System Design for a Computational-RAM Logic-In-Memory Parallel-Processing Machine

by

Peter M. Nyasulu, B.Sc., M.Eng.

A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Ottawa-Carleton Institute for Electrical and Computer Engineering,
Department of Electronics,
Faculty of Engineering,
Carleton University,
Ottawa, Ontario, Canada

May, 1999

© Peter M. Nyasulu, 1999

National Library of Canada

Bibliothèque nationale du Canada

Acquisitions and Bibliographic Services

Acquisitions et services bibliographiques

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-42803-6

Canada

The undersigned recommend to the Faculty of Graduate Studies and Research

acceptance of this thesis

"System Design for a Computational-RAM Logic-in-Memory

Parallel-Processing Machine"

submitted by Peter M. Nyasulu (B.Sc., M.Eng.) in partial fulfillment of the

requirements for the degree of Doctor of Philosophy

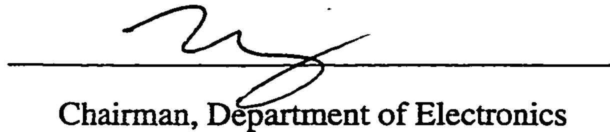

Chairman, Department of Electronics

Professor Jim Wight

Thesis Co-supervisor

Professor Ralph Mason

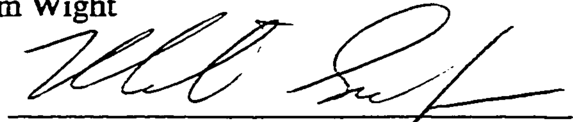

Thesis Co-supervisor

Professor Martin Snelgrove

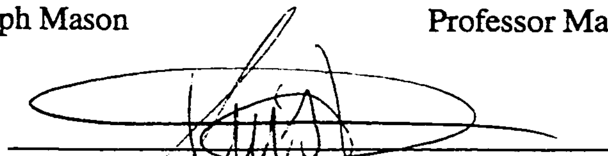

External Examiner

Dr. Lluis Paris

MOSAID Technologies Incorporated

Carleton University

May, 1999

Abstract

Integrating several 1-bit processing elements at the sense amplifiers of a standard RAM improves the performance of massively-parallel applications because of the inherent parallelism and high data bandwidth inside the memory chip. However, implementing such a logic-in-memory system on a host computer poses several challenges because of the small bandwidth at the host system buses, and the different data formats used on the two systems. In this thesis, solutions to these system design issues, including control of the processing elements, interface to the host, data transposition, and application programming, are considered.

A minimal-hardware controller provides high utilization of processing elements while using a simple and general-purpose architecture. A buffer-based host interface unit enhances external data transfers, and minimizes the effect of the host on the performance of the logic-in-memory system. A parallel array-based corner-turning scheme reduces the time to convert data between bit-serial and bit-parallel formats. High-level programming tools, implemented with and using the standard C++ language, hide low-level architectural details of the system, allowing software developers and system analysts to concentrate on implementation details.

Two controller prototypes that interface to the PCI and ISA host buses, and one system prototype, with 64 processing elements and implemented as an ISA expansion card, demonstrate working models of this logic-in-memory system. Simulations using systems with a larger number of processing elements show that the logic-in-memory system yields significant performance speedup over uniprocessor systems when executing massively-parallel applications. Comparisons with other logic-in-memory systems and conventional supercomputers show comparable speed while using less and simpler hardware.

Acknowledgments

I am very grateful to the many people who have contributed to the successful completion of my doctoral studies.

I thank my thesis supervisors, Professor Martin Snelgrove and Professor Ralph Mason, for their guidance and support. I also thank my colleagues, Duncan Elliott, Thinh Le, Christian Cojocaru, Robert Mackenzie, and Philip Lauzon, for their cooperation throughout my research work.

I have also benefited greatly, both academically and financially, from working on the microelectronics bridge camps with Professor John Knight. I am therefore very grateful to him for giving me this opportunity. I also thank Professor Tadeusz Kwasniewski for his support and advice during my early years as a graduate student at Carleton.

I thank the sponsors of my scholarship, the Canadian Commonwealth Scholarship and Fellowship Plan, and acknowledge the help I received from my program managers at Canadian Bureau for International Education and International Council for Canadian Studies.

Studying so far away from home can be very stressful. For this, I am very grateful to the Malawian community in Ottawa and Montreal for the wonderful times we shared together.

I thank my parents and relatives in Malawi for their prayers and love. Most of all I thank my wife and best friend, Stella Nyakasambara, for her unfailing confidence in me, for her patience and love, and for all the help she has provided and all the sacrifices she has made.

Dedicated to my parents,

wadada wa Maynard Nyasulu na wamama wa Bella Nyamkandi

for everything they have done for me

Table of Contents

Abstract i
Acknowledgments ii
Table of Contents iv
List of Figures viii
List of Tables x
List of Abbreviations xi

Chapter 1: Introduction 1
1.1 Computational RAM 2
1.2 Application to Parallel-Processing 3
1.3 System Design 4
1.4 Thesis Scope 6
1.5 Thesis Contributions 7
1.6 Thesis Organization 8

Chapter 2: Logic-In-Memory Systems 9
2.1 Introduction 10
2.2 Logic-in-Memory Systems 13
2.2.1 Integrated Memory Array Processor (IMAP) 13
2.2.2 Terasys Processor In Memory (PIM) Array 15
2.2.3 MIT Pixel-Parallel Image Processing System 16
2.2.4 Intelligent RAM (IRAM) 18
2.2.5 An Integrated Graphics Accelerator and Frame Buffer 18
2.2.6 A Memory-Based Parallel Processor for Vector Quantization 19
2.3 Other Common SIMD Control Path Strategies 20
2.3.1 Hierarchical Controllers 20
2.3.2 Standard Microprogram Sequencers 21
2.3.3 High-Performance Processors 21
2.4 Summary 22

Chapter 3: Computational RAM 23
3.1 Architecture 24
3.1.1 RAM with SIMD Processors 24
3.1.2 Processing Element (PE) 24
3.1.3 CRAM Computations 26
3.2 Prototypes 28
3.2.1 A 64-PE, 128 bits/PE CRAM (C64p128) 28
3.2.2 A 64-PE, 1 Kbits/PE CRAM (C64p1K) 28
3.2.3 A 512-PE, 480 bits/PE CRAM (C512p480) 31

3.2.4 A 1024-PE, 16 Kbits/PE CRAM (C1Kp16K) 31
3.3 Summary 32

Chapter 4: CRAM Controller 33
4.1 Introduction 34
4.2 CRAM-Host Interface Unit 38
4.2.1 CRAM-ISA Interface Unit 39
4.2.2 CRAM-PCI Interface Unit 40
4.3 Instruction Queue Unit 41
4.3.1 Effect of IQU on Performance of PE Array Controller 41
4.3.2 Instruction Fetch Unit and FIFO 46
4.4 Instruction Execution Unit 47
4.4.1 Instruction Word 48
4.4.2 Microinstruction Word 49
4.4.3 Microprogram Sequencer 51
4.4.4 A Simplified Microprogram Sequencer 52
4.4.5 Control Store 54
4.4.6 Address Unit (ADU) 57
4.5 Read/Write Buffers and Constant Broadcast Unit 58
4.5.1 Effect of Read/Write Buffers on Variable Accesses 60
4.5.2 Constant Broadcast Unit 64
4.6 User-Accessible Registers 68
4.7 Memory Map 69
4.8 Summary 70

Chapter 5: CRAM System Implementation 71
5.1 Controller Logic Design and Verification 72
5.1.1 Design Flow 72
5.1.2 VHDL Behavioral Description and Simulation 72
5.1.3 Logic synthesis, Timing Analysis, and Gate-Level Simulation 74
5.1.4 Xilinx FPGA Design Flow 75
5.1.5 ASIC Layout Design Flow 76
5.1.6 CRAM System VHDL Simulation Model 76
5.2 CRAM Controller Implementation and Testing 78
5.2.1 Xilinx Implementation 78
5.2.2 ASIC Simulation 80
5.2.3 CRAM Controller Prototype Testing 81
5.3 ISA CRAM System Prototype 83
5.4 Summary 84

Chapter 6: CRAM System Software Tools 85
6.1 Introduction 86
6.2 CRAM Compiler (CRAM C++ Library) 87
6.2.1 Using a C++ Compiler 87
6.2.2 CRAM Classes (Data Types) 88

6.2.3 Memory Allocation 91
6.2.4 Overloaded Operators 92
6.2.5 Library Functions 93
6.2.6 Data-Parallel Conditional Statements 94
6.2.7 Operations with Scalar Constants 95
6.2.8 Operand Extension 97
6.2.9 Support Classes 97
6.2.10 CRAM System Objects and Their Initialization 99
6.3 Corner-Turning: Host Access of CRAM Variables 100
6.3.1 Host-Based Data Transposition 101
6.3.2 Parallel Array-Based Data Transposition 103
6.3.3 Effect on Performance 106
6.4 CRAM Assembler 107
6.4.1 Introduction 107
6.4.2 CRAM Assembly Language 108
6.4.3 Using cvar Addresses and C++ Integer Variables 109
6.5 CRAM Microcode Assembler 110
6.6 Grouping of Microroutines 110
6.7 CRAM C++ Simulator 113
6.7.1 Introduction 113
6.7.2 Simulation Model 114
6.7.3 Host Objects 114
6.7.4 CRAM Controller 115
6.7.5 CRAM PEs and RAM 116
6.7.6 Timing Information 116
6.7.7 Debugging Information 118
6.7.8 Program Execution Statistics 119
6.8 Summary 120

Chapter 7: Applications and Performance Analysis 121
7.1 Analysis Methodology and Parameters 122
7.2 Basic Arithmetic, Logic and Memory Operations 123
7.3 Low-Level Image Processing 126
7.3.1 Brightness Adjustment 127
7.3.2 Spatial Average Low-Pass Filtering 127
7.3.3 Edge Detection and Enhancement 130
7.3.4 Segmentation 131
7.4 Database Applications 132
7.4.1 Basic Searches 133
7.4.2 Least Mean Squared (LMS) Match 134
7.5 Image and Video Compression 135
7.5.1 Vector Quantization 135
7.5.2 MPEG-2 Motion Estimation 138
7.6 Performance Analysis 141
7.6.1 Applications Performance Summary 141
7.6.2 Controller and System Performance 142

7.6.3 Degree of Parallelism 145
7.6.4 Comparison with other SIMD Systems 147
7.7 Summary 149

Chapter 8: Conclusions and Future Work 150
8.1 Summary 150
8.2 Future Work 152
8.2.1 An On-Chip Controller and a MIMD-SIMD CRAM System 152
8.2.2 A Pipelined Constant-Sensitive PE 153
8.2.3 Other Work 154

References 155
Appendix A: CRAM Controller Architectural Details 159
Appendix B: CRAM System PCBs and Pinouts 166
Appendix C: CRAM Software Details 168
Appendix D: Applications Source Code 173

List of Figures

Figure 1.1: CRAM Architecture and Processing Element 2
Figure 1.2: Implementation of Image Inversion on CRAM 3
Figure 1.3: CRAM in a Typical Computer System 5
Figure 2.1: Memory Bandwidth in a Computer System 11
Figure 2.2: Integrated Memory Array Processor Configuration 13
Figure 2.3: RVS-2 Configuration 14
Figure 2.4: RVS-2 Controller (RVSC) 14
Figure 2.5: A 16K Processor Terasys Workstation 15
Figure 2.6: MIT Image Processing System 16
Figure 2.7: MIT Controller Architecture 17
Figure 2.8: Instruction Selection 17
Figure 2.9: Organization of an IRAM Vector Processor 18
Figure 2.10: Block Diagram of the FMPP-VQ64 19
Figure 2.11: VASTOR Controller Hierarchy 20
Figure 2.12: LUCAS Control Unit 21
Figure 3.1: CRAM Architecture 24
Figure 3.2: Baseline Processing Element 25
Figure 3.3: Extended Processing Element 26
Figure 3.4: Setting PE Opcodes (TTOP and COP) 27
Figure 3.5: PE 4-bit Addition 27
Figure 3.6: PE Cycle External Timing 29
Figure 3.7: PE Cycle Internal Timing 30
Figure 4.1: CRAM Controller Architecture 37
Figure 4.2: CRAM-ISA Interface Unit 39
Figure 4.3: CRAM-PCI Interface Unit 40
Figure 4.4: Instruction Queue Unit 41
Figure 4.5: Effect of IQU on Controller Performance 44
Figure 4.6: Effect of FIFO Size on Performance 47
Figure 4.7: instruction Format 48
Figure 4.8: Microinstruction Word 49
Figure 4.9: Microprogram Sequencer 51
Figure 4.10: A Simplified Microprogram Sequencer 53
Figure 4.11: Effect of Control Store Size on Performance 56
Figure 4.12: Address unit 58
Figure 4.13: Read/Write Buffer and Constant Broadcast Unit 59
Figure 4.14: Data Buffers Organization 60
Figure 4.15: Effect of R/W Buffers on Variable Initialization 63
Figure 4.16: Constant Broadcast Unit 65
Figure 4.17: Effect of Buffer and Constant Unit on Operate-Immediate Instructions 67
Figure 4.18: User-Accessible Registers 69
Figure 4.19: CRAM Controller Memory Map 69
Figure 5.1: Design Flow 73

Figure 5.2: CRAM System VHDL Simulation Model 77
Figure 5.3: CRAM Controller Test Fixture 81
Figure 5.4: CRAM System ISA Card 83
Figure 6.1: CRAM System Software Tools 86
Figure 6.2: Definitions of CRAM Data Types 88
Figure 6.3: Host Corner-Turning 101
Figure 6.4: Parallel Array-Based Data Transposition 103
Figure 6.5: Array-Based vs. Host-Based Transposition 105
Figure 6.6: Transposing N-Byte Variables 106
Figure 6.7: Data Transpose Time as Percentage of Total I/O Overhead 107
Figure 6.8: Grouping of Microroutines 111
Figure 6.9: CRAM C++ Simulation Model 115
Figure 7.1: Effect of Degree of Parallelism on Basic Operations 124
Figure 7.2: Effect of Operand Size on Basic Operations 125
Figure 7.3: Computing in a 3 x 3 Pixel Neighborhood 129
Figure 7.4: The Laplacian Edge Detection Operator 130
Figure 7.5: CRAM Implementation of Vector Quantization 137
Figure 7.6: Motion Estimation 138
Figure 7.7: CRAM Implementation of Motion Estimation 140
Figure 7.8: Effect of Controller on Applications Execution Times 143
Figure 7.9: Effect of Degree of Parallelism on Applications Performance 145
Figure 7.10: CRAM Speedup for Small Number of PEs 146
Figure 8.1: A MIMD-SIMD CRAM System 153
Figure A.1: PCI Device's Configuration Header 162
Figure B.1: ISA CRAM System PCB Layout 166

