


# Alpha Architecture Reference Manual

EDITED BY  
**RICHARD L. SITES**



# Alpha Architecture Reference Manual



## *Contributing Authors*

Richard Witek  
Alpha co-architect

and

Ellen M. Batbouta  
Richard A. Brunner  
Wayne M. Cardoza  
Daniel W. Dobberpuhl  
Robert A. Giggi  
Henry N. Grieb  
Richard B. Grove  
Robert H. Halstead, Jr.  
Michael S. Harvey  
Nancy P. Kronenberg  
Raymond J. Lanza  
Stephen J. Morris  
William B. Noyce  
Charles G. Nylander  
Mary H. Payne  
Audrey R. Reith  
Robert M. Supnik  
Benjamin J. Thomas  
Catharine Van Ingen

Edited by

**Richard L. Sites**  
Alpha co-architect

digital

D I G I T A L P R E S S

Copyright © 1992 by Digital Equipment Corporation

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 prior written permission of the publisher.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

Order number EY-L520E-DP

ISBN 1-55558-098-X

Technical Writer: Charles Greenman

Production Editor: Kathe Rhoades

Technical Illustrator: Lynne Kenison

Cover Design: Marshall Henrichs

The following are trademarks of Digital Equipment Corporation: DEC, the Digital logo, OpenVMS, PALcode, PDP-11, VAX, VMS, and ULTRIX. Cray is a registered trademark of Cray Research, Inc. IBM is a registered trademark of International Business Machines Corporation. OSF/1 is a registered trademark of Open Software Foundation, Inc. UNIX is a registered trademark of UNIX System Laboratories, Inc.

Digital believes the information in this book is accurate as of its publication date; such information is subject to change without notice. Digital is not responsible for any inadvertent errors.

**Alpha  
Architecture  
Reference  
Manual**

# Foreword

---

In the foreword to the *VAX Architecture Reference Manual*, Sam Fuller, Digital's Vice President for Research and Architecture, wrote, "Computer design continues to be a dynamic field; I expect we will see more rather than less change and innovation in the decades ahead." *The Alpha Architecture Reference Manual* demonstrates the accuracy of that prediction.

Alpha follows VAX by about fifteen years. Those fifteen years have witnessed a torrent of change in computer technology, one that shows no sign of abating:

- More than a 1000-fold increase in the performance of microprocessors
- More than a 1000-fold increase in the density of semiconductor memories
- More than a 500-fold increase in the density of magnetic storage devices
- More than a 100-fold increase in the speed of network connections

During the same period, the internal organization of computer systems has changed as well, based on developments such as RISC architecture, symmetric multiprocessing, and coherent distributed systems. Moreover, the fundamental paradigms of computing have changed not once, but several times, with the introduction of personal computers, graphics workstations, local area networks, and client/server computing.

These developments present an enormous challenge for computing in the 21st century. Future computers will be called upon to solve problems of great scale and complexity, worldwide, in a distributed manner. They will have to provide unprecedented performance, flexibility, reliability, and scalability in order to implement a global infrastructure of information, and to give users an untrammelled window on the world.

Alpha is Digital's response to the challenges of 21st-century computing. It represents the culmination of the company's knowledge and belief about how the next generations of computers should be built. Alpha is based on a decade's experimental and engineering work in RISC architecture, high-speed implementation, software compatibility and migration, and system serviceability. It provides the foundation for implementations ranging from mobile computing units to massively parallel supercomputers.

Alpha is designed to handle the largest computing problems of today and tomorrow. When the Alpha architecture is compared to its predecessor, the VAX architecture, two differences stand out immediately. First, Alpha is a 64-bit architecture; VAX is a 32-bit architecture. This means that Alpha's virtual address extends to a 64-bit linear range of bytes in memory. Supporting this extended virtual address space are an extended maximum physical address range (up to 48 bits) and larger pages (8KB to 64KB). Alpha's extended virtual address range allows direct manipulation

of the gigabytes and terabytes of data produced in electrical and mechanical design, database and transaction processing, and imaging.

Second, Alpha is a RISC architecture; VAX is a CISC architecture. RISC stands for Reduced Instruction Set Computer, CISC for Complex Instruction Set Computer. RISC architectures are characterized by simple, fixed-length instruction formats; a small number of addressing modes; large register files; a load-store instruction set model; and direct hardware execution of instructions. CISC architectures are characterized by variable-length instruction formats; a large number of addressing modes; small-to-medium-sized register files; a full set of register-to-memory (or even memory-to-memory) instructions; and microcoded execution of instructions. Alpha's streamlined organization facilitates high-speed implementation in a variety of technologies, while providing strong compatibility with today's programs and data.

