8.P1438H67 1989 004'.35 88-8191

ISBN 0-201-18492-3

Original Japanese edition "PAX Computer"

Copyright © 1985 by Tsutomu Hoshino, published by Ohmsha, Ltd.,
Tokyo, Japan

English edition copyright © 1989 by the Addison-Wesley Publishing Company.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

ABCDEFGHIJ-HA-89

PAX COMPUTER

**High-Speed Parallel Processing
and
Scientific Computing**

PREFACE

ON THE PUBLICATION OF *PAX COMPUTERS*

Recently, the extremely rapid progress in science and technology has demanded marked improvement in the speeds at which large-scale scientific computations can be performed. It is to meet these demands that so-called “supercomputers” are being developed, and there is no telling how great such demands may become. In order to develop the equipment necessary to handle larger and larger calculations, Japan has established a project to construct supercomputers whose scope is scientific and engineering applications of high-speed computing systems. This undertaking is on a national scale, resulting in much interesting discussion concerning the problems that arise during the research and development of such systems. Naturally, this has increased the demand for both extremely high speed devices and methods for highly parallel processing. The range of important problems that must be solved if we are to satisfy both of these requirements is very broad, including both theoretical and practical research topics.

The present work addresses the revolutionary results of parallel-processing research in easily understood terms. It is believed that developments in VLSI technology and parallel-processing algorithms will be actively used toward the solution of many scientific and engineering problems that occur in the natural world. On the other hand, as technology advances, the scientific and engineering applications of computers will undoubtedly change. To meet this need, beginning in

1977, Professor Hoshino and his colleagues have been engaged in research to develop the parallel computer that they subsequently named the *PAX*.

This book presents a detailed description of *PAX* hardware and software, and the issues arising from parallel-processing methods and algorithms for such methods, with references to pertinent historical trends. In order to demonstrate the utility of the *PAX*, it details the results of many representative examples of scientific and engineering problems that have been solved using an experimental version of the machine.

This text can be regarded as a reference work that provides a simple explanation of useful research results, but we should not overlook the many unique ideas of this research team. Moreover, they develop these ideas meticulously and cogently into an actual system, with the result that we see the *PAX* as a concrete example of a parallel computer. We cannot fail to recognize their skillful efforts through which the value of the *PAX* has been thoroughly proven, together with the identification of the research problems yet unsolved that need to be addressed in the next generation of machines.

In the research laboratories of Japan's universities, where funding is rarely available for projects like *PAX*, it is unusual to find a research effort in which not only have the ideas been carefully explored and developed, but an actual prototype has been built and put into operation. It is all the more unusual, therefore, to see these results presented in book format, in a form in which the collection of information provides the foundations to help solve future problems.

Professor Hoshino and his colleagues began this research as experienced computer users, and as such they have dealt directly with the problems other users face in attempting to solve large-scale problems. Therefore, they could point out the importance of parallel processing and the problems to be solved in the future, as well as the significance of the development of *PAX*.

Because the book has self-contained explanations of the most advanced topics, only a basic knowledge of computer systems is required to understand the ideas expressed in the book. Specialists in most engineering or scientific fields, as well as specialists in computer-architecture research will find this work a valuable reference tool.

It is my wish that the reader will recognize not only the scientific value but also the hard work and enthusiasm of the researchers represented here. With this in mind, hopefully young students will adopt equally powerful approaches and attack research problems with equal zeal.

Finally, in the near future, I look forward to seeing this research group achieve their dream of a new world record for processing speed with a new generation of *PAX* supercomputers.

Hideo Aiso, Ph.D.
Professor
Faculty of Science and
Technology
Keio University

AUTHOR'S INTRODUCTION

Each computer constitutes a universe of its own. This book introduces the world of the PAX, a highly parallel computer whose construction began in 1977. We have chosen to make our central focus an examination of the scientific and engineering applications of parallel processing, an area that has yet to be understood by most people.

Only forty some years have passed since the original vacuum-tube computers of the 1940's were built, and the rate of progress during that time is mind-boggling. As we will see when we trace the earliest periods in its history, the computer liberated people from the arduous tasks of computation. Such computations met socioeconomic needs for census taking and accounting and scientific needs for the calculation of trajectories and the construction of numerical tables. Although computational speed has increased 100,000 times during the period from the initial ENIAC machines to the large-scale computers of today, the demands placed on these new devices have increased enormously in terms of both quantity and quality. It would be difficult therefore to regard existing machines as sufficient.