List of Tables

Table 4.1: Percentage of Short-Sequence Instructions 45
Table 4.2: Minimum Buffer Size for $T_{exe} \geq T_{xload}$ during Variable Initialization 62
Table 4.3: CRAM Parameter Registers 68
Table 5.1: CRAM Controller FPGA Device Utilization 79
Table 5.2: CRAM Controller Functional Blocks FPGA Utilization 79
Table 5.3: Area of CRAM Controller in TSMC 0.35 mm Technology 80
Table 6.1: CRAM C++ Operators 92
Table 6.2: CRAM Library Functions 93
Table 6.3: Optimizations for Operations with Constants 96
Table 6.4: Instruction Groups 112
Table 7.1: CRAM Basic Operations 123
Table 7.2: Performance of Basic Operations 124
Table 7.3: Performance of Basic Searches 133
Table 7.4: Applications Performance Summary 141
Table 7.5: Controller and System Performance 144
Table 7.6: System Comparisons 148
Table A.1: Next Address Instructions 159
Table A.2: Microsequencer Conditions 159
Table A.3: CRAM Functions 160
Table A.4: CRAM Controller Instructions 160
Table A.5: Auxiliary Controller Instructions 161
Table A.6: Command Register 161
Table A.7: Status Register 162
Table A.8: CRAM PCI Command Registers 163
Table A.9: CRAM PCI Status Register 164
Table A.10: CRAM PCI Base Address Register 0 165
Table A.11: PCI Command Types 165

List of Abbreviations

ACU ....Array Control Unit
ALU ....Arithmetic Logic Unit
ASIC ....Application-Specific Integrated Circuit
BiCMOS ....Bipolar Complementary Metal Oxide Semiconductor
BPE ....Baseline Processing Element
CAD ....Computer-Aided Design
CAM ....Content-Addressable Memory
CASM ....CRAM Assembler
CD ....Compact Disc
CLB ....Configurable Logic Block
CMASM.........CRAM Microcode Assembler
CMC...............Canadian Microelectronics Corporation
CMOS ....Complementary Metal Oxide Semiconductor
COP................Control Opcode
CPLD ....Complex Programmable Logic Device
CPU................Central Processing Unit
CRAM............Computational Random Access Memory
DRAM............Dynamic Random Access Memory
DRC ....Design Rule Check
EDIF...............Electronic Design Interchange Format
EISA...............Extended Industry Standard Architecture
EPROM..........Erasable Programmable Read-Only Memory
FIFO ....First In First Out
FIR ....Finite Impulse Response
FPGA ....Field Programmable Gate Array
FSM ....Finite State Machine
GIPS ....Giga Instructions Per Second
IMAP..............Integrated Memory Array Processor

IC....................Integrated Circuit

IO (I/O) ....Input/Output

IOB.................Input/Output Block

IQU ....Instruction Queue Unit

IR....................Instruction Register

IRAM ....Intelligent Random Access Memory

ISA ....Industry Standard Architecture

LMS ....Least Means Squared

LSI..................Large Scale Integration

LVS ....Layout Versus Schematic

MAE...............Mean Absolute Error

MCK ....Memory Clock

MIMD ....Multiple Instruction Multiple Data

MIPS ....Million Instructions Per Second

MPEG ....Moving Picture Expert Group

MPP ....Massively Parallel Processing/Processor

MSE ....Mean Squared Error

MUX ....Multiplexer

n-D ....*n*-dimensional ($n = 1, 2, 3,...$)

NNS................Nearest Neighbor Search

OPS ....Operate Strobe

PC...................Personal Computer

PCB ....Printed Circuit Board

PCI ....Peripheral Component Interconnect

PE...................Processing Element

PIM ....Processor In Memory

PIPRAM.........Parallel Image Processing Random Access memory

PPR ....Partition, Place and Route

PROM ....Programmable Read-Only Memory

RAM ....Random Access Memory

Reg ....Register

RISC...............Reduced Instruction Set Computer

ROM...............Read-Only Memory

RTL................Register Transfer Level

RW (R/W)......Read/Write

SDF................Standard Delay Format

SSD................Sum of Squared Differences

SIMD..............Single Instruction Multiple Data

SRAM............Static Random Access Memory

TMPA.............Temporary or Mask Address

TTOP..............Truth-Table Opcode

VHDL............VHSIC Hardware Description Language

VHSIC............Very High-Speed Integrated Circuit

VME...............Versa Module Eurocard

VQ..................Vector Quantization

VRAM............Video Random Access Memory

VSS................VHDL System Simulator

XNF................Xilinx Format

XPE................Extended Processing Element

PCI Bus Signals:

AD..................Address/Data bus

CBE................Command and Byte Enable

DEVSEL........Device Select

FRAME..........Cycle Frame

IDSEL............Initialization Device Select

IRDY..............Initiator Ready

PAR................Parity

PERR.............Parity Error

SERR.............System Error

STOP..............Stop Transaction

TRDY.............Target Ready

ISA Bus Signals:

BALE ....Bus Address Latch Enable

CHRDY..........Channel Ready

IRQ*n*...............Interrupt Request Line *n*

LA ....Higher address bus (23:17)

M16................Memory 16-bit Chip Select

MRDC............Memory Read Control

MWTC...........Memory Write Control

NOWS............No Wait State

SA ....Lower address bus (19:0)

SBHE ....System Byte High Enable

SD ....Data bus

Chapter 1

Introduction

The idea of merging logic and memory on a single chip was conceived as a correction measure for the wide performance gap between the CPU (microprocessor) and its main memory (DRAM) [1]. Traditionally, microprocessors and DRAM are fabricated on totally different technology processes. Microprocessor fabrication lines are usually optimized to yield fast logic, whereas DRAM processes are designed to reduce leakage current and increase cell density. This has resulted in a processor-memory performance gap that increases by more than 50% every year [2]. While a number of solutions, such as sophisticated cache schemes and pipelined processors, have been used to correct this gap, these have not been totally effective and the processor-memory performance gap continues to be a major obstacle to improved computer system performance. In the continuing effort to improve system performance, some recent research work has concentrated on equipping standard memories with some processing power. The primary objective of these logic-in-memory or processing-in-memory systems is to utilize the high data bandwidth and inherent parallelism that is available inside the memory chip. This not only improves system performance, but also improves power consumption and in some cases reduces system cost.

Computational RAM (CRAM) is one of the pioneering research projects in logic-in-memory systems. The following sections give a brief overview of CRAM, and outlines the scope, contributions and organization of this thesis.

1.1 Computational RAM

CRAM is a SIMD-memory hybrid in which very simple 1-bit processing elements (PEs) are integrated at the sense amplifiers of a standard RAM [3]. CRAM is designed to increase the speed of executing massively-parallel applications by utilizing the high bandwidth available at the sense amplifiers. Several PEs have access to the data at the sense amplifiers and operate on it without the need to read the data out of the RAM and transmit it over long buses to the processor. This improves performance and reduces power consumption. Figure 1.1 shows the architecture of CRAM and the 1-bit PE.

The computational logic of a PE consists of an 8-to-1 multiplexer and 3 registers (Y, X, M). The PE computes the result of an operation by using the 3 bits from the registers to select any one bit of the 8-bit instruction. Therefore, a PE instruction is simply the multiplexer truth table output for all 8 possible register combinations for that particular operation. For example, to have a PE result of '1', the instruction is set to 0xFF so that any combination of Y-X-M selects '1' as the output of the multiplexer. The W register is used for conditional write-back of the result to the local memory of the PE. The shift ports and the bus tie are used for inter-PE communication.

Since the PEs are only 1-bit wide, a single operation generally requires several instructions. For example, an n-bit addition requires $(6n + 1)$ instructions. However, by making the PEs small, a number of them can be integrated in the pitch of one or a few sense amplifiers in order to increase the degree of computational parallelism.

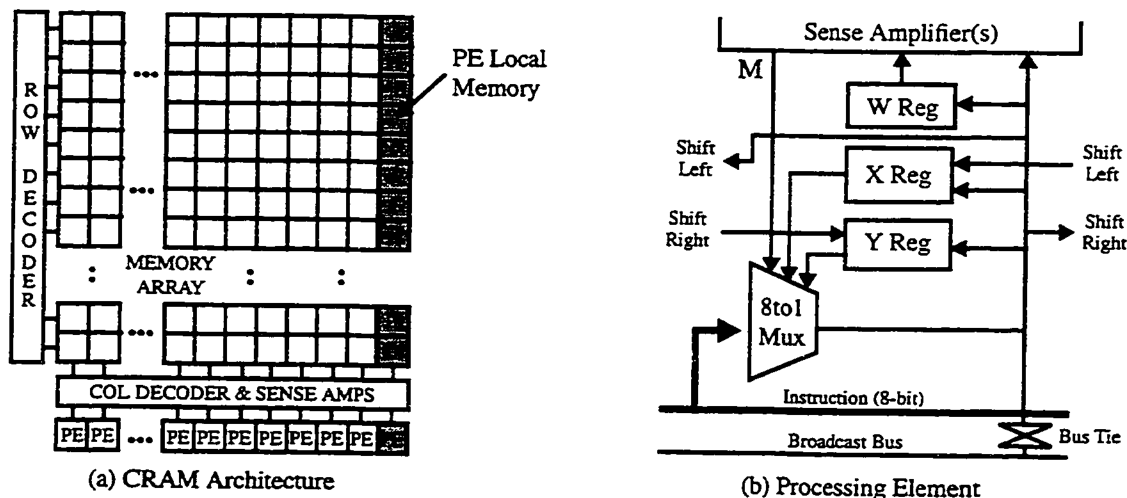

Figure 1.1 CRAM Architecture and Processing Element

1.2 Application to Parallel-Processing

Applications most suited for CRAM, like most massively parallel SIMD machines, are those that have fine grain parallelism and regular communication patterns. Such applications can be found in numerous fields including image processing, databases, video and image compression, digital signal processing, computer-aided design, graphics, and numerical analysis [4], [5], [6], [7]. Specific applications that have been studied for CRAM implementation include image convolution, FIR filters, data mining, fault simulation, and the satisfiability problem [4], as well as general image processing techniques, discrete cosine transform, run-length encoding, scalable and hierarchical vector quantization, and other MPEG-2 algorithms [5]. A few CRAM applications are described in this thesis. These include low-level image processing, basic database applications, vector quantization, and motion estimation.

Figure 1.2 illustrates the implementation of image inversion for a 256 x 256 8-bit image on a 64K-PE CRAM (image inversion is useful in the display of medical images and in producing negative prints of images). The image is spread out across the PEs, with one pixel per PE. All PEs perform the inversion operation in parallel. While each PE takes 24 instructions to invert its pixel, i.e. 1.2 μs for a 20 MHz CRAM, the fact that this is done in parallel for all the 65536 pixels still results in high speedup over a high-performance uniprocessor system. For example, a 400 MHz Pentium processor theoretically takes 7.5 ns to invert one pixel (two memory accesses and one computation). But since it has to

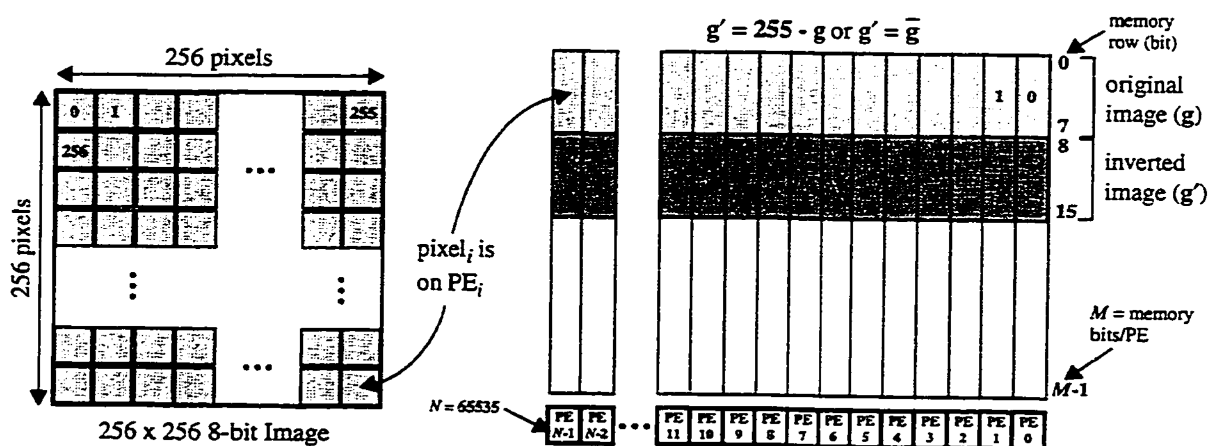

Figure 1.2 Implementation of Image Inversion on CRAM

process 64K pixels sequentially, its total execution time of 0.5 ms is more than 400 times slower than that of the CRAM system. Note that this is just a theoretical number used for illustration. As shown later in the thesis, practical speedups may be slightly lower or higher because of CRAM system design issues. Also, a uniprocessor memory access may not be executed in one CPU cycle because memory systems usually have slower cycle times.

1.3 System Design

As shown later in the thesis, most logic-in-memory systems are designed for specific applications and are implemented mostly on the VME bus or the Sbus. On the other hand, CRAM is designed to be a general-purpose parallel-processing system that can be used on a variety of platforms. One platform specifically targeted in our initial prototypes is the widely-used personal computer (PC) environment. Figure 1.3 shows possible implementations and use of CRAM systems in a typical PC-like environment. The dotted CRAM system is the case where CRAM either replaces or coexists with the standard RAM as the computer main memory or as video RAM. The implementations on the PCI local bus and the expansion buses (shaded boxes) are the ones developed in this work. The main use of the CRAM controller is to allow CRAM application programs to be run from the host computer. The controller acts as a PE array controller as well as an interface to the host bus.

Implementing a CRAM system on a host computer such as the PC poses several challenges. Considering that even a simple operation such as addition requires several instructions to be issued to the 1-bit PEs, the first challenge is to control the PEs and attain high PE utilization while using the small bandwidth of the host buses. Otherwise the PEs would be idle most of the time, thus reducing the overall advantage of CRAM over a uniprocessor system. Second, unlike most SIMD systems that are implemented using several PCBs and housed in large cabinets, a PC-based CRAM system has to be implemented on a single standard-size PC card. This means using as few external components as possible. The third challenge in using CRAM on a standard computer is due to the different data formats used on the two systems. CRAM is bit-serial, while most

conventional computer systems are bit-parallel. Therefore, there is need for format conversion or data transposition when data is transferred between the two systems. This difference in data formats also means that programming tools on the host can not be directly used to write programs for CRAM. Therefore, new software tools have to be developed for CRAM, or existing tools on the host have to be modified or enhanced to support CRAM data types. Finally, there is need to minimize any extra CRAM hardware in order to reduce system costs as well as facilitate future integration of all CRAM hardware on a single chip.