The following tabulation contrasts the architectural differences between VAX and Alpha:

	<b>VAX</b>	<b>Alpha</b>
Architecture	CISC	RISC
Virtual address range	32 bits	Up to 64 bits
Physical address range	Up to 32 bits	Up to 48 bits
Page size	512 bytes	8KB–64KB
Instruction lengths	1–51 bytes	4 bytes
General registers	16 × 32 bits	64 × 64 bits
Addressing modes	21	3
Instruction set architecture	General	Load-store
Directly supported data types	Integer, floating, bit field, queue, character string, decimal string	Integer, floating

This book is the culmination of an effort begun three years ago. In that time, Alpha has grown from a paper specification to a cohesive set of chips, systems, and software, spanning the computer spectrum. This achievement is due to the efforts of many hundreds of people in Engineering, Marketing, Sales, Service, and Manufacturing. This book is documentation of, and a tribute to, the outstanding work they have done.

Bob Supnik  
Corporate Consultant,  
Vice President

# Preface

---

The Alpha architecture is a RISC architecture that was designed for high performance and longevity. Following Amdahl, Blaauw, and Brooks,<sup>1</sup> we distinguish between architecture and implementation:

- Computer architecture is defined as the attributes of a computer seen by a machine-language programmer. This definition includes the instruction set, instruction formats, operation codes, addressing modes, and all registers and memory locations that may be directly manipulated by a machine-language programmer.
- Implementation is defined as the actual hardware structure, logic design, and data-path organization.

This architecture book describes the required behavior of all Alpha implementations, as seen by the machine-language programmer. The architecture does not speak to implementation considerations such as how fast a program runs, what specific bit pattern is left in a hardware register after an unpredictable operation, how to schedule code for a particular chip, or how to wire up a given chip; those considerations are described in implementation-specific documents.

Various Alpha implementations are expected over the coming years, starting with the Digital 21064 chip.

## Goals

When we started the Alpha project in the fall of 1988, we had a small number of goals:

1. High performance
2. Longevity
3. Run VMS and UNIX
4. Easy migration from VAX (and soon-to-be MIPS) customer base

As principal architects, Rich Witek and I made design decisions that were driven directly by these goals.

We assumed that high performance was needed to make a new architecture attractive in the marketplace, and to keep Digital competitive.

We set a 15–25 year design horizon (longevity) and tried to avoid any design elements that we thought would become limitations during this time. The design horizon led directly to the conclusion that Alpha could not be a 32-bit architecture: 32-bit addresses will be too small within 10 years. We thus adopted a full 64-bit

---

1. Amdahl, G.M., G.A. Blaauw, and F.P. Brooks, Jr. "Architecture of the IBM System/360." *IBM Journal of Research and Development*, vol. 8, no. 2 (April 1964): 87–101.

architecture, with a minimal number of 32-bit operations for backward compatibility. Wherever possible, 32-bit operands are put in registers in a 64-bit canonical form and operated upon with 64-bit operations.

The longevity goal also caused us to examine how the performance of implementations would scale up over 25 years. Over the past 25 years, computers have become about 1000 times faster. This suggested to us that Alpha implementations would need to do the same, or we would have to bet that the industry would fall off the historical performance curve. We were unwilling to bet against the industry, and were unwilling to ignore the issue, so we seriously examined the consequences of longevity.

We thought that it would be realistic for implementors to improve clock speeds by a factor of 10 over 25 years, but not by a factor of 100 or 1000. (Clock speeds have improved by about a factor of 100 over the past 25 years, but physical limits are now slowing down the rate of increase.)

We concluded that the remaining factor of 100 would have to come from other design dimensions. If you cannot make the clock faster, the next dimension is to do more work per clock cycle. So the Alpha architecture is focused on allowing implementations that issue many instructions every clock cycle. We thought that it would be realistic for implementors to achieve about a factor of 10 over 25 years by using multiple instruction issue, but not a factor of 100. Even a factor of 10 will require perhaps a decade of compiler research.