Both Professors Hoshino and Kawai, the senior members of the original PAX research team, have had many years of experience using scientific calculations in the field of nuclear engineering. In the later half of the 1970s, our attention was drawn to the rapid developments in the area of LSI, and we wondered whether it

would be possible to use available computer power to bridge the gap between supply and demand. In 1977 we decided that if typical scientific problems were to be computed at high speeds, they would have to be assigned to parallel computers designed with strict objectives in mind—the machine that resulted from that decision is the PAX computer.

We believe that the computer is nothing more than a tool used to solve concrete, practical problems that cannot be solved otherwise; its existence is based on its utility. Were the computer no longer a useful tool, it would become nothing more than a large piece of junk that would be best thrown away. With this in mind, Chapter 1 traces the relationship between the development of the computer and the scientific and engineering calculations that can be done with computers in order to put these topics in a historical framework.

The basic configuration of the computer, known as its architecture, is viewed as the hardware and software of the actual machine. Users execute applications programs that perform scientific computations. Our viewpoint is that optimally the entire system must be designed for the ultimate goals of the specific scientific calculations to be performed. Chapter 2, therefore, discusses the significant features of scientific calculations and how those features influence the design of architectures. Chapter 3 sets forth the basic principles that underlie parallel-processing architectures, Chapter 4 describes how these ideas have been implemented as hardware and software within the PAX architecture, and Chapter 5 discusses the basic methods used for parallel processing.

Chapter 6 describes representative applications examples that have been possible to implement thus far. Although these examples are typical scientific calculations, in practical terms, they are little more than simplified problems. Here they are used for expository purposes to make the point that it has been possible to use the PAX to speed up these calculations, because such experiments were, in fact, successful.

Chapter 7 reexamines parallel processing from the point of view of the algorithms employed. The chosen topics have been selected because of their general applicability.

Finally, Chapter 8 outlines the vision we have for the future development of the PAX and describes what technical problems must be overcome if we are to see the fulfillment of this “dream.”

The book has been written with great care to make it readable, although we have assumed that the reader has a basic knowledge of computers and of scientific calculations. Naturally, the main focus of a book of this type has to be technical, but in the final chapter, we have introduced the viewpoints and reactions of the many people whom we have encountered during the course research. We think that a work of this type can represent a case study, demonstrating how a new set of scientific theories or a new technology will be accepted by the Japanese community, even though we only outline this in general terms. Therefore, we suggest that those who are not interested in the technical details of the PAX skip to the Epilogue and concentrate on it only.

The development of the PAX, beginning in 1977, involved the efforts of a great number of people. Here we give a brief biography of some of them. In that year, when electric appliance stores began to sell kits that allowed the purchaser to assemble 8-bit microprocessors, Professor Kawai was a scientist working at Hitachi's nuclear power division in Hitachi City. At that time, his company's senior officials had requested that all employees consider possible applications for microprocessors, and report their ideas to the firm. Professor Kawai, who had spent many years working on nuclear reactor core calculations, believed that if he could allocate portions of the calculations to numerous microprocessors connected adjacently, he could increase processing speed by many orders of magnitude. Furthermore, it was his idea that should the implementations of very high speed calculations be possible, the standardization of solution methods would allow the computer to generate programs with these methods built in and thus the computer could take over much of the programming and coding ordinarily performed by the users. He hoped that this would free a large number of able young people from the "slavery of FORTRAN." However, because there was no initial reaction to his proposal it took some time before serious working groups were formed.

In 1977, Hoshino, who was educated as an electrical engineer, was working at the Institute of Atomic Energy, Kyoto University, where he was involved on a daily basis in large-scale computer simulations for nuclear engineering applications. Although he was not a computer scientist, he was well acquainted with computers, and Professor Kawai's proposal reminded him of the famous ILLIAC IV supercomputer, which had been constructed at the University of Illinois. He was not the sort of person to make snap decisions about beginning research in this direction; however, he gradually became possessed with an increasingly strong desire to build an ILLIAC IV-type machine. An article entitled "Microelectronics and Computer Science" [Sutherland and Mead, 1977] that appeared just then in *Scientific American* greatly encouraged his interest. The article's main stress was the great importance attached to the implementation of interprocessor communications in the design of parallel computers.

The first concepts of the architecture of the present PAX, whose most characteristic features are near-neighbor communications among processors and high-speed global synchronization, were formed at the end of 1977. In February 1978, a COSMO Terminal D was purchased from the Aster International Corporation, and PAX hardware construction was started. Jun'ichi Higashino, then a Kyoto University graduate student and now employed at Hitachi Central Research Laboratories, had worked on microcomputers as a hobby and joined the development team as a very active participant. Hoshino, Kawai, and Higashino formed the team that created the basic architecture and hardware for the PAX.