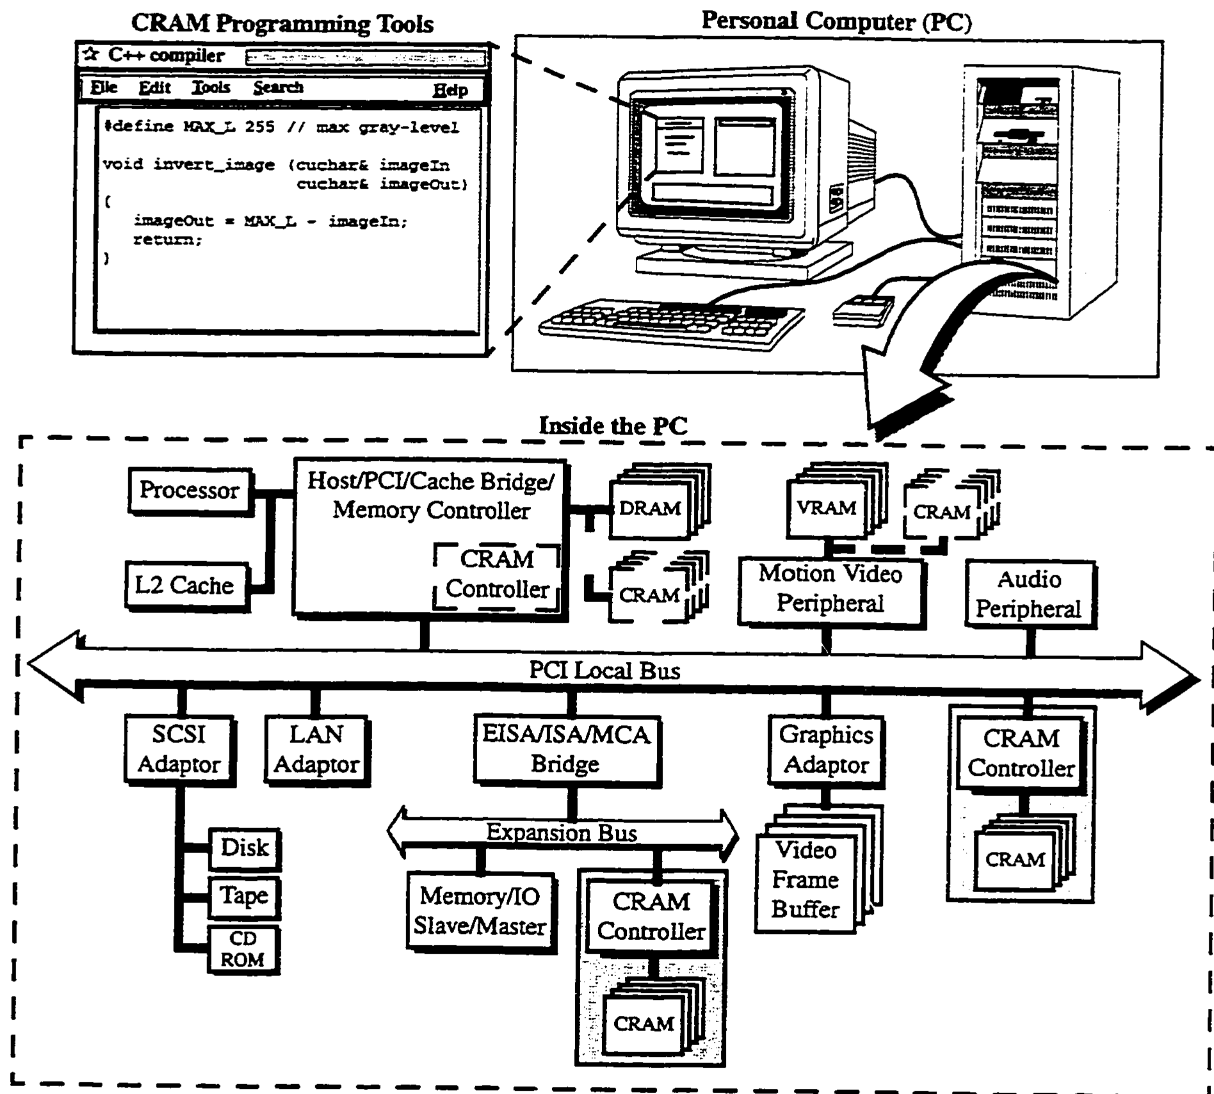

Figure 1.3 CRAM in a Typical Computer System

1.4 Thesis Scope

The scope of this thesis encompasses the study of the system design challenges outlined in Section 1.3. The thesis argues and proves that even with the limitations imposed by the host computer, i.e. small bandwidth, bit-parallel data format, and constrained PCB size, it is still possible to build a minimal-hardware and general-purpose CRAM system that yields significant performance speedup over conventional uniprocessor systems when executing massively-parallel applications. The main system design issues covered in this work include the following:

- Design and implementation of a minimal-hardware and high-performance PE controller that also enhances the use of CRAM as a general-purpose parallel-processing system.
- Design and implementation of a CRAM-host interface that minimizes the effect of the host on the performance of a CRAM system. This allows the implementation of CRAM on a wide variety of platforms.
- Design and implementation of data transposition schemes that offer reasonable compromise between area/hardware and performance.
- Implementation of a CRAM system prototype to demonstrate a working model of the whole CRAM concept.
- Analysis of the general implementation and performance of CRAM applications.
- Analysis of the effect of different CRAM architectural features on the overall performance of a CRAM system.

A secondary objective of this thesis is to develop a set of high-level software tools that can be used for application programming, system simulation, architectural analysis, and other CRAM development work. Since the CRAM prototypes implemented so far are very small, almost all the analysis work reported in this thesis is based on simulations using the C++ CRAM system simulator developed in this work.

1.5 Thesis Contributions

This thesis contributes to the general system design, implementation, and analysis of a logic-in-memory parallel-processing machine. The major design contributions include:

- A novel constant broadcast unit that improves the performance of operations with constant values by more than 35%. It also simplifies the use of variable-size constants, and reduces the size of the control store, thus reducing the area of the CRAM controller.
- A FIFO-based instruction queue that improves the performance of short-sequence instructions by a factor as high as tenfold. It also allows the CRAM controller to approach an ideal PE controller (100% PE utilization) at a small number of microinstructions per instruction. This translates in increased performance for a wider range of applications.
- The use of read/write buffers for data transfers from the host to CRAM that simplifies synchronization, eases electrical and physical loading of the host bus, and may reduce the time of loading data onto CRAM by as much as 80%.
- A new parallel array-based data transposition approach that is more than a hundred times faster than host-based software transposition, and contributes less than 10% of the total I/O overhead for CRAM systems with more than 4K PEs.
- A simple but innovative approach of grouping microroutines that reduces the required size of the microprogram memory by more than 50%. This small size (less than 256 32-bit words) makes an on-chip control store feasible, even in standard ASIC technologies and FPGAs, and hence reduces the number of components on a CRAM PCB.
- A general-purpose architecture for the controller that enhances the use of CRAM as a general-purpose parallel-processing system.

The major implementation contributions include:

- A CRAM C++ library that is used to write CRAM application programs using the standard C++ language and standard C++ compilers.
- A CRAM C++ simulator that can be used by both hardware and software designers to analyze CRAM architectural features as well as the performance of applications.
- Implementation of two controller prototypes and an ISA-based 64-PE CRAM system prototype to demonstrate working models of the CRAM concept.

In terms of system analysis,

- this work demonstrates that using the design features described above minimizes the effect of the host computer on the performance of a CRAM system. This is important for CRAM as a general-purpose system because it means that CRAM can be implemented on a variety of platforms, including slow host systems such as ISA-based computers and embedded systems that use slow microcontrollers.
- It is also shown that even with a bandwidth-limited host, a CRAM system still yields reasonably high performance for a variety of massively-parallel applications because of the performance-enhancement features of the CRAM controller.
- Using practical applications, this work highlights architectural bottlenecks of CRAM, especially the inter-PE communication network.

Finally, the thesis proposes two new CRAM ideas: an on-chip CRAM controller and a MIMD-SIMD CRAM system, and a pipelined constant-sensitive CRAM PE.

1.6 Thesis Organization

This chapter has outlined the background, motivation, scope and contributions of this thesis. Chapter 2 outlines the reasons behind logic-in-memory systems and describes work related to CRAM. Chapter 3 summarizes the architectural details of CRAM and lists CRAM prototype chips implemented so far. Chapter 4 to Chapter 7 discuss the primary contributions of this thesis. Chapter 4 describes the architecture of the CRAM controller and its interface to the host processor. Chapter 5 discusses the design, implementation, and testing of CRAM prototypes. Chapter 6 describes the different CRAM system software tools that have been developed in this work. Applications and performance analysis of a CRAM system are described in Chapter 7. Chapter 8 summarizes major accomplishments and suggests ideas for future work.

Appendix A provides architectural and implementation details of the CRAM controller. Appendix B provides PCB layout and pinouts of the CRAM system prototypes. Details of CRAM software tools are described in Appendix C. Appendix D lists the C++ source code for all applications described in this thesis.

Chapter 2

Logic-In-Memory Systems

This chapter discusses the background of logic-in-memory systems, and gives examples of CRAM-related work. Section 2.1 discusses the rationale for logic-in-memory systems and outlines the major advantages of these systems. Section 2.2 presents some of the more advanced and well-known logic-in-memory systems, with emphasis on the architecture and size of the processing elements (PEs), as well as system design issues (such as control of the PEs and the interface to the host). Since most logic-in-memory systems are SIMD systems, a brief review of other common SIMD PE control strategies is also given in Section 2.3. The architecture of CRAM itself is presented in Chapter 3.

2.1 Introduction

The idea of logic-in-memory has been around for years. One of the earliest documented cases is the logic-in-memory computer proposed by Harold Stone [1]. Stone argued that since the cost of components was to become heavily dependent on the number of package pins and not on the chip gate count, it was logical to equip the memory with some processing power so that operations can be performed directly in memory and take advantage of the inherent parallelism. This enhanced logic-in-memory cache would then be used as a high-speed buffer between the CPU and the main memory. It can therefore be seen that much of the initial motivation for a logic-in-memory system was to bridge the speed gap between a CPU and its memory. This holds true even today.

The main reason for the differences in speed between the CPU (microprocessor) and its memory (DRAM) has been due to the fact that different fabrication processes are used for the two. Microprocessor fabrication processes are usually optimized for fast transistors and have many metal layers. This results in fast logic, accelerates communication, and simplifies power distribution. On the other hand, DRAM processes have more polysilicon layers in order to have small memory cells. They also have high threshold voltages and thicker oxides in order to reduce leakage current. This increases the density of the DRAM, and also reduces its refresh rate. Because of these differences in processes, microprocessor performance has been improving at a rate of 60% per year, while DRAM access time has been improving at less than 10% a year [2]. This widening performance gap between the processor and the memory is now the major obstacle to improved computer system performance. Using sophisticated multi-level caches has only partially solved the problem since these caches increase memory latency and also do little to increase memory bandwidth because they (SRAMs) are traditionally constrained by their narrow data bus widths. SRAMs are also very expensive when compared to DRAMs.

One solution to the processor-memory performance gap is to combine logic and memory on the same chip. The complexity of the logic that is merged into memory depends on the desired system performance and cost. On one end, a small number of complex full-featured processors or CPUs are integrated into memory. This allows the logic-in-memory system to perform all computations done by a standard microprocessor.

However, this type of system is expensive because of the size of the processors, and it is also difficult to implement a complex logic system in a DRAM process. The other type of logic-in-memory system, such as Computational RAM (CRAM) [3], integrates very simple single-bit processing elements (PEs) at the sense amplifiers of a standard RAM. The PEs add about 5-10% to the area of the RAM, and because of their simplicity, it is easier to implement them even in a DRAM process. Also, since the PEs are small, a big number of them can be integrated in a RAM chip, thus increasing the degree of computational parallelism. The following are advantages of logic-in-memory systems, with particular emphasis on CRAM-type systems:

- **High Memory bandwidth:** Typically, the data available at the RAM sense amplifiers is more than a 1000 times the width of the RAM external bus (which is typically 1, 8, 16, or 32 bits wide). Therefore, processing elements integrated at the sense amplifiers can utilize this high bandwidth to speed up the execution of parallel applications. Figure 2.1 shows the bandwidths available at different points in a computer system. This comparison [8] is based on a system with 256 MBytes of 16 Mb, 50ns DRAM chips, and a 100 MHz CPU with a 64-bit bus.

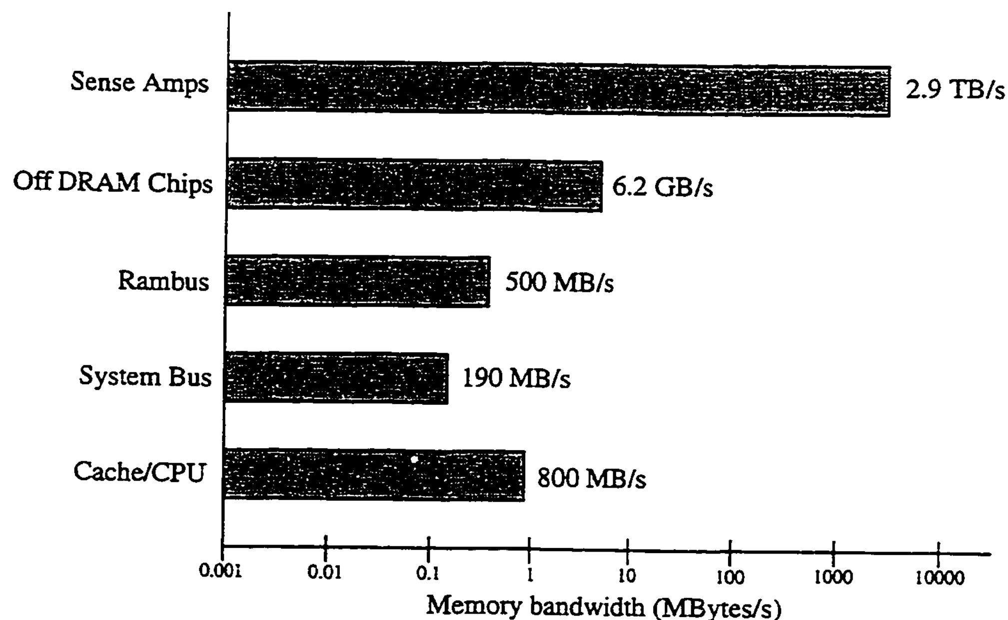

Figure 2.1 Memory Bandwidth in a Computer System

- **Lower Power Consumption:** Logic-in-memory systems improve power consumption. First, since most of the computations are done on the memory chip, there are fewer external memory accesses which would otherwise consume energy in driving the high-capacitance off-chip buses. Second, in a standard RAM chip, a single memory access drives several hundreds of bitlines even though a few tens of data bits are made available to the external RAM bus. This is an obvious wastage of energy. In a logic-in-memory system that integrates PEs at the RAM sense amplifiers, almost all the data bits driven onto the sense amplifiers are used in the computation, thus improving the efficiency of power consumption. For example, in a comparison done in [9], a 200 MHz Pentium accessing its cache transfers data over 10 cm of 32-bit buses. Assuming a 20 pF capacitance per data bit, and a 3.3 V power supply, this data transfer results in energy consumption of about 110 pJ per data bit. On the other hand, the PE's of an equivalent CRAM system with 16 x 1024b CRAM chips and a 50 ns cycle time, access their data over 5 mm metal lines with a total of about 1 pF per line. This results in power consumption of about 5.5 pJ.
- **Reduced Board Size:** A logic-in-memory system may result in fewer chips than a system implemented with discrete RAM and CPU chips. Also, for systems that require modest computation power but whose board area is precious, such as in portable electronic devices, a single chip with merged logic and memory is more attractive.