We concluded that the remaining factor of 10 would have to come from some other design dimension. If you cannot make the clock faster, and cannot do more work per clock, the next dimension is to have multiple clocked instruction streams, that is, multiple processors. So the Alpha architecture is focused on allowing implementations that apply multiple processors to a single problem. We thought that it would be realistic for implementors to achieve the remaining factor of 10 over 25 years by using multiple processors.

Overall, the factor-of-1000 increase in performance looked reasonable, but required factor-of-10 increases in three different dimensions. These three dimensions therefore formed part of our design framework:

- Gracefully allow fast cycle-time implementations
- Gracefully allow multiple-instruction-issue implementations
- Gracefully allow multiple-processor implementations

The cycle-time goal encouraged us to keep the instruction definitions very simple, and to keep the interactions between instructions very simple. The multiple-instruction-issue goal encouraged us to eliminate specialized registers, architected delay slots, precise arithmetic traps, and byte writes (with their embedded read-modify-write bottleneck). The multiple-processor goal encouraged us to consider the memory model and atomic-update primitives carefully. We adopted load-locked/store-conditional sequences as the atomic-update primitive, and eliminated strict read-write ordering between processors.

All of the above design decisions were driven directly by the performance and



longevity goals. The lack of byte writes, precise arithmetic traps, and multiprocessor read/write ordering have been the most controversial decisions, so far.

### **Clean Sheet of Paper**

To run both OpenVMS and UNIX without burdening the hardware implementations with elaborate (and sometimes conflicting) operating system underpinnings, we adopted an idea from a previous Digital RISC design. Alpha places the underpinnings for interrupt delivery and return, exceptions, context switching, memory management, and error handling in a set of privileged software subroutines called PALcode (privileged architecture library code). PALcode subroutines have controlled entries, run with interrupts turned off, and have access to real hardware (implementation) registers. By having different sets of PALcode for different operating systems, the architecture itself is not biased toward a specific operating system or computing style.

PALcode allowed us to design an architecture that could run OpenVMS gracefully without elaborate hardware and without massively rewriting the VMS synchronization and protection mechanisms. PALcode lets the Alpha architecture support some complex VAX primitives (such as the interlocked queue instructions) that are heavily used by OpenVMS, without burdening a UNIX implementation in any way.

Finally, we also considered how to move VAX and MIPS code to Alpha. We rejected various forms of “compatibility mode” hardware, because they would have severely compromised the performance and time-to-market of the first implementation. After some experimentation, we adopted the strategy of running existing binary code by building software translators. One translator converts OpenVMS VAX images to functionally identical OpenVMS Alpha images. A second translator converts MIPS ULTRIX images to functionally identical DEC OSF/1 Alpha images.

Fundamentally, PALcode gave us a migration path for existing operating systems, and the translators (and native compilers) gave us a migration path for existing user-mode code. PALcode and the translators provided a clean sheet of design paper for the bulk of the Alpha architecture. Other than an extra set of VAX floating-point formats (included for good business reasons, but subsettable later), no specific VAX or MIPS features are carried directly into the Alpha architecture for compatibility reasons.

These considerations substantially shaped the architecture described in the rest of this book.

### **Organization**

The first part of this book describes the instruction-set architecture, and is largely self-contained for readers who are involved with compilers or with assembly language programming. The second and third parts describe the supporting PALcode routines for each operating system—the specific operating system PALcode architecture.

### **Acknowledgments**

My collaboration with Rich Witek over the past few years has been extremely rewarding, both personally and professionally. By combining our backgrounds and viewpoints, we have produced an architecture that is substantially better than either of us could have produced alone. Thank you, Rich.

A work of this magnitude cannot be done on a shoestring or in isolation. Rich and I were blessed with a rich environment of dozens and later hundreds of bright, thoughtful, and outspoken professional peers. I thank the management of Digital Equipment Corporation for providing that rich environment, and those peers for making the architecture so much more robust and well-considered.

Three people have especially influenced my views of computer architecture, through personal interaction and landmark machine design: Fred Brooks, John Cocke, and Seymour Cray. This work is built directly upon theirs, and could not exist without them.

The organization, editing, and production of this text in final form is largely the work of Charlie Greenman, whose clear writing is much appreciated.

Richard L. Sites  
May 1992

# A Note on the Structure of This Book

---

The *Alpha Architecture Reference Manual* is divided into three parts, three appendixes, and an index. Each part describes a major portion of the Alpha architecture. Each contains its own table of contents.