Purchasing computer parts with the limited funds of a university budget and trying to do parallel processing of the neutron diffusion equations for a nuclear reactor are both difficult tasks to accomplish. We began by buying nothing more than a soldering iron and a circuit tester. However, we were fortunate in obtaining the help of Kazuo Sawada, now employed at Fujitsu's Fanac. He was then a

senior in college, but he was a computer "maniac" so skilled that he could easily have been regarded as a semiprofessional. With the cooperation of Higashino and Sawada, the long journey began. The author himself was of the 1960s generation, a generation that was comfortable with the old machine language for the KDC-1 (Hitachi's first transistor computer), and his experience with machines at the most primitive levels made it possible (and not unpleasant) for him to embark on a lengthy research project that would tap that experience.

In February 1979, we completed the PACS-9,¹ a parallel computer that connected nine microprocessors, and successfully programmed it to calculate nuclear diffusion equations. However, the speed of the computation reached only 0.01 MFLOPS for the nine microprocessors as a whole because the floating-point operations had to be executed in software at that time. In the spring of 1979, after Higashino and Sawada graduated and left the university, Akira Yamaoka, then a graduate student and now also at Hitachi's Central Research Laboratories, joined our group and helped with the design of the PACS-32 (subsequently renamed the PAX-32) and the development of its hardware. In order to broaden the applications for the machine, a high-level language was an absolute necessity. Fortunately, we were able to employ Takashi Sato, a researcher at Kyoto University and presently at Intech, to work on the compiler for the language that was later called SPLM. The PACS-32 achieved a speed of 0.5 MFLOPS in March 1980, and in September, thanks to the successful completion of the SPLM compiler, we were able to begin the three-dimensional simulation of the boiling water reactor core. Afterward, the hardware for the PACS-32, which allowed the execution of Monte Carlo methods, was completed by Hachidai Ito, now at Toshiba's Fuchu Works. The Monte Carlo methods were in development through March 1981.

During this period, PACS development fell somewhat within the framework of nuclear engineering research, and we gratefully acknowledge the overall support we received from Professor Wakabayashi of the Atomic Energy Research Institute.

In the spring of 1981, Hoshino moved to the Institute of Engineering Mechanics of the new University of Tsukuba, and as a consequence, the PACS-32 was also transferred there. At about the same time, Professor Kawai left Hitachi for the Faculty of Science and Engineering at Keio University, and Tomonori Shirakawa came to Tsukuba's Institute of Engineering Mechanics from Osaka Prefecture University. The PACS was in a new environment and it entered a new phase. The year 1981 saw the transfer of the PACS and the construction of the hardware diagnostic program by Sumito Arakawa, who was then a senior at the University of Tsukuba and is now at Sony. In 1982, we were joined by Takahisa Kageyama, then also a senior at the University of Tsukuba and now at Toshiba's Fuchu Works; Hidehiko Abe, who is presently employed by Matsushita Electric

1. The machine was called a Processor Array for Continuum Simulation, or PACS, but the name was subsequently changed to PAX.

TABLE I.1

Chapter	Author	Additional research collaborators and those who provided data
Introduction	Tsutomu Hoshino	
1	Tsutomu Hoshino	
2	Tsutomu Hoshino, Toshio Kawai*	
3	Tsutomu Hoshino	Toshio Kawai,* Jun'ichi Higashino, Tomonori Shirakawa
4	Tomonori Shirakawa, Takahisa Kageyama	Jun'ichi Higashino,** Kazuo Sawada,** Akira Yamaoka,** Hachidai Ito,** Takashi Sato**
5	Tsutomu Hoshino	Sumito Arakawa, Hidehiko Abe, Yoshiaki Kaminaga, and others
6		
6.1	Tsutomu Hoshino	
6.2	Tomonori Shirakawa	Tsutomu Hoshino
6.3	Takeshi Kamimura	Tsutomu Hoshino
6.4	Tsutomu Hoshino	Satoshi Sekiguchi, Manami Ejiri, Sumiko Majima
6.5	Tsutomu Hoshino	Kiyo Takenouchi
6.6	Yoshio Oyanagi	Tsutomu Hoshino, Kiyo Takenouchi, Sumiko Majima
6.7	Takeshi Kamimura	Tsutomu Hoshino
6.8	Yoshiyuki Sato	Takeshi Kamimura, Tsutomu Hoshino
6.9	Sumiko Majima Toshiyuki Miyake	Hiroyoshi Shiigai,
6.10	Tomonori Shirakawa	Katsunori Tanibuchi
6.11	Tsutomu Hoshino	Toshio Kawai,* Tomonori Shirakawa, Masahiro Wakatani,** Hachidai Ito,** Sumiko Majima
6.12	Tsutomu Hoshino	
7	Tsutomu Hoshino	
7.8	Toshio Kawai*	
8	Tsutomu Hoshino	Tomonori Shirakawa
Epilogue	Tsutomu Hoshino	
General editor:	Tsutomu Hoshino	

Note that at the time when the research was being undertaken everyone mentioned in the table was affiliated with the University of Tsukuba except those names marked * were affiliated with Keio University and those marked ** were affiliated with Kyoto University.