The CRAM project is one of the pioneering logic-in-memory research efforts. However, in the past five years, a number of universities and companies have pursued this idea. In this chapter, we present some of the more advanced and well-known logic-in-memory systems. This provides a basis of comparison with CRAM, and also gives an overview of the direction of research in such systems. Since most logic-in-memory systems are SIMD systems, a brief review of other common SIMD controllers is also given. The architecture of CRAM itself is presented in Chapter 3.

2.2 Logic-in-Memory Systems

2.2.1 Integrated Memory Array Processor (IMAP)

The integrated Memory Array Processor (IMAP) chip [10] integrates 64 processing elements (PE) with a 2-Mb SRAM. A 7000-transistor PE consists of an 8-bit ALU/Shifter, 14 8-bit registers, 4 1-bit registers, and a 4-bit exchange unit. Each PE has 4K x 8-bit of memory. IMAP was designed by NEC specifically for image processing, hence the 8-bit PEs. It was fabricated in a 0.55 μm BiCMOS double layer metal technology and contains 11 million transistors in a 15.1 x 15.6 mm^2 die area. The chip operates at 40 MHz and has a peak processing performance of 2.56 GIPS and a peak memory bandwidth of 1.28 GBytes/s. Figure 2.2 shows the configuration of the IMAP chip.

As an improvement to IMAP, NEC developed the Parallel Image-Processing RAM (PIP-RAM) [12]. This logic-in-memory image processor integrates 128 PE's and 16 Mb DRAM in a 64 Mb DRAM process technology. As in IMAP, the PE is 8-bits and consists mainly of an 8-bit ALU, a shifter, 24 general-purpose registers, and 5 special-purpose registers. Each PE has 128 Kb of DRAM. The PIP-RAM chip has a die size of 18.8 x 16.7 mm^2 and operates at 30 MHz. It has a peak processing performance of 7.68 GIPS and a peak memory bandwidth of 3.84 GB/s.

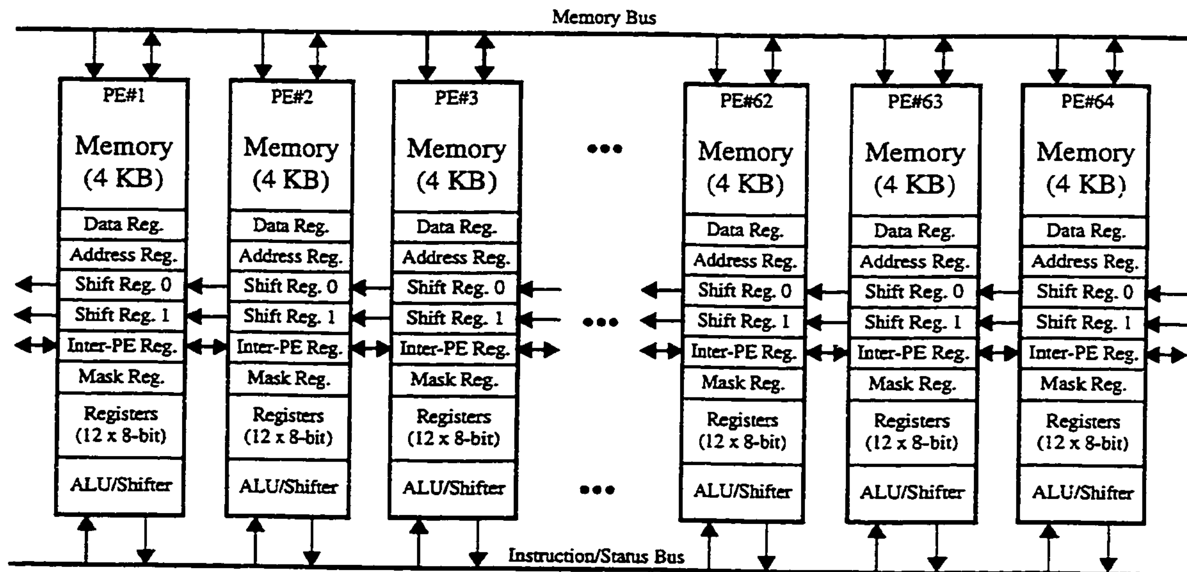

Figure 2.2 Integrated Memory Array Processor Configuration

A real-time vision system (RVS-2) [11] was designed using eight IMAP chips and a custom controller, and it interfaces to a host workstation using the VME bus. Figure 2.3 shows the block diagram of the RVS-2 system. It consists of an IMAP board, a video board, and a host workstation board. The IMAP board consists of eight IMAP LSI's, a controller LSI (RVSC), program memory, data memory, and VME interface.

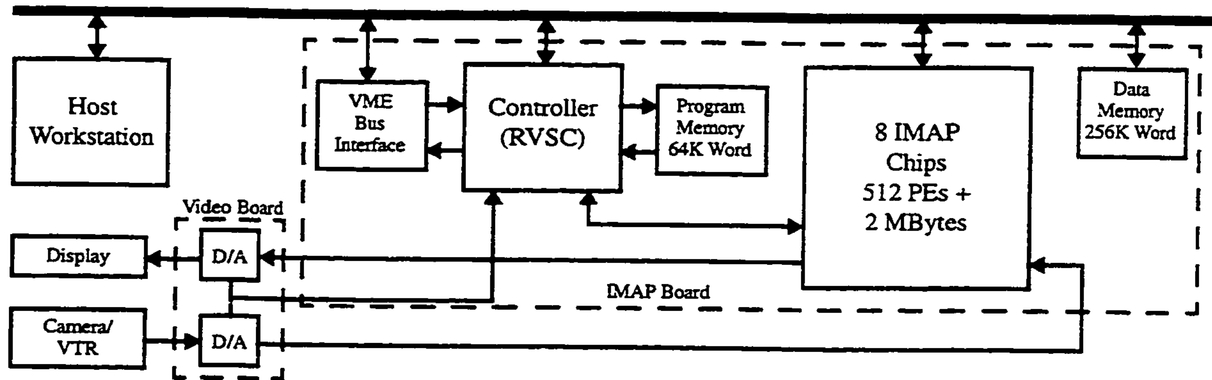

Figure 2.3 RVS-2 Configuration

Figure 2.4 shows the block diagram of the controller. The main functions of the controller are to sequence program execution, arbitrate memory accesses from the host, the PE array, and the controller itself, and to process global data. In order to meet the last requirement, the RVSC has a 16-bit processor, a special function to execute CAM-like functions, and selective access to data or status from an arbitrary PE.

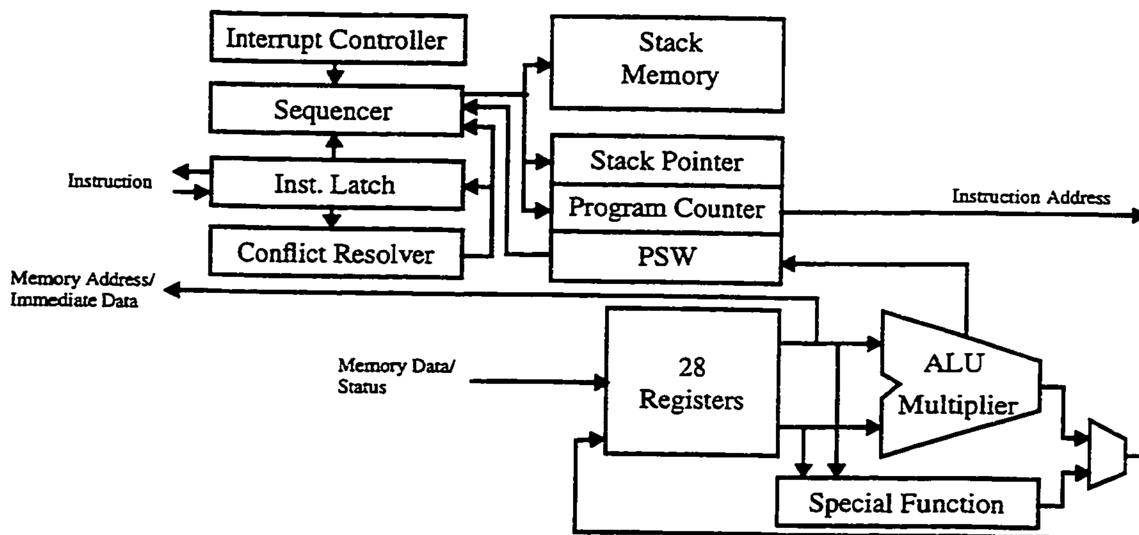

Figure 2.4 RVS-2 Controller (RVSC)

2.2.2 Terasys Processor In Memory (PIM) Array

The Terasys PIM system [13] was designed by researchers at the Supercomputing Research Center (SRC). It contains 2K x 64 bits of SRAM with 64 bit-serial processors. Processors are connected in a linear network and can communicate using a global-OR, partitioned OR and a parallel prefix network. A PIM array unit contains 64 PIM chips (4K processors) spread over two boards. A half dozen Terasys workstations were built at SRC. Each workstation consists of a Sun Sparc-2 workstation, an SBus interface card residing in the sparc cabinet, a Terasys interface board, and 8 PIM array units (32K processors). At an instruction issue rate of 100 ns, the system delivers 3.2×10^{11} peak bit operations per second. Figure 2.5 shows the organization of a 16K processor Terasys workstation. Applications on Terasys are programmed in a custom-designed language called dbc (data bit C), which is a superset of ANSI C.

The interface board is the controller for the PIM array and memory. A read and write operation to PIM memory are sent to the interface board which decodes whether it is a normal memory operation or a PIM array command. Each command is 12-bit wide, and indexes into a 4K lookup table. The lookup table contains 25-bit microcode instructions. These microinstructions are generated by a microcode assembler on a per-program basis, and are loaded into the controller board when a Terasys program is initiated. The interface board also provides PIM timers, selects one of two PIM commands based on the current value of the global OR signal, and registers 31 bits of global OR history.

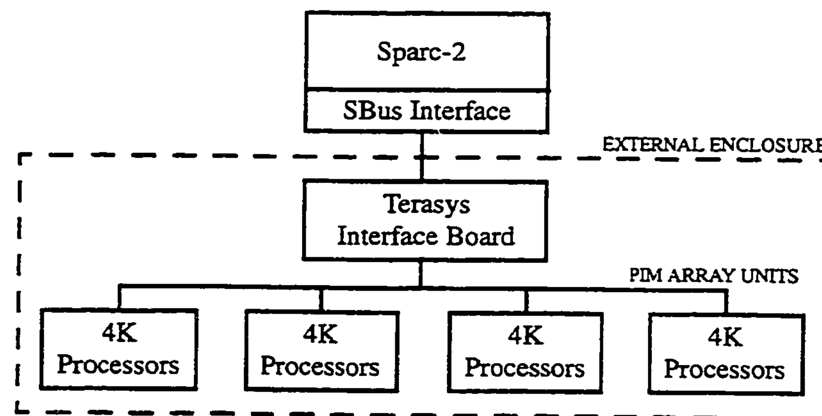

Figure 2.5 A 16K Processor Terasys Workstation

2.2.3 MIT Pixel-Parallel Image Processing System

The MIT image processing system [14] (Figure 2.6) integrates 64 x 64 bit-serial processing elements with 128 bits/PE of DRAM in a 0.6 μm HP CMOS14TB technology. The PE's use a 256 function generator and are connected to its four nearest (South, East, North, and West) neighbors. This system was designed purely for pixel-processing, and hence most of its architectural features (such as data format converters, array controller, and pixel/PE array configuration) were specifically tailored for 8-bit gray scale pixels. The interface to the host computer is through the VME bus. System performance, tested on a limited number of pixel-processing applications (Smoothing and Segmentation, Median Filtering, and Optical Flow) exceeds thirty frames per second.

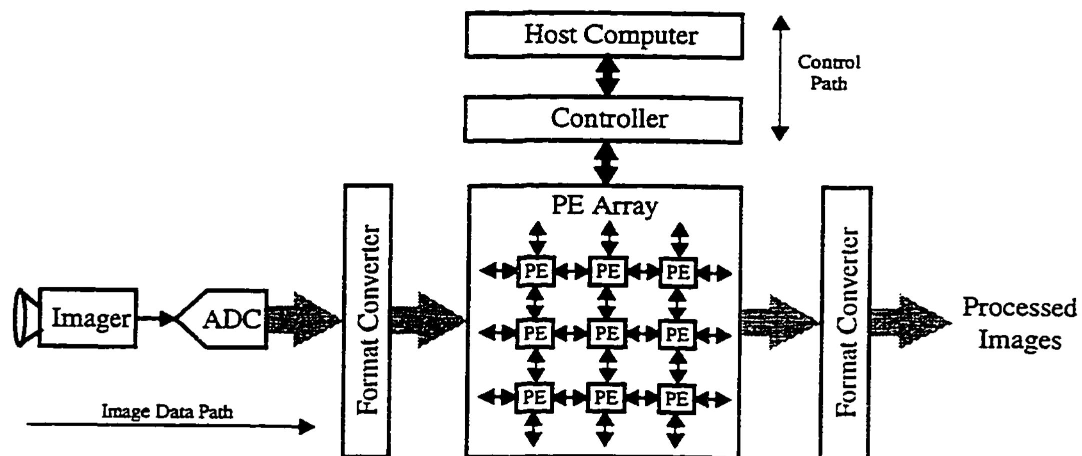

Figure 2.6 MIT Image Processing System

Figure 2.7 shows the controller architecture. Sequences of microinstructions are generated by the host and are stored in the control store. The sequences perform basic arithmetic, comparison, and data movement operations. To initiate a sequence of microinstructions, the host computer writes the starting address of the sequence into the opcode register. The sequencer steps through the control store, producing one array instruction every clock cycle (100 ns).

The select register and the associated multiplexer are used for operations with scalar variables. Figure 2.8 shows how instruction selection is employed to add 4-bit scalar variable, *b*, to a parallel variable, *A*. Before processing begins, pairs of instruction

sequences are stored in the controller. During program execution, the value of b is loaded into the select register, which is a parallel-in, serial-out shift register. As the sequencer steps through the control store, the shift register output selects one array instruction from each microinstruction pair.