The following tabulation outlines the book's contents:

<b>Name</b>	<b>Contents</b>
Part I	Common Architecture This part describes the instruction-set architecture that is common to and required by all implementations.
Part II	OpenVMS Alpha Software This part describes how the OpenVMS operating system relates to the Alpha architecture.
Part III	DEC OSF/1 Alpha Software This part describes how the DEC OSF/1 operating system relates to the Alpha architecture.
Appendixes	The appendixes describe implementation considerations, IEEE floating-point conformance, and instruction encodings.
Index	Index entries are called out by the symbol (I), (II), or (III). Each symbol is associated with the corresponding Part. Index entries for the appendixes are called out by appendix name and page number.

# Part I

# Common Architecture

This part describes the common Alpha architecture and contains the following chapters:

1. Introduction
2. Basic Architecture
3. Instruction Formats
4. Instruction Descriptions
5. System Architecture and Programming Implications
6. Common PALcode Architecture
7. Console Subsystem Overview
8. Input/Output



# Contents

---

Foreword  
Preface

## **Part I / Common Architecture**

- 1 Introduction
- 2 Basic Architecture
- 3 Instruction Formats
- 4 Instruction Descriptions
- 5 System Architecture and Programming Implications
- 6 Common PALcode Architecture
- 7 Console Subsystem Overview
- 8 Input/Output

## **Part II / OpenVMS Alpha Software**

- 1 Introduction to OpenVMS Alpha
- 2 OpenVMS PALcode Instruction Descriptions
- 3 OpenVMS Memory Management
- 4 OpenVMS Process Structure
- 5 OpenVMS Internal Processor Registers
- 6 OpenVMS Exceptions, Interrupts, and Machine Checks

## **Part III / DEC OSF/1 Alpha Software**

- 1 Introduction to DEC OSF/1 Alpha
- 2 OSF/1 PALcode Instruction Descriptions
- 3 OSF/1 Memory Management
- 4 OSF/1 Process Structure
- 5 OSF/1 Exceptions and Interrupts

## **Appendixes**

- A Software Considerations**
- B IEEE Floating-Point Conformance**
- C Instruction Encodings**

**Index**

# Contents

---

## Common Architecture (I)

---

### Chapter 1 Introduction (I)

---

1.1	The Alpha Approach to RISC Architecture .....	1-1
1.2	Data Format Overview .....	1-3
1.3	Instruction Format Overview .....	1-4
1.4	Instruction Overview .....	1-5
1.5	Instruction Set Characteristics .....	1-6
1.6	Terminology and Conventions .....	1-7
1.6.1	Numbering .....	1-7
1.6.2	Security Holes .....	1-7
1.6.3	UNPREDICTABLE and UNDEFINED .....	1-7
1.6.4	Ranges and Extents .....	1-8
1.6.5	ALIGNED and UNALIGNED .....	1-8
1.6.6	Must Be Zero (MBZ) .....	1-9
1.6.7	Read As Zero (RAZ) .....	1-9
1.6.8	Should Be Zero (SBZ) .....	1-9
1.6.9	Ignore (IGN) .....	1-9
1.6.10	Implementation Dependent (IMP) .....	1-9
1.6.11	Figure Drawing Conventions .....	1-9
1.6.12	Macro Code Example Conventions .....	1-9

### Chapter 2 Basic Architecture (I)

---

2.1	Addressing .....	2-1
2.2	Data Types .....	2-1
2.2.1	Byte .....	2-1
2.2.2	Word .....	2-1
2.2.3	Longword .....	2-2
2.2.4	Quadword .....	2-2
2.2.5	VAX Floating-Point Formats .....	2-3
2.2.5.1	F_floating .....	2-3
2.2.5.2	G_floating .....	2-5
2.2.5.3	D_floating .....	2-6
2.2.6	IEEE Floating-Point Formats .....	2-7
2.2.6.1	S_Floating .....	2-8
2.2.6.2	T_floating .....	2-10

2.2.7	Longword Integer Format in Floating-Point Unit .....	2-11
2.2.8	Quadword Integer Format in Floating-Point Unit .....	2-12
2.2.9	Data Types with No Hardware Support .....	2-13

