


*THE*  
*ARCHITECTURE*  
*OF*  
*SUPERCOMPUTERS*  
TITAN, A CASE STUDY



DANIEL P. SIEWIOREK

PHILIP JOHN KOOPMAN, JR

# THE ARCHITECTURE OF SUPERCOMPUTERS

TITAN, A CASE STUDY

---

---

---

---

This page intentionally left blank

# THE ARCHITECTURE OF SUPERCOMPUTERS

TITAN, A CASE STUDY

---

---

---

---

DANIEL P. SIEWIOREK

*Department of Computer Science  
Carnegie-Mellon University  
Pittsburgh, Pennsylvania*

PHILIP JOHN KOOPMAN, JR.

*Department of Electrical and Computer Engineering  
Carnegie-Mellon University  
Pittsburgh, Pennsylvania*



ACADEMIC PRESS, INC.  
Harcourt Brace Jovanovich, Publishers

Boston San Diego New York  
London Sydney Tokyo Toronto

This book is printed on acid-free paper. (∞)

Copyright © 1991 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.

1250 Sixth Avenue, San Diego, CA 92101

*United Kingdom Edition published by*

ACADEMIC PRESS LIMITED

24–28 Oval Road, London NW1 7DX

Library of Congress Cataloging-in-Publication Data

Siewiorek, Daniel P.

The architecture of supercomputers : Titan, a case study / Daniel P. Siewiorek, Philip John Koopman, Jr.

p. cm.

Includes bibliographical references (p. ) and index.

ISBN 0-12-643060-8 (alk. paper)

1. Supercomputers. 2. Computer architecture. I. Koopman, Phil, 1960– II. Title.

QA76.5.S536 1991

004.1'1—dc20

90-1276

CIP

Printed in the United States of America

91 92 93 94 9 8 7 6 5 4 3 2 1

# CONTENTS

FOREWORD	ix
PREFACE	xv
CHAPTER ONE	
TITAN ARCHITECTURE OVERVIEW	1
1.1 Visualization of Scientific Computations	1
1.2 System Overview	4
1.3 General Operation Principles	9
1.4 From Concept to Realization	10
1.5 The Next Generation	26
CHAPTER TWO	
ARCHITECTURAL EVOLUTION—ISSUES AND SOLUTIONS IN HIGH-PERFORMANCE DESIGN	27
2.1 In the Beginning: Uniprocessors	27
2.2 Memory Performance Hierarchies	30
2.3 Speedup through Specialization	38
2.4 Concurrency through Overlap of Operations	39
2.5 Concurrency through Replication	49
2.6 Vector Operations	61
2.7 Limits to Speedup: Amdahl's Law	68
2.8 Summary	71

<b>CHAPTER THREE</b>	
<b>PERFORMANCE AND BALANCE</b>	<b>73</b>
3.1 Hardware Performance Measures	75
3.2 Kiviat Graphs	81
3.3 System Performance Measures	83
3.4 The Importance of Balance	95
 <b>CHAPTER FOUR</b>	
<b>HIERARCHICAL MODEL OF COMPUTATION</b>	<b>97</b>
4.1 A High-Level Model of Computation	97
4.2 A More Detailed Model of Computation	98
4.3 Primitive Operations	102
 <b>CHAPTER FIVE</b>	
<b>TITAN ARCHITECTURE AND ITS IMPLEMENTATION</b>	<b>105</b>
5.1 System Bus	105
5.2 Memory	107
5.3 Integer Unit	111
5.4 Vector Processing Unit	118
5.5 Graphics Processing Unit	135
5.6 I/O	143
5.7 Reliability Features	144
5.8 An Example of System Operation	148
5.9 Feature Summary	152
 <b>CHAPTER SIX</b>	
<b>ARCHITECTURAL ANALYSIS OF TITAN</b>	<b>153</b>
6.1 System Bus	153
6.2 Memory	153
6.3 Integer Unit	157
6.4 Vector Processing Unit	158
6.5 Graphics Processing Unit	160
6.6 I/O	160
6.7 Analysis of System Throughput and Plumbing Diagram	161
 <b>CHAPTER SEVEN</b>	
<b>BENCHMARKS AND PERFORMANCE MEASUREMENT</b>	<b>167</b>
7.1 LINPACK and LAPACK	167
7.2 Elementary Functions	171
7.3 The P3 Modifications for Improved Performance	173
7.4 The Eleven Rules of Supercomputer Design	177

CONTENTS	vii
APPENDIX A	
THE TITAN BUS	181
A.1 Introduction	181
A.2 Write and Read Bus Transactions	184
A.3 Other Transaction Types	186
A.4 Error Detection and Handling	186
APPENDIX B	
TITAN P2 SPECIFICATION SUMMARY	189
APPENDIX C	
GLOSSARY OF TERMS	191
REFERENCES	195
INDEX	197



*To Benjamin, Gail, and Nora*

# FOREWORD

In December 1985, when the founders were deciding to form Ardent Computer Corporation, I argued that we simply use multiple, MIPS Computer System Inc.'s microprocessors in a multiprocessor configuration in order to get a high performance computer for the technical marketplace. Such a computer would also be useful for the traditional timesharing user who needed lots of mips.