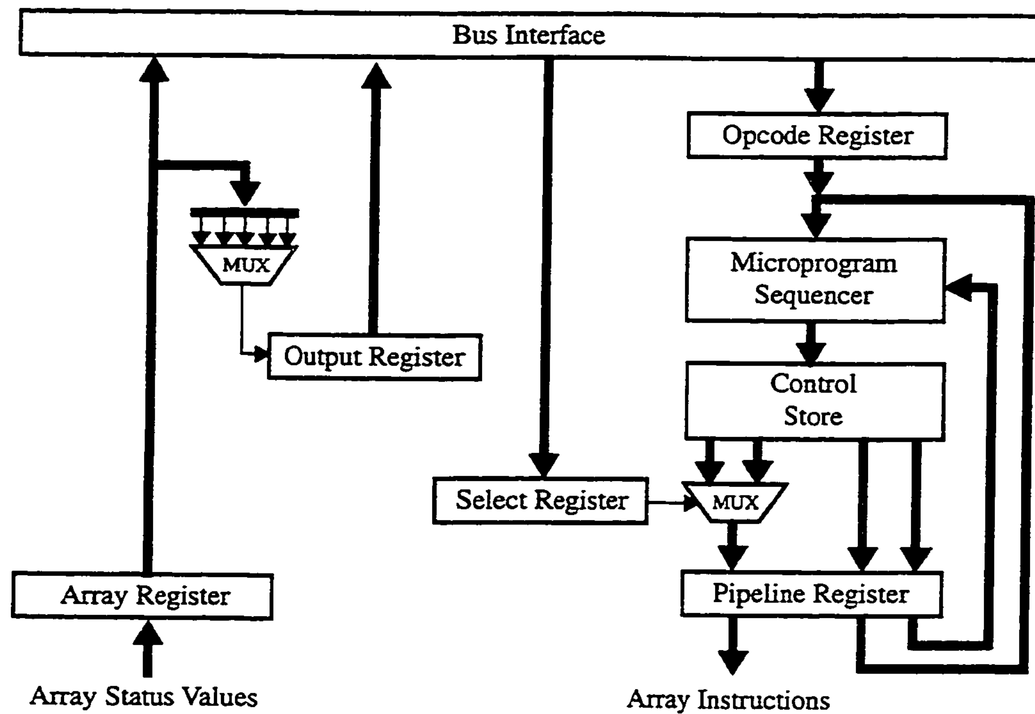

Figure 2.7 MIT Controller Architecture

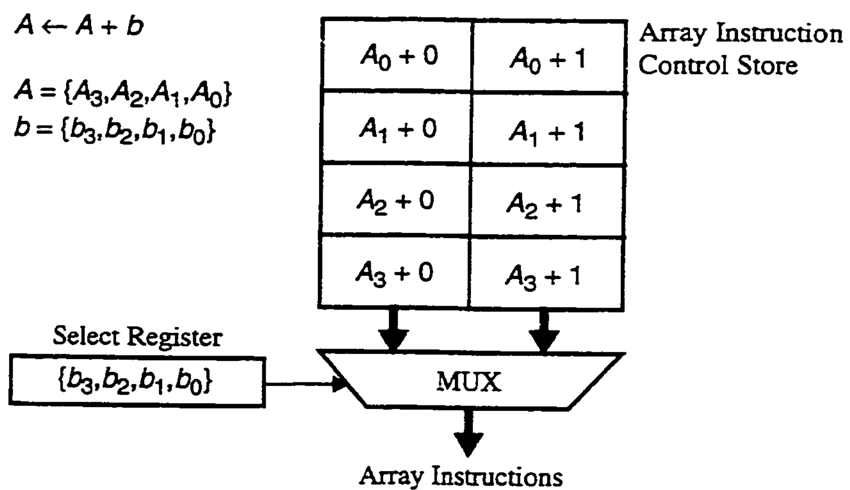

Figure 2.8 Instruction Selection

2.2.4 Intelligent RAM (IRAM)

IRAM [15] is another CRAM-related research effort that aims at merging processing and memory into a single chip to lower memory latency, increase memory bandwidth, and improve energy efficiency. This research is conducted by Patterson's group at the University of California at Berkeley. Though they have not yet built any prototype chips, they have proposed an IRAM vector processor shown in Figure 2.9. The processor includes sixteen 1024-bit-wide ports on the IRAM, thirty-two 64-element vector registers, pipelined vector units for floating-point add, multiply and divide, integer operations, load/store, and multiplication operations. It is projected that in a 0.18 μm DRAM process with a 600 mm^2 chip, a high performance IRAM-based vector accelerator would have eight add-multiply units running at 1000 MHz and sixteen 1-Kbit buses running at 50 MHz, resulting in system performance of 16 Gflops and 100 GB/s.

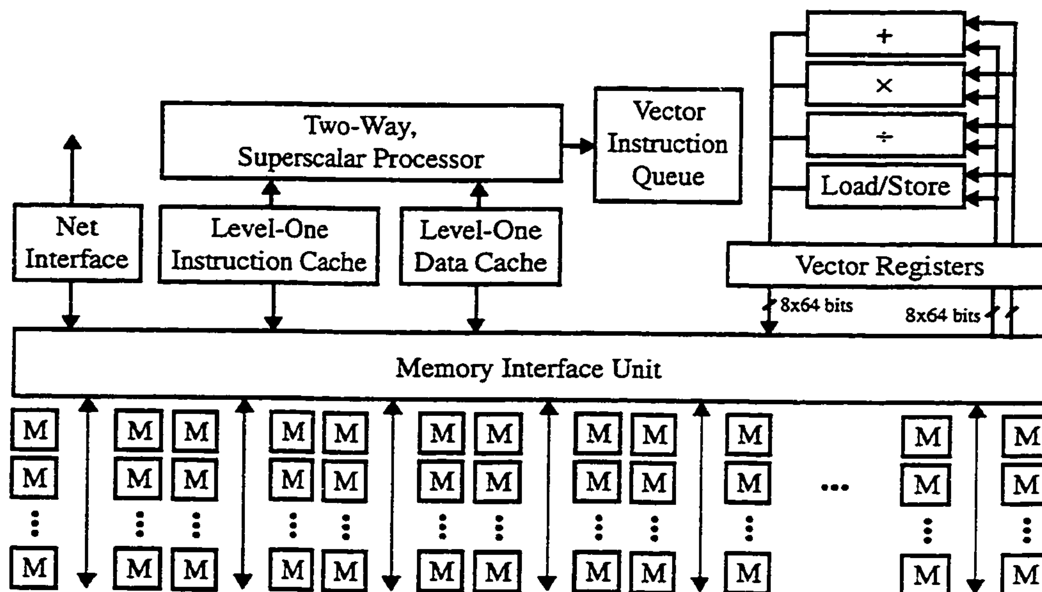

Figure 2.9 Organization of an IRAM Vector Processor

2.2.5 An Integrated Graphics Accelerator and Frame Buffer

This is a commercial chip [18] designed by Mosaid and Accelerix using some ideas from early CRAM research work [3]. The 33 GB/s, 13.4 Mb chip integrates parts of the graphics processor with DRAM. Other on-chip units include a PCI interface, VGA core

and video input block, and custom-designed pixel output path (POP) logic. The 4K pixel-processing units (PPUs) are each pitch-matched to 4 DRAM columns. The chip is implemented in a 0.35 μm/0.5 μm blended DRAM and logic process. Because of a Non-Disclosure Agreement, the architecture of the chip will not be discussed.

2.2.6 A Memory-Based Parallel Processor for Vector Quantization

The Memory-Based Parallel Processor for Vector Quantization (FMPP-VQ64) [16], [17] is a memory-based parallel processor containing 64 PEs. It is designed to accelerate nearest neighbor search (NNS) in Vector Quantization (VQ). Almost all its features are specifically tuned and fixed for VQ processing. Each PE has sixteen 8-bit SRAM words that are used to store sixteen codebook vectors. The PE ALU can only compute the absolute distance between an input and a codebook vector. The ALU is 12-bits because the distance between an input vector and a codebook vector may only grow up to a 12-bit word. The FMPP-VQ64 was fabricated in a 0.7 μm CMOS process and has a die size of 55 mm^2. It operates at 25 MHz and dissipates 20 mW of power at 3.0 V. It can perform 53,000 nearest neighbor searches per second. Figure 2.10 shows the FMPP-V64 block diagram.

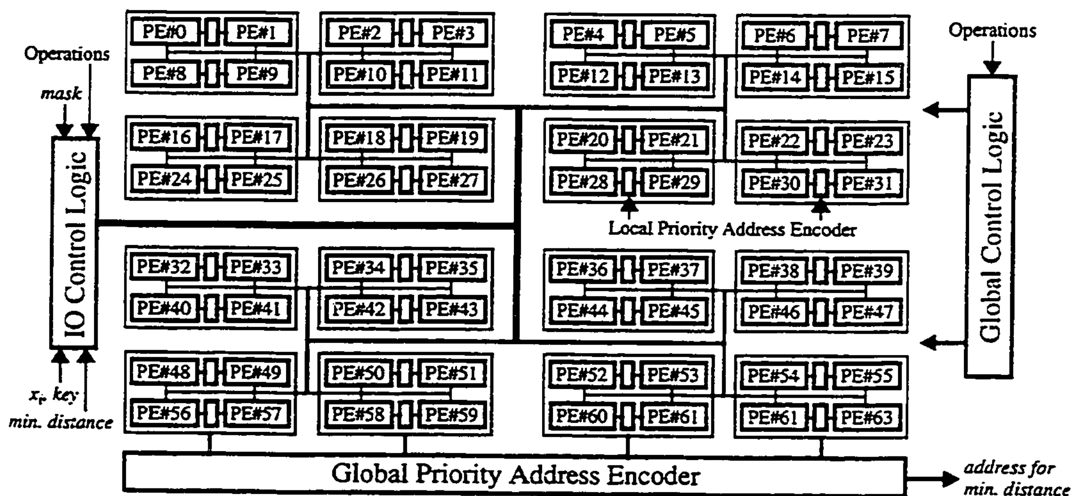

Figure 2.10 Block Diagram of the FMPP-VQ64

2.3 Other Common SIMD Control Path Strategies

2.3.1 Hierarchical Controllers

Hierarchical controllers are used in SIMD machines to meet the high bandwidth of control required by the PE array. Examples of such machines include VASTOR [26], Associative String Processor (ASP) [27], and Massively Parallel Processor (MPP) [28]. Figure 2.11 show the controller hierarchy in VASTOR. The microcontroller executes sequences of microcode stored in an internal read-only memory. The starting address for a specific microcode sequence is loaded from the buffer memory. The buffer memory is divided into 16 separate task control blocks that are filled by the microprocessor. When ready, the microcontroller requests the address of the next control block by interrupting the microprocessor. The microprocessor further reduces the control bandwidth by translating high-level operations from the host computer into sequences of microcontroller tasks. It also handles storage management.

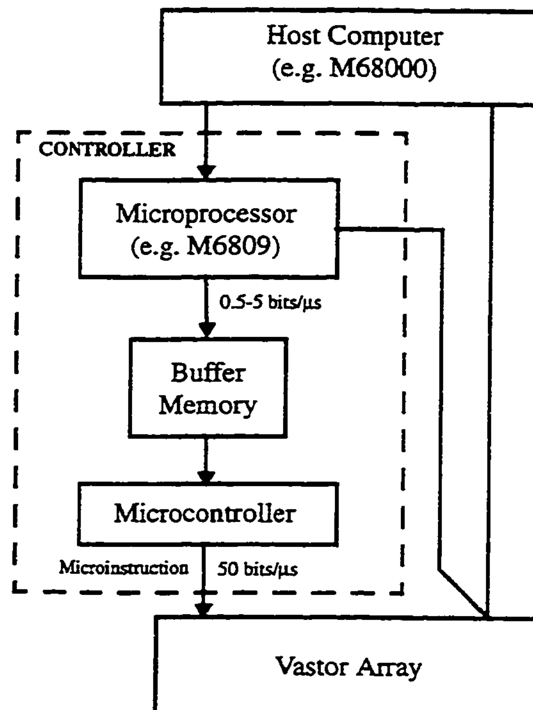

Figure 2.11 VASTOR Controller Hierarchy

2.3.2 Standard Microprogram Sequencers

Standard microsequencers, especially the Am2910 [30], have also been used as PE and memory controllers of SIMD machines. In this case, the microsequencer is the main functional block of a non-hierarchical control unit, with a few other components (such as an address processor) added to compliment its functions. Figure 2.12 shows the block diagram of the controller used in LUCAS Associate Array Processor [31]. An instruction is loaded from the master processor (in this case a Z80 microcomputer) into the instruction register (IR). This is then used by the Am2910 microsequencer to step through a 4K word microprogram memory. The address processor computes (increment, decrement, add constant, compare, etc.) the address to the bit-slice memory depending on the values in the parameter registers PR1-PR4.

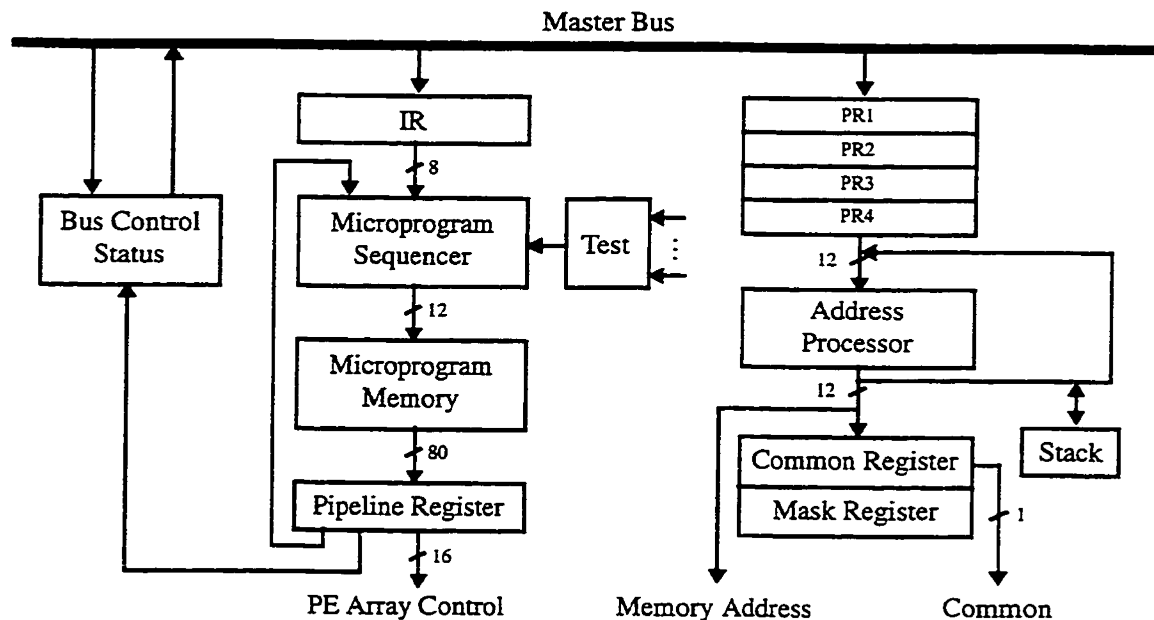

Figure 2.12 LUCAS Control Unit

2.3.3 High-Performance Processors

Some SIMD machines employ either standard or custom-built high-performance processors to control the PEs, as well as perform independent program execution. For example, the MasPar MP-1 array control unit (ACU) [32] is a 14 MIPS scalar processor with a RISC-like instruction set and a demand-paged instruction memory. The ACU

fetches and decodes MP-1 instructions, computes addresses and scalar data values, issues control signals to the PE array, and monitors the status of the PE array. It is implemented with a microcoded engine and occupies one PCB.