## Chapter 3 Instruction Formats (I)

---

3.1	Alpha Registers .....	3-1
3.1.1	Program Counter .....	3-1
3.1.2	Integer Registers .....	3-1
3.1.3	Floating-Point Registers .....	3-2
3.1.4	Lock Registers .....	3-2
3.1.5	Optional Registers .....	3-2
3.1.5.1	Memory Prefetch Registers .....	3-2
3.1.5.2	VAX Compatibility Register .....	3-2
3.2	Notation .....	3-2
3.2.1	Operand Notation .....	3-3
3.2.2	Instruction Operand Notation .....	3-4
3.2.3	Operators .....	3-5
3.2.4	Notation Conventions .....	3-8
3.3	Instruction Formats .....	3-8
3.3.1	Memory Instruction Format .....	3-9
3.3.1.1	Memory Format Instructions with a Function Code .....	3-9
3.3.1.2	Memory Format Jump Instructions .....	3-10
3.3.2	Branch Instruction Format .....	3-10
3.3.3	Operate Instruction Format .....	3-10
3.3.4	Floating-Point Operate Instruction Format .....	3-12
3.3.4.1	Floating-Point Convert Instructions .....	3-12
3.3.5	PALcode Instruction Format .....	3-13

## Chapter 4 Instruction Descriptions (I)

---

4.1	Instruction Set Overview .....	4-1
4.1.1	Subsetting Rules .....	4-2
4.1.1.1	Floating-Point Subsets .....	4-2
4.1.2	Software Emulation Rules .....	4-2
4.1.3	Opcode Qualifiers .....	4-3
4.2	Memory Integer Load/Store Instructions .....	4-4
4.2.1	Load Address .....	4-5
4.2.2	Load Memory Data into Integer Register .....	4-6
4.2.3	Load Unaligned Memory Data into Integer Register .....	4-7
4.2.4	Load Memory Data into Integer Register Locked .....	4-8
4.2.5	Store Integer Register Data into Memory Conditional .....	4-11
4.2.6	Store Integer Register Data into Memory .....	4-13
4.2.7	Store Unaligned Integer Register Data into Memory .....	4-14
4.3	Control Instructions .....	4-15



4.3.1	Conditional Branch .....	4-17
4.3.2	Unconditional Branch .....	4-19
4.3.3	Jumps .....	4-20
4.4	Integer Arithmetic Instructions .....	4-22
4.4.1	Longword Add .....	4-23
4.4.2	Scaled Longword Add .....	4-24
4.4.3	Quadword Add .....	4-25
4.4.4	Scaled Quadword Add .....	4-26
4.4.5	Integer Signed Compare .....	4-27
4.4.6	Integer Unsigned Compare .....	4-28
4.4.7	Longword Multiply .....	4-29
4.4.8	Quadword Multiply .....	4-30
4.4.9	Unsigned Quadword Multiply High .....	4-31
4.4.10	Longword Subtract .....	4-32
4.4.11	Scaled Longword Subtract .....	4-33
4.4.12	Quadword Subtract .....	4-34
4.4.13	Scaled Quadword Subtract .....	4-35
4.5	Logical and Shift Instructions .....	4-36
4.5.1	Logical Functions .....	4-37
4.5.2	Conditional Move Integer .....	4-38
4.5.3	Shift Logical .....	4-40
4.5.4	Shift Arithmetic .....	4-41
4.6	Byte-Manipulation Instructions .....	4-42
4.6.1	Compare Byte .....	4-44
4.6.2	Extract Byte .....	4-46
4.6.3	Byte Insert .....	4-50
4.6.4	Byte Mask .....	4-52
4.6.5	Zero Bytes .....	4-55
4.7	Floating-Point Instructions .....	4-56
4.7.1	Floating Subsets and Floating Faults .....	4-56
4.7.2	Definitions .....	4-57
4.7.3	Encodings .....	4-58
4.7.4	Floating-Point Rounding Modes .....	4-59
4.7.5	Floating-Point Trapping Modes .....	4-60
4.7.5.1	Imprecise /Software Completion Trap Modes .....	4-62
4.7.5.2	Invalid Operation Arithmetic Trap .....	4-63
4.7.5.3	Division by Zero Arithmetic Trap .....	4-63
4.7.5.4	Overflow Arithmetic Trap .....	4-63
4.7.5.5	Underflow Arithmetic Trap .....	4-63
4.7.5.6	Inexact Result Arithmetic Trap .....	4-64
4.7.5.7	Integer Overflow Arithmetic Trap .....	4-64
4.7.6	Floating-Point Single-Precision Operations .....	4-64
4.7.7	FPCR Register and Dynamic Rounding Mode .....	4-64
4.7.7.1	Accessing the FPCR .....	4-66
4.7.7.2	Default Values of the FPCR .....	4-67