As the company formed in January 1986 and moved forward, the design increased in complexity with the addition of the vector processing unit, turning it into a product Ardent called the graphics supercomputer.<sup>1</sup> The likelihood greatly increased that a product could be designed within the 18 months from the company's formation at a sales price of \$50,000 as described in the business plan. It turned out that Titan took 30 months to design (a schedule fantasy factor of about 1.7) and the selling price was \$80,000 (a factor of 1.6 over the plan). The first Titan provided substantially more power as a supercomputer with 4 vector processors than a sim-

<sup>1</sup> A Supercomputer has evolved to mean the fastest technical computer that can be built at a given time and using the Cray-style, multiple, vector processor architecture. Mini-supercomputers were introduced to mean the fastest computer for about \$1 million. A graphics supercomputer was supposed to be a personal supercomputer costing around \$50,000.

I place the following constraint on anything called a supercomputer—such a computer must be able to deliver the same computation to a user that the user would obtain on a “true” shared, supercomputer.

ple, multiprocessor would have delivered. In fact, to obtain the same peak power for scientific users would require a factor of 4 more, or 16 scalar processors. The second version of Titan won the 1990 prize for being the most cost-effective computer as measured by the University of Illinois Perfect Club Benchmarks. I'm proud to be associated with the design team and their product. They're probably the best team of this size I've ever worked with.<sup>2</sup>

When I came back to head R&D at Ardent Computer Company in the fall of 1987 to help finish the product design, I contacted Professor Siewiorek about writing a book on the design of Titan, the first graphics supercomputer. In January 1986, after helping Ardent form, I went off to start up the National Science Foundation's Computer and Information Science and Engineering Directorate as its first assistant director. This directorate was also responsible for NSF's five, national supercomputer centers, and I spent two years at NSF also being concerned with providing computing and networking for the research community. This is not unlike the past 25 years I spent at Digital Equipment Corporation, where I was responsible for DEC computers, most of which were designed for the technical marketplace.

If one characterizes computers as MISC (minimal instruction set computer, such as the first computers and early minicomputers), RISC (reduced instruction set), or CISC (complex or complete), the modern supercomputer should be characterized as an UCISC (ultra-complex, since it has operations on vector data-types). A supercomputer implementation is even more complex than the architecture would imply because of the need to execute multiple instructions at once, each of which are operating on multiple independent data elements. When something goes wrong while a program is executing (say a page fault), several hundred data items may be somewhere in transit associated with the context of a processor. It's easy to see why Seymour Cray doesn't provide paging or memory management hardware beyond a single relocation register. Thus, another reason to study Titan is the complexity necessary for vector processing, caching, and paging. Only DEC's VAX 9000 is more complex in its architecture (a CISC with a vector instruction set), its implementation (because it has to go fast to compete with the RISC microprocessors), and its packaging technology. Yet the 9000 provides no graphics and not a great deal more computation capability, even though it costs almost an order of magnitude more to buy and probably a factor of 50 times more to design.

<sup>2</sup>The VAX team, though larger, was probably the most talented team. However, VAX and the first implementation, the VAX 11/780, was considerably simpler and highly evolutionary based on the PDP-11.

The motivation for the book was that I was especially proud of the design team and what they had accomplished. In only two and a half years since the beginning of the company, the team designed and shipped a state of the art multiple vector, processor computer (i.e., a supercomputer) that, in addition, had the ability to transform and display almost 100,000 3D polygons per second. Thus, a team of under 50 had designed and built a product that a large company would take at least 4 years and almost 10 times the staff to accomplish. Titan included not only Glen Miranker's architecture, Richard Lowenthal and Jon Rubenstein's hardware, with a dozen custom chips, Tom Diede's graphics subsystem, but Randy Allen's parallelizing, vectorizing compiler, Way Ting's enhancements to UNIX to support multithread operation so that a loop could be executed on multiple processors, Mike Kaplan's great graphics environment, Doré, that enabled a user to exploit the graphics hardware, and Bill Worley's insight and code on numerous parts of the design where speed was essential. I also felt that a book might also help encourage the sale and use of Titan.

In addition to self-serving motivations, I felt strongly that a book on a really good product design is the best way to help others understand the design process and a particular design.<sup>3</sup> I like to read such books (e.g., Bucholz's book on the IBM Stretch, and Thornton's book on the CDC 6600) and write papers, and to read books about design too (Bell and Newell, *Computer Structures: Readings and Examples*, 1971; Siewiorek, Bell, and Newell, *Computer Structures: Principles and Examples*, 1982; and Bell, Mudge, and McNamara, *Computer Engineering: A DEC View of Hardware Design*). There are no books on supercomputer design, and I feel this one is essential for future designers. All modern supercomputers of Convex, Cray, Fujitsu, Hitachi, IBM (with its vector attachment to the 370 architecture), and NEC are organized according to the principles given in this book.

Finally, this book has turned out to be a really good book about architecture generally, simply because Titan uses nearly all the techniques that one would expect in a modern, high speed computer. I hope engineers

<sup>3</sup>The reader might look at my own critical view of Titan, which is given in a video lecture on Titan: 11 Rules for Supercomputer Design and available through University Video Communication, Box 2666, Stanford, Ca. 94309. Ardent and Titan are also described in the book, *High Technology Ventures*, Addison-Wesley, Reading Massachusetts (Bell and McNamara, 1991) about start-up companies. This view includes the various marketing constraints and a discussion of why the first design, P2 lacked the critical scalar performance that made it vulnerable in the market place until the second processor, P3 was introduced. In fact, given the complexity of a supercomputer, it is essential to do a second one following the first, to fix all the flaws inherent in doing a complex design "right" the first time. All "first" supercomputers by an organization have been flawed in some fashion and have required redesign.