2.4 Summary

In this chapter, the motivation behind logic-in-memory systems has been presented. A number of CRAM-related systems have been described to indicate the state of research in this area. Emphasis has been on system design issues, especially controller design strategies. This forms the basis of comparison for the controller and system design work presented in this thesis.

Most of the logic-in-memory systems described in this chapter are designed and implemented for specific applications, especially image processing and graphics. CRAM, on the other hand, is meant to be a general-purpose parallel-processing system targeted for a variety of applications and implemented on a number of standard platforms. The PC has been chosen as the major CRAM implementation platform because of its wide usage. This thesis describes the system design issues for a general-purpose logic-in-memory CRAM system. First, a brief overview of the architecture of CRAM is described in Chapter 3. Thereafter, the system design issues are presented. This includes the PE controller, the interface to the host computer, controller and system prototypes, programming and other software tools, applications, and performance analysis.

Chapter 3

Computational RAM

This chapter describes the architecture of Computational RAM (CRAM), and lists CRAM prototype chips implemented so far. Section 3.1 describes how logic is added to a standard RAM to form CRAM. It also gives details about the CRAM processing element (PE) array, including PE architecture, global and inter-PE communication networks, and how SIMD computations are performed on CRAM. Section 3.2 gives brief descriptions of prototype chips that have been implemented since the CRAM project began. One of the prototype chips, the C64p1k, has been described in more detail because it forms the basis of most of the system prototyping and performance analysis presented in this thesis.

3.1 Architecture

3.1.1 RAM with SIMD Processors

Computational RAM (CRAM) is a SIMD-memory hybrid architecture [19], [20], with single-bit processing elements (PEs) integrated at the sense amplifiers of a standard RAM. CRAM is designed to improve the speed of executing massively-parallel applications by utilizing the high bandwidth available at the sense amplifiers. Typically, the data width at the sense amplifiers is more than 1000 times the width of the RAM external bus (which is usually 1-, 8-, 16- or 32-bit wide). Several PE's can therefore have access to this data and operate on it without the need to read the data out of the RAM chip and transmit it over long high-capacitance buses to the processor. This improves performance and also reduces power consumption. Figure 3.1 shows the architecture of CRAM. The PE is laid out in the pitch of a few sense amplifiers (typically four or eight for DRAM, and one for SRAM).

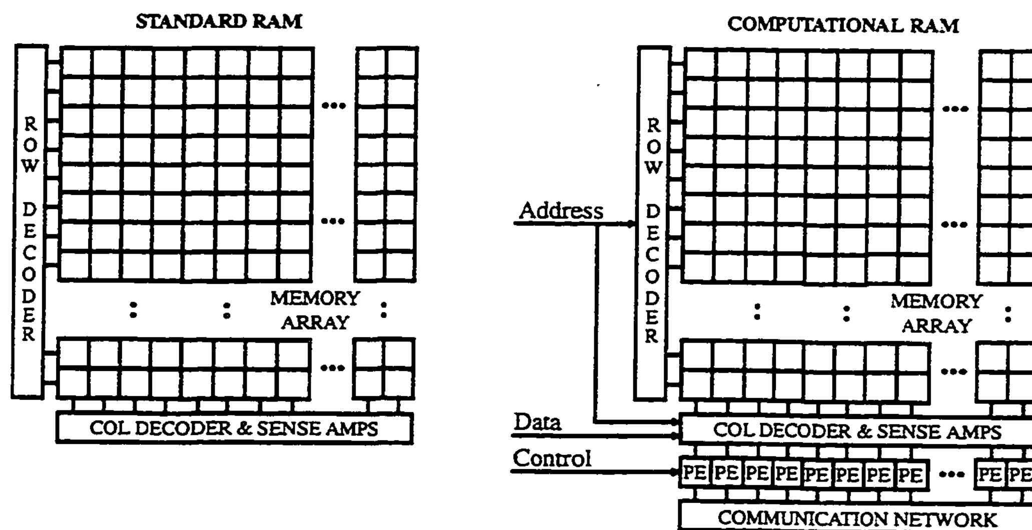

Figure 3.1 CRAM Architecture

3.1.2 Processing Element (PE)

A 1-bit PE architecture was chosen in order to achieve the highest performance per silicon area [21]. For applications with a high degree of parallelism, a trade-off can be made between higher processor complexity and a greater number of simple processors. In CRAM, two types of processing elements have been implemented.

- **Baseline PE (BPE):** The baseline PE shown in Figure 3.2 is the simpler of the two and consists basically of three registers (W, X, and Y) and an 8-to-1 multiplexer (mux). A fourth source of operands is the memory (M). The broadcast bus implements a global-OR (or bus-tie) of all the PEs outputs. Communication between adjacent PEs is through a shift-left/shift-right network. A PE instruction consists of an 8-bit multiplexer truth-table opcode (TTOP) and a 6-bit control opcode (COP). TTOP determines the actual operation performed by the PE depending on the contents of Y, X and M (the multiplexer select inputs). Three bits of COP (WY, WX, and WW) determine whether the output of the PE mux should be written to the PE registers Y, X and W registers respectively. The next two COP bits (SLX and SRY) control the PE shift operation by enabling the write of a PE mux output into the X or Y register of its left or right neighbor respectively. The sixth bit of COP enables the global-OR of the PE outputs. Register W controls the ability of a PE to write to its own memory. This is used in conditional code execution. TTOP and COP are multiplexed with data and address on the RAM data and address buses, respectively. Each PE operates on a single element of a vector, one bit at a time. The baseline PE is the architecture used in the bit-serial CRAM prototype chips C64p128, C64p1K, and C1Kp16K (Section 3.2). The BPE consists of 75 transistors in dynamic logic.

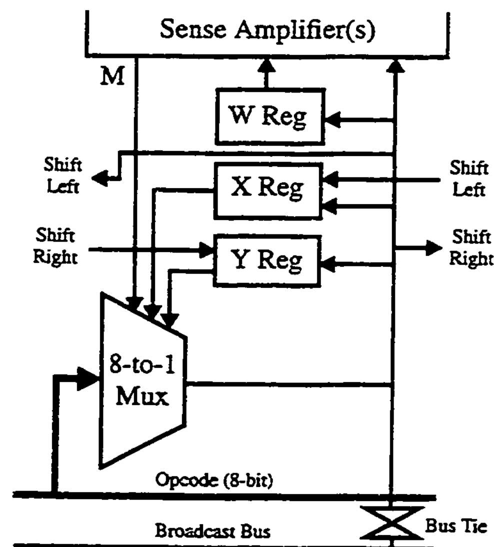

Figure 3.2 Baseline Processing Element

- **Extended Processing Element (XPE):** Shown in Figure 3.3, the extended PE is an enhancement of the baseline PE, with more registers to reduce the number of RAM cycles per computation. It also has a carry-chain to speed up addition, and has PE-chain segmentation features (such as the bus-tie segmentation register T, and the carry-in selection/segmentation registers S and B) to allow the grouping of adjacent PEs to work on multi-bit data. This data-parallel operation offers substantial improvement in the performance of the more complex arithmetic operations such as multiplication. The XPE can also be configured to function as a bit-serial baseline PE. This PE architecture has been used in the C512p512 CRAM chip and has 147 transistors.

Figure 3.3 Extended Processing Element

3.1.3 CRAM Computations

The PE ALU (multiplexer) computes any function of 3-bit variables Y, X, and M. TTOP defines the operation, with Y:X:M selecting which bit of TTOP should be the output of the ALU. Figure 3.4 shows the values that have to be set on TTOP in order for the multiplexer to compute the functions O=M, O=M, $O=Y\oplus X\oplus M$ (addition) and O=YX+YM+XM (carry). The destination register of the operation is specified by setting

the corresponding bit of COP to 1. To illustrate how each PE does the computation, consider 4-bit addition. Figure 3.5 shows how the data (operands) for the addition are arranged in each PE memory. The addition is then performed one bit at a time using the algorithm shown in the figure. Bit 0 (a_0, b_0, r_0) is the least significant bit. The looping and the setting of the opcodes are performed by a controller external to the CRAM chip.

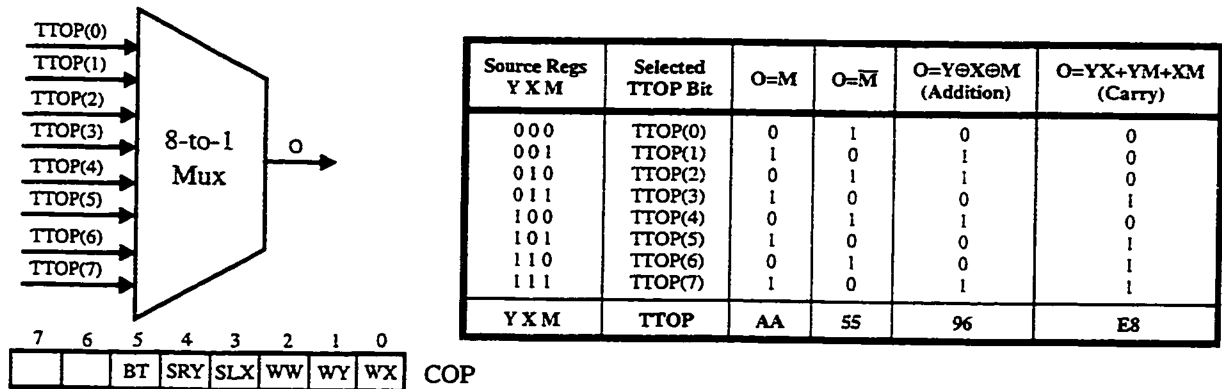

Source Regs Y X M	Selected TTOP Bit	O=M	O=$\overline{M}$	O=Y⊕X⊕M (Addition)	O=YX+YM+XM (Carry)
0 0 0	TTOP(0)	0	1	0	0
0 0 1	TTOP(1)	1	0	1	0
0 1 0	TTOP(2)	0	1	1	0
0 1 1	TTOP(3)	1	0	0	1
1 0 0	TTOP(4)	0	1	1	0
1 0 1	TTOP(5)	1	0	0	1
1 1 0	TTOP(6)	0	1	0	1
1 1 1	TTOP(7)	1	0	1	1
Y X M	TTOP	AA	55	96	E8

Figure 3.4 Setting PE Opcodes (TTOP and COP)

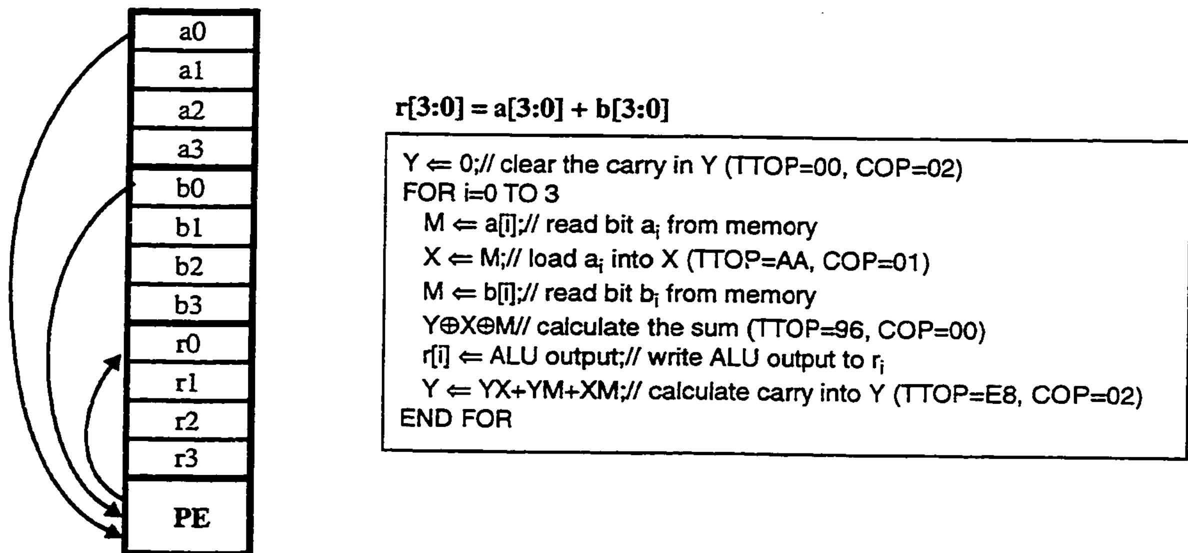

r[3:0] = a[3:0] + b[3:0]

```
Y ⇐ 0;// clear the carry in Y (TTOP=00, COP=02)
FOR i=0 TO 3
   M ⇐ a[i];// read bit a_i from memory
   X ⇐ M;// load a_i into X (TTOP=AA, COP=01)
   M ⇐ b[i];// read bit b_i from memory
   Y⊕X⊕M// calculate the sum (TTOP=96, COP=00)
   r[i] ⇐ ALU output;// write ALU output to r_i
   Y ⇐ YX+YM+XM;// calculate carry into Y (TTOP=E8, COP=02)
END FOR
```

Figure 3.5 PE 4-bit Addition

3.2 Prototypes

This section describes the CRAM prototype chips that have been implemented since the CRAM project began. These are very brief descriptions. Detailed descriptions can be found in the referenced theses of the students who implemented them. The prototype chip C64p1K has been described in slightly more detail because it is the chip that has been used in the CRAM prototype system designed and implemented in this thesis.

3.2.1 A 64-PE, 128 bits/PE CRAM (C64p128)

This CRAM was designed and implemented by Duncan Elliott [4]. It is an 8 Kbit CRAM, with 64 baseline PEs and 128 SRAM memory bits per PE. Each PE fits in the pitch of one sense amplifier. The PEs occupy only 9% of the total area, with the 6-transistor SRAM cells dominating the chip area. The PE registers are designed as 5-transistor static CMOS latches. A PE cycle has three phases: precharge, ALU (mux) operation, and write (registers and memory). The C64p128 was implemented in a 1.2 µm CMOS process and has a read-operate-write cycle time of 114 ns. With the exception of an error in one of the 64 column decoders, the chips are functional.

3.2.2 A 64-PE, 1 Kbits/PE CRAM (C64p1K)

The C64p1K was implemented by Christian Cojocaru [22]. It is basically an enhancement of the C64p128, with more memory (1 Kbits) per PE and implemented in a faster BNR/NT's 0.8 µm BiCMOS technology. BNR synchronous SRAM modules were used as the PEs memory in order to speed up the implementation. These single-port SRAM modules use a standard 6-transistor cell and can operate up to 180 MHz. The C64p1K has an 8-bit data bus and a 13-bit address bus.

C64p1K memory accesses are clocked by the memory clock (MCK) signal. On the positive edge of MCK, the RW and RAM signals (which shows whether the access is read/write, or is external or between the PE and its memory) are sampled. The memory address, and the write data, are also latched. The positive edge then triggers a self-timing circuitry that generates all the timing signals necessary to complete the memory cycle.

The PE operate cycle is controlled by six signals, namely:

1. ALUPRE : Precharges the PE dynamic logic prior to evaluation. It is active LOW.
2. TTOPCK : The active HIGH evaluation clock.
3. COPCK : When HIGH, the results of the evaluation are written to the PE registers.
4. BTENCK : When HIGH, it enables the bus-tie (global-OR) of all the PE outputs.
5. TTOPLD : Its positive edge latches TTOP from the data bus into internal registers.
6. COPLD : Its positive edge latches COP from the address bus into internal registers.

These signals can either be generated externally by a CRAM controller, or they can be generated internally by the CRAM chip when the external control pin (XCTRL) is set to zero. These two timing schemes are described below.

- **External Timing:** In this approach, all the six PE timing signals must be generated by the CRAM controller. While this slightly reduces the area of the CRAM, it greatly increases the complexity of the controller, especially since a high clock rate must be used to generate the intermediate phases of the PE cycle. Figure 3.6 shows the timing of these signals. The signals ALUPRE, TTOPCK, and COPCK must be mutually non-overlapping. Also, BTENCK can not overlap ALUPRE. Unfortunately, the external timing scheme is not functional on the C64p1K chip due to a suspected layout error.

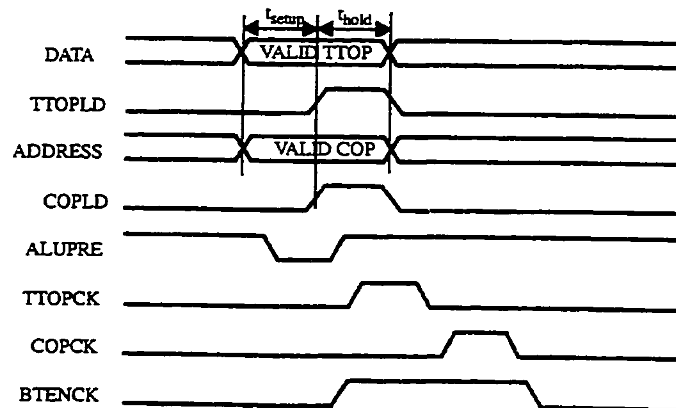

Figure 3.6 PE Cycle External Timing

- **Internal Timing:** For this scheme, the PE timing signals are generated on chip using a finite state machine (FSM) controlled by two external signals: OPS (Operate Strobe) and CCK (CRAM Clock). A PE cycle involving a bus-tie operation is one CCK period

longer than a standard cycle to allow for worst case bus-tie propagation delays. Figure 3.7 shows the timing for the two PE cycles. OPS positive edge signals the beginning of a PE cycle. Thereafter, transitions in the timing signals are controlled by both edges of CCK, which was designed to be a free-running clock. RESET is only asserted at power-up to initialize the FSM. As attractive as it is for the design of the CRAM controller, the internal timing scheme implemented on the C64p1K chip has two bugs. The first one, not evident at the time the chip was designed, concerns the timing of operate-write cycles. Since ALUPRE was designed to be normally-active (instead of the more standard and more logical normally-inactive scheme), and since the second positive edge of CCK re-activates ALUPRE after being deactivated by CCK first positive edge, the results of a PE operation can not be correctly written back to memory since the PE output will already have been precharged again when the memory-write begins. This bug was discovered during the design of the controller. To correct this for an operate-write cycles, the controller no longer issues a free-running CCK. But rather

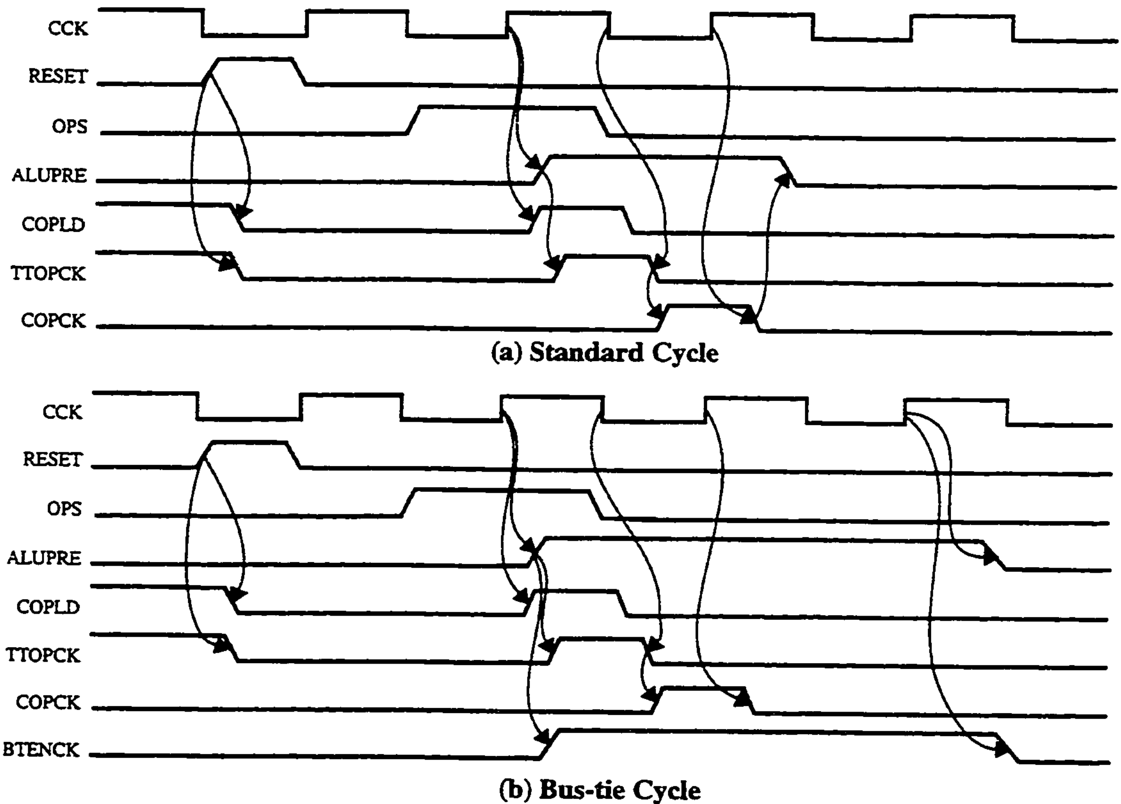

(a) Standard Cycle

(b) Bus-tie Cycle

Figure 3.7 PE Cycle Internal Timing

CCK is now a second control signal that is delayed (issued after MCK) if an internal memory cycle follows a PE-operate cycle. This correction has added to the complexity of the controller. The second bug, reported in the chip documentation [22], concerns the timing of the bus-tie operation. Even though the bus-tie is extended by one CCK period, the write-register (COPCK) signal is still asserted at exactly the same point as in a standard PE cycle. This means that the extra length of the bus-tie cycle is useless if the results of the bus-tie operation are to be written to a PE register (which is usually the case) instead of the PE memory. In this case, the extra CCK cycle simply complicates the timing sequencing, especially now that CCK is not free-running.

3.2.3 A 512-PE, 480 bits/PE CRAM (C512p480)

The C512p480 [22] is a bit-parallel/bit-serial CRAM based on the extended PE shown in Figure 3.3. Though initially designed to have 512 bits of memory per PE, it was implemented with 512 x 480 bits due to silicon area constraints. Like the C64p1K, the C512p480 was fabricated in the BNR/NT 0.8 μm BiCMOS technology. The C512p480 bit-parallel CRAM architecture simplifies system design and programming since the data doesn't have to be transposed from its bit-parallel format in the host computer. Also, the increased number of registers in the XPE reduces the number of memory accesses per computation. While organizing PEs in groups to work on multi-bit data in bit-parallel format speeds up the rate of computation for a single data item, it does not necessarily increase the overall performance since the degree of parallelism is reduced. Also, some algorithms, especially those requiring searches or comparisons on vector items, are difficult, if not slower, to implement on a bit-parallel architecture. Unfortunately, the C512p480 chips are not fully functional; only the RAM part is functional. It is suspected that there is a flaw in the PE timing circuitry or in the layout of the PEs themselves.

3.2.4 A 1024-PE, 16 Kbits/PE CRAM (C1Kp16K)

The design and implementation of the C1Kp16K CRAM chip is still in progress [9]. This is the first CRAM prototype that is based on DRAM. It integrates 1024 baseline PEs in a 16 Mb IBM DRAM [23]. Each PE in a DRAM octant is connected to 32 sense

amplifiers and has 16 Kbits of local memory. A vacant test mode available on the original DRAM is used for CRAM operation. Therefore, no extra control pins are required for CRAM, thus maintaining JEDEC compatibility [24].

3.3 Summary

This chapter has described the architectural details of Computational RAM (CRAM). CRAM integrates 1-bit processing element (PEs) at the sense amplifiers of a standard RAM. These PEs work in parallel in SIMD style, and improve the performance of massively-parallel applications by utilizing the inherent parallelism and high data bandwidth inside the memory chip. The PEs are 1-bit to reduce their area so that a number of them can be implemented in a single memory chip. This increases the degree of computational parallelism and performance.

Prior to this work, the CRAM project had concentrated on the design and implementation of individual CRAM prototype chips described in Section 3.2. The next obvious step was therefore to explore system design issues, and that is the focus of this thesis. Specific issues studied include the design of the PE controller, interface to the host computer, CRAM system software tools, prototyping, and system performance analysis. The next chapter describes the design of the controller and the interface to the host computer.

Chapter 4

CRAM Controller

This chapter describes the architectural details of the CRAM controller and how they affect the performance of a CRAM system. The requirements and characteristics of the controller are laid out in Section 4.1. Section 4.2 describes the interface of the controller to the host bus. Four host buses (PCI, VME, EISA, ISA) are studied. Section 4.3 describes the instruction queue unit and how it affects the execution time of CRAM instructions issued from the host computer. After this, the execution unit, which comprises of a microprogram sequencer and an address unit, is described. The use of read/write buffers to improve the performance of data transfers and operate-immediate instructions is discussed in Section 4.5. The last two sections describe the mode of access and memory-map of all user-accessible units on the CRAM controller. All performance results reported in this chapter are based on simulations of CRAM systems hosted on a 133 MHz Pentium PC.

4.1 Introduction

The main use of the CRAM controller is to allow CRAM application programs to be run from a front-end host computer. The controller acts as a PE array controller as well as an interface to the host bus. There are two reasons why the PEs have to be controlled from a dedicated controller. First, in order to minimize their area, CRAM PEs do not have control structures. Instead, they are designed such that control, data, and address information is broadcast to them in an SIMD style from an external controller. Second, since the PEs are bit-serial, a single standard operation such as addition takes several PE instructions to execute. Rather than issue such low-level (micro) instructions from the host processor, only high-level (macro) instructions are passed from the host to the controller, which in turn generates the PE microinstructions. This frees the system bus from congestion with the high traffic microinstructions, and also increases CRAM system performance since the host cannot issue PE microinstructions fast enough to correspond to the rate at which the PEs execute them.

As described in Chapter 2, a number of approaches have been used to implement controllers for SIMD machines in general, and logic-in memory systems in particular. However, the architecture and implementation of the CRAM controller [25] has been based mainly on the following requirements:

- **Simple and Minimal Hardware:** Typically, SIMD machines use high performance microprocessors as their controllers. For example, VASTOR uses an M6809 microprocessor and a microcontroller [26]. The Massively Parallel Processor (MPP) and the MasPar MP-1 use custom processors as their controllers [28], [29]. Even some logic-in-memory systems, such as the IMAP [11], use full scale processors as controllers. In CRAM, the need for very minimal hardware arises for three reasons. The first one is to reduce system cost. Secondly, unlike all related systems, CRAM is designed to be a future replacement of standard RAM as computer main memory. For this to be a realizable idea, both the CRAM and its controller must add minimal hardware to the existing memory system (i.e. to RAM and RAM controller). The third reason for a simple and small controller is that one of the future efforts suggested in this thesis is a CRAM system in which both the CRAM and its controller are implemented on the same chip

(see Chapter 8). Therefore, to maintain the small additional area of CRAM over standard RAM, the controller must use as few gates as possible. Also, since a combined RAM/logic system will generally mean logic being implemented in a slow RAM process [9], [18], a simple controller would yield higher speed.

- **High PE Utilization:** One of the underlying requirement of an SIMD PE controller is to make sure that the PEs are always kept busy. In other words, instructions to the PEs must be issued fast enough to match the rate at which the PEs execute them. To achieve this while maintaining a simple and minimal architecture, the CRAM controller has been equipped with performance-enhancement features such as a FIFO-based instruction queue unit, read/write buffers, and a new constant unit. Its sequencing features have also been trimmed to only those required to control the rather simple single-bit CRAM PEs.

- **General-Purpose Architecture:** In most logic-in-memory systems, the PE array and the controller are designed for specific applications. For example, the PE and the controller for IMAP[10], PIPRAM[12], and MIT-PP[14], are designed specifically for 8-bit pixel processing. FMPP-VQ[16] can only run Vector Quantization, while the logic-in-memory graphics accelerator and frame buffer [18] has a controller that is hard-wired for graphics operations only. On the other hand, the CRAM controller had to be designed to enhance the use of CRAM as a general-purpose parallel-processing system. For this reason, the controller supports operands of different word-lengths, and it has a microprogrammed architecture to allow different applications to be optimized in software (using the microcode development tools). Other architectural features, such as the organization of the data buffers and the isolation of the external bus interface unit from the controller core logic, allow the CRAM system to be easily implemented on different platforms (three system bus interface units have been designed).

- **Easy Programming:** No special programming language or instruction syntax is required for the controller. In fact, since the controller does not execute the standard processor instructions, but rather just sequences the PE microinstructions, the programmer need only worry about writing software for the variables in the CRAM mem-

ory. A CRAM C++ compiler (see Section 6.2) has been designed for this. The few instructions that need to be executed on the controller are limited to loading, incrementing and decrementing its registers. These are automatically handled by the CRAM C++ compiler, unless the programmer chooses to program in CRAM assembly code.

- **Constrained PCB Size:** Most SIMD machines occupy many PCBs housed in very large cabinets [13], [6], [73]. However, for a system to be implemented as an add-in card in a PC environment, it must be on a single standard-size board. Therefore, in order to fit the controller on the same small board as the CRAM chips it controls, it must use as few external chips as possible. Features of SIMD controllers that are commonly implemented with external chips include the microprogram control store, data-transposing hardware (sometimes referred to as data format converters), and program and data memory. The control store of the CRAM controller has been made small enough to be easily implemented on-chip. This has been achieved by storing microinstructions of basic operations only, and then using an innovative method of grouping microroutines that takes advantage of the unique architecture of the CRAM PEs and its multiplexed TTOP/Data bus to reduce the number of required microinstructions. Data-transposition is not done on the controller. Instead, for accesses involving a small number of elements, the conversion of data from the CRAM (vertical) format to the host format is done on the host. For large accesses, an array-based data-transposition, that takes advantage of the high degree of parallelism in the PE array, has been proposed (Section 6.3). No program or data memory is provided on the controller. All programs and data are store on the host computer.

Figure 4.1 shows the architecture of a CRAM controller that meets the above requirements. It is designed mainly as a PE microinstruction sequencing machine. All other sequential operations in an application program are executed on the host computer. The rest of the chapter describes the architectural features of the controller and how they affect the overall performance and design of a CRAM system.

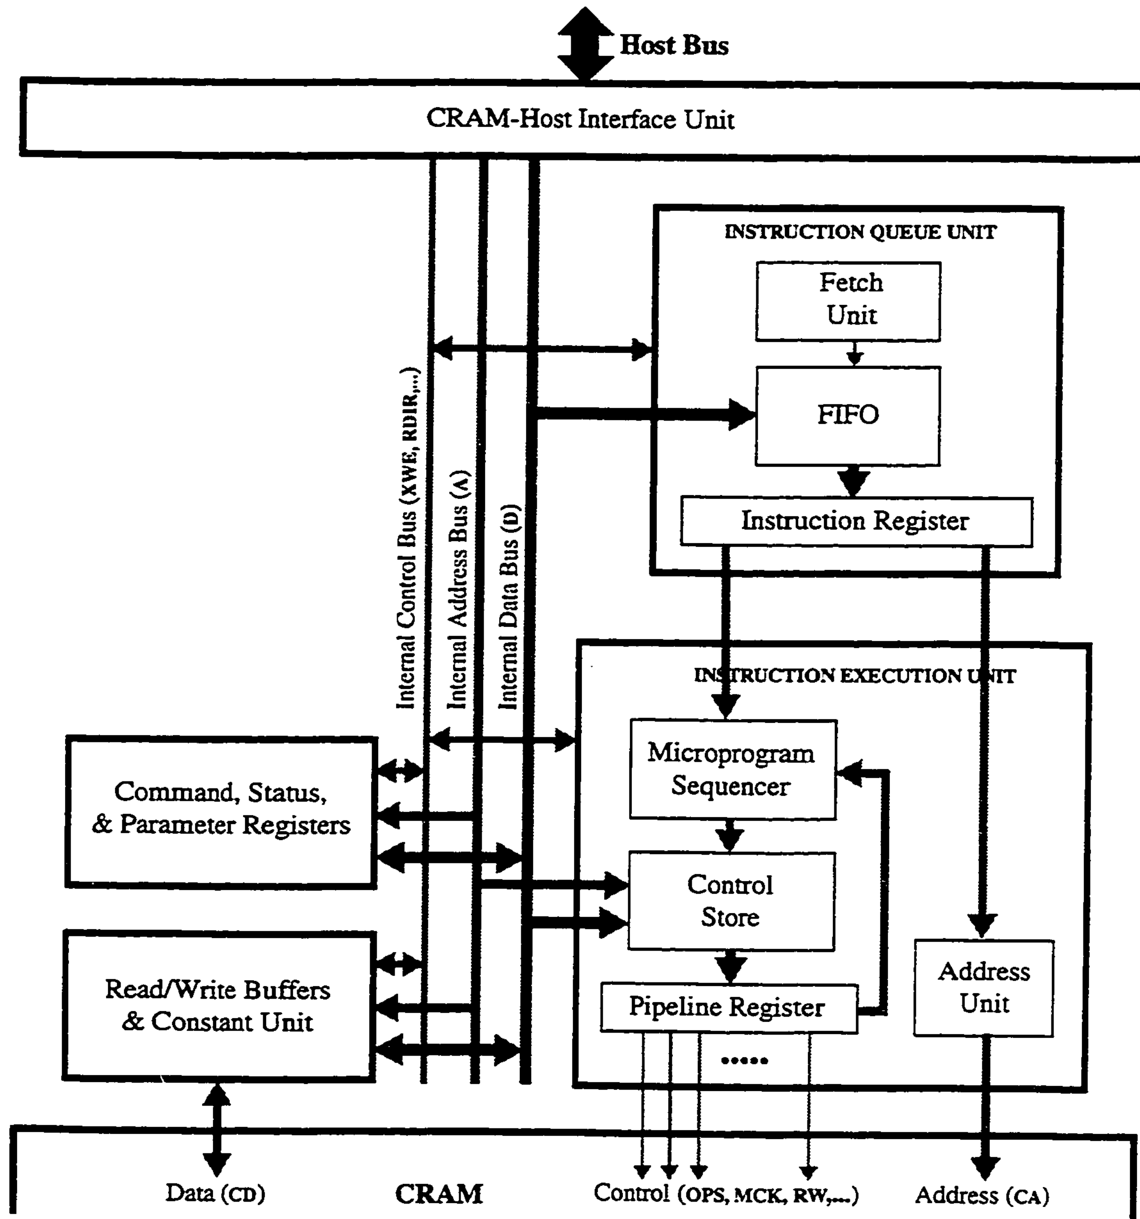

Figure 4.1 CRAM Controller Architecture

4.2 CRAM-Host Interface Unit

The CRAM-Host Interface unit connects the CRAM controller to the host processor. This allows the transfer of CRAM instructions and data from the host processor environment to the CRAM system, and the transfer of CRAM status and data to the host.

In this work, the major emphasis has been on the study of the effect of the CRAM-host interface on the performance, area, and design of a CRAM system, rather than on the implementation of a specific interface unit. For this reason, a number of host interfaces have been studied to allow a wide range of experimentation. For our main target environment, which is the Personal Computer (PC), we selected three buses: the PCI bus, the ISA bus, and the EISA bus. The PCI bus was chosen because of its high speed [34], and it is also becoming a standard in most PCs [35]. Its major disadvantage is that it has a very complicated bus protocol and requires significant configuration resources. In this work, we use the 33 MHz, 32-bit PCI target (slave) bus protocol compliant with the PCI Specification Revision 2.1 [36]. The ISA bus is the simplest of the PC expansion buses. While its data rate is low (16-bit at 8 MHz), unlike the PCI bus, its protocol [38] requires no configuration registers, it does not support burst cycles, and it has a very simple and slow timing scheme. The only configuration resources are jumpers for setting the address of the ISA card. The EISA bus [38], which is an extension of the ISA bus, offers a compromise between the ISA bus and the PCI bus. While the EISA bus runs at the same clock speed (8 MHz) as the ISA bus, the former supports burst transfers and has a 32-bit data bus allowing data transfers on the EISA to reach a maximum rate of 33M bytes/s [64]. Also, the EISA bus has a 32-bit address bus and has more advanced configuration resources than the ISA bus. This makes EISA devices easier to use in a plug-and-play system (automatic assignment of 32-bit address space) than ISA devices. It is not as widely used as the other two buses, however, the CRAM-EISA interface offers a speed improvement over the ISA bus while retaining the simplicity of the ISA bus protocol. Also, the EISA bus is still offered in many systems that have decided to drop the ISA bus as the secondary bus to the PCI bus. Finally, to allow CRAM to interface to systems other than the PC, a VME bus [39] interface has also been studied. VME is an asynchronous bus with data transfers of up to 40M bytes/s for 32-bit buses and 80M bytes/s for 64-bit buses.

As a proof-of-concept for the CRAM controller, and to enable a more realistic study of the effect on performance, area, and design effort, two interface units based on the PCI and ISA buses have been designed and implemented in two separate Xilinx FPGA controller prototypes. While the EISA and VME interfaces have not been implemented, their bus data transfer parameters have been used in the study of the effect of the CRAM-host interface on the performance of a CRAM system. The following sections describe the PCI and ISA interface units.

4.2.1 CRAM-ISA Interface Unit

Because of the relatively small amount of hardware and low design effort required to implement the ISA bus protocol, it was decided to use the ISA interface in the first controller prototype. This presented an ideal starting point for a Xilinx FPGA controller prototype because of the area and speed constraints in programmable devices. Figure 4.2 shows the block diagram of the CRAM-ISA interface unit. The address decoder compares the address on the LA bus with the value set by the CRAM address jumpers. If there is a HIT, the data transfer is executed by the ISA bus control logic. The CRAM interface unit, like the one in the CRAM-PCI interface unit (Figure 4.3), generates read/write control to the CRAM controller units and passes their status to the ISA bus control logic.

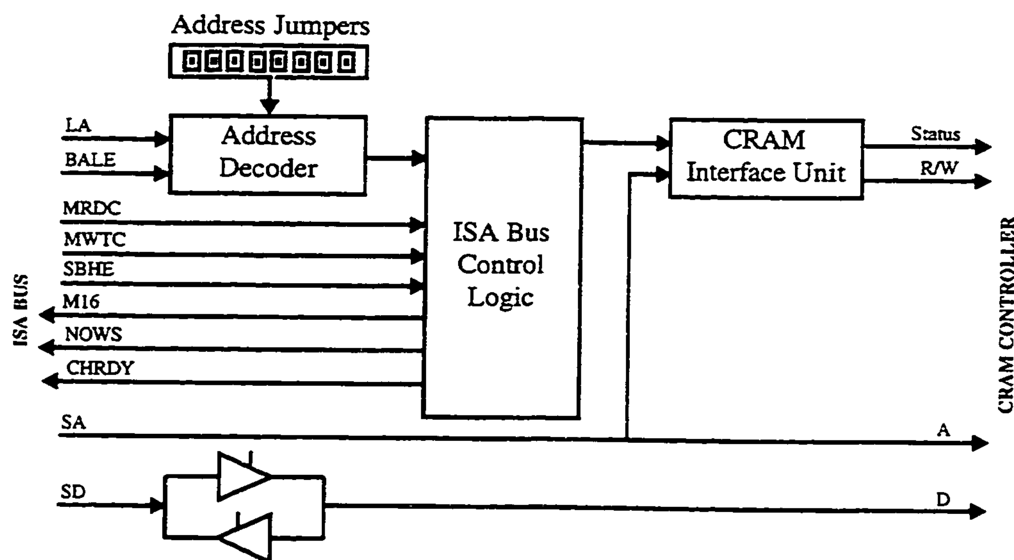

Figure 4.2 CRAM-ISA Interface Unit

4.2.2 CRAM-PCI Interface Unit

The PCI interface unit is shown in Figure 4.3. The address comparator and command decoder decode the PCI cycle and passes it to the PCI Target FSM/Control logic, where it is executed. The parity checker/generator generates even parity on the 32-bit Address/Data bus and the 4-bit Command/Byte Enable bus during PCI read cycles. This is implemented on the CRAM controller because it is mandatory for all PCI devices. Parity checking on the AD and CBE buses during write cycles is not mandatory, but it has also been implemented on the controller to guard against errors especially when loading microinstructions and macroinstructions from the host. This was also influenced by the minimal hardware required to implement parity checking. To minimize the area of the CRAM controller, only the PCI-required configuration registers have been implemented. These registers, together with other implementation details of the CRAM-PCI interface unit, are described in Appendix A.

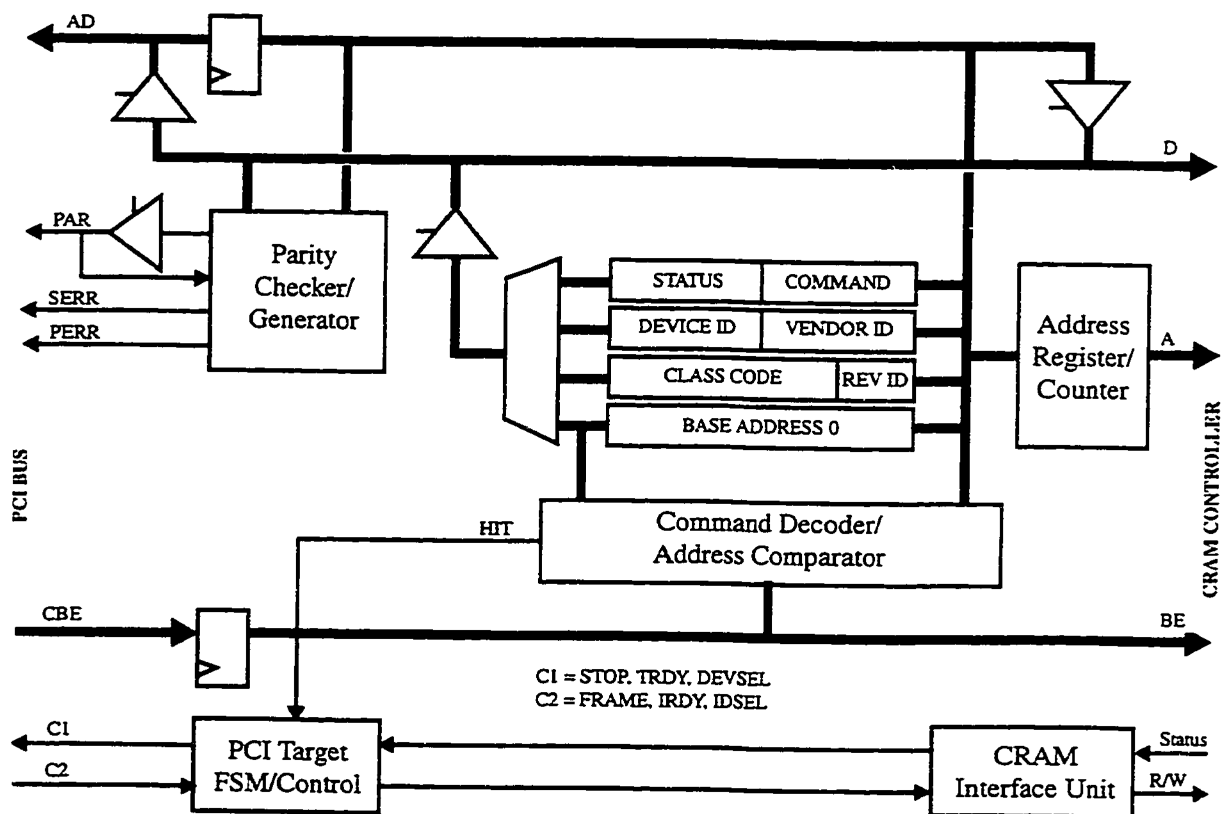

Figure 4.3 CRAM-PCI Interface Unit

4.3 Instruction Queue Unit

All CRAM instructions from the host processor are first loaded into the Instruction Queue Unit (IQU). Once in the IQU, instructions are executed by the CRAM controller completely independent of the host processor. The IQU is shown in Figure 4.4, and consists of an instruction FIFO, an instruction fetch unit, and the instruction register (IR). The FIFO removes the need to synchronize the external (host) bus transfers of CRAM instructions to the internal execution of these instructions. It also allows the host processor to transfer the instructions in advance without waiting for the previous instruction to finish executing, thus improving PE utilization. Otherwise, if instructions were loaded from the external bus directly into the instruction register, there would be big breaks in the instruction flow (several idle clock cycles between instructions) due to bus transaction delays. Also, for buses such as the PCI, EISA and VME, the instruction FIFO allows the host processor to transfer instructions at increased speed using burst cycles.