Company; and Tomonori Shirakawa, mentioned earlier. Work was begun on the PAX-128 machine, an array with 128 processors and a speed of 4 MFLOPS. The PAX-128 was finished in June 1983. Of course, the PAX-128 would not have been possible without the ongoing contributions, albeit at different points in time, of all of the persons previously mentioned.

The PAX had left its initial phase in 1983 and was then in what might be called the "harvest" period in which the researchers began to see wider applications resulting from applied research. In addition, research efforts were simplified by the concentration of the participants within the university. The University of Tsukuba provided a very favorable environment within which we could pursue the broad academic research needed to solve questions of scientific calculation in general. In that atmosphere, the many applications described in Chapter 6 were executed by the PAX, and its performance was evaluated by the following people:

Faculty members—Hiroyoshi Shiigai (Institute of Engineering Mechanics, University of Tsukuba), Toshio Kawai (Faculty of Science and Engineering, Keio University), and Yoshio Oyanagi (Institute of Information Sciences, University of Tsukuba);

Researchers—Sumiko Majima (Institute of Engineering Mechanics, University of Tsukuba); and

Seniors and Graduate students—Takeshi Kamimura (NEC C&C Systems Research Laboratories), Kiyo Takenouchi (Toshiba), Satoshi Sekiguchi (ETL), Manami Ejiri (Chase Manhattan Bank), Katsunori Tanibuchi (Furukawa Electric), Yoshiyuki Sato (Toshiba Fuchu Works), and Toshiyuki Miyake (Asian and Pacific Centre for the Transfer of Technology).

When the PAX was first being developed, although it was thought of as a machine for the user, the most favorable reactions actually came from computer engineers because of progress being made at home and abroad toward the development and marketing of the supercomputer, which led to much greater interest among the general public as well. Absolute performance in terms of very high speeds, however, was not always the major objective. Rather the essential aim of the PAX was to furnish the user with a machine to do computations economically. We expect research on this aspect to continue to move very actively.

Table I.1 lists all of the contributing authors, together with their affiliations, and research collaborators whose data have been consulted. We emphasize here that we are deeply indebted to all parties named in Table I.1.

Finally, we would like to express our great appreciation to all of those who furnished us with programs and data: Tetsuo Kamimura from Nagoya University; Hideharu Amano from Keio University; Isao Katanuma, Makoto Okazaki, Satoshi Ito, Osamu Watanabe, Yasuhiko Ikebe, and Toshiyuki Inagaki from the University of Tsukuba; and Ken'ichi Miura from Fujitsu.

We would also like to thank Toshihiro Iida and Tomomi Hasegawa of the University of Tsukuba who completed the organization of this manuscript, and

finally, we wish to express our deep appreciation to Professor Hideo Aiso of Keio University who wrote the Preface to this book, and to all of the people at Ohmsha Publishing Company whose hard work brought this book from manuscript stage to final publication.

Tsutomu Hoshino
University of Tsukuba
Tsukuba, Japan

EDITOR'S INTRODUCTION

In an era when computer technology in Japan is progressing in giant steps, rarely has the English-speaking world had an opportunity to see a leading research project in great depth. This book is one of those infrequent instances. Researchers can learn many things from the research contributions of Professor Hoshino and his colleagues and can experience the birth and growing pains of an influential supercomputer project.

At first glance, this book seems to be about a second-generation ILLIAC IV. But the details of machine structure in Chapter 4 reveal that the PAX architecture is much more powerful than an ILLIAC IV, and the applications studies in Chapter 6 demonstrate that the machine is extremely efficient over a broad class of problems that are solved by a variety of programming paradigms. Indeed, there are years of experience reflected in these pages that convey a host of successful techniques for capturing the power of parallel computers.

Two major reasons why English-speaking readers will value this book are

1. It shows how to formulate and solve model problems on parallel machines.
2. It describes a specific parallel machine with novel features that are powerful, and yet have been largely overlooked in the West.

Solving parallel problems is a skill that is growing in the scientific community as small-scale parallel machines become available. But Professor Hoshino, nevertheless, through his extensive experience delineates some basic techniques

for solving problems that are occasionally missed in the research community. The most important issue for his architecture (and for any architecture in which certain processor-to-processor communication links are preferred) is that physical problems be mapped so that nearby physical points are mapped to the same processor or to processors connected by a preferred link. In fact, Hoshino demonstrates that many important algorithms for solving physical systems work well on a two-dimensional mesh-connected architecture. They will work at least as well on a hypercube processor, and possibly on other architectures. His algorithms, therefore, are of interest to all parties who write parallel programs, regardless of the machine architecture. The algorithms reflect the underlying physical models that are being solved. Professor Hoshino's adaptations show how to map the physical models to one class of supercomputers. Programmers who cannot use the adaptations directly, such as those writing for a vector architecture, will still find them useful because they reveal how a physical model can be mapped to fit a specific machine.

For this reason the book is ideally suited to engineers and scientists who are trying to solve real problems on parallel supercomputers. But the second point raised earlier is that Professor Hoshino's machine has features that have not been used on the highly parallel machines designed in the West. Consequently, another segment of the book audience are those professional designers of parallel computers and computer architecture researchers who are trying to build a generation of newer, more powerful, parallel machines.

Perhaps the most significant contribution of the PAX machine is its fast, low-cost, global-synchronization logic. The PAX demonstrates that barrier synchronization can be dealt with simply and effectively by means of the PAX design. The bus-based multiprocessors popular in the West perform poorly when executing code that contains barrier synchronization. Yet, in scientific and engineering codes, barrier synchronization seems to be the most prevalent means for synchronizing processors. The West has largely been building parallel machines whose millions of synchronizations per second (MSYPS) capacity does not grow with the number of processors. For barrier synchronization, the MSYPS capacity of the PAX increases almost linearly with the number of processors so that barrier synchronizations are not generally a bottleneck for PAX. Given the experience related in this book, a designer cannot ignore the power of the global synchronizer.

While learning to write better parallel programs and to build better parallel machines are excellent reasons for reading this book, the reader will also appreciate Professor Hoshino's skill at isolating the essential principles so that the reader learns about what truly counts. In several instances, Professor Hoshino quotes widely held beliefs from the literature and provides evidence from the PAX experience why those beliefs are misleading or false. Finally, the reader will be caught up in Professor Hoshino's dream to build the world's fastest supercomputer. Ironically, if this book is successful in influencing the West, it could shatter his dream by creating a host of competitors in the race spurred by his ideas. But

meanwhile Professor Hoshino is forging ahead with the next generation of PAX, the QCDPAX, which is planned to be operational by the time this book appears.

The newest PAX is among the world's fastest machines; it has 480 33.3-MIPS processors organized in a 20×24 two-dimensional torus array for a total peak computational capacity of just under 16 GFLOPS. It is essentially the same architecture as its predecessors described in this book, except that each processor has an added capability to process vectors of data by means of a local floating-point accelerator chip. While the capabilities of this machine seem enormous, the quantum chromodynamics (QCD) codes that the QCDPAX will run converge extremely slowly, and will take months, possibly a full year, to execute. Professor Hoshino and his colleagues are taking on one of the great challenges in computation. Yet, because the QCDPAX is the fifth in a series of machines, there is very little risk that the project will be delayed or defeated by design flaws that are usually detected late in the development effort of a large-scale computer.

The QCDPAX, if successful in performing the QCD calculations, will demonstrate the importance of having a yet more powerful machine, a TERAFLIPS machine, that can perform those calculations in a day. We can speculate on when a TERAFLIPS machine may be built, what it will look like, and who will build it. But we know with certainty that Professor Hoshino and his colleagues will be contenders in the race to build that machine, and hopefully so will you.

Harold S. Stone
Chappaqua, New York

CONTENTS

AUTHOR'S INTRODUCTION ix
EDITOR'S INTRODUCTION xvii

1

THE HISTORY OF PARALLEL PROCESSING 1

1.1	The Need for Scientific Computers	1
1.2	Parallel Processing	3
1.3	Richardson's Dream	6
1.4	The Development of the ILLIAC IV	8
1.5	A Taxonomy of Parallel Computers	9
1.6	The Development of the ILLIAC IV (Continued)	10
1.7	Pipeline Supercomputers	12
1.8	Software Technology	14
1.9	Complexity and Specialization	16
1.10	Parallel Processors with Highly Parallel Arrays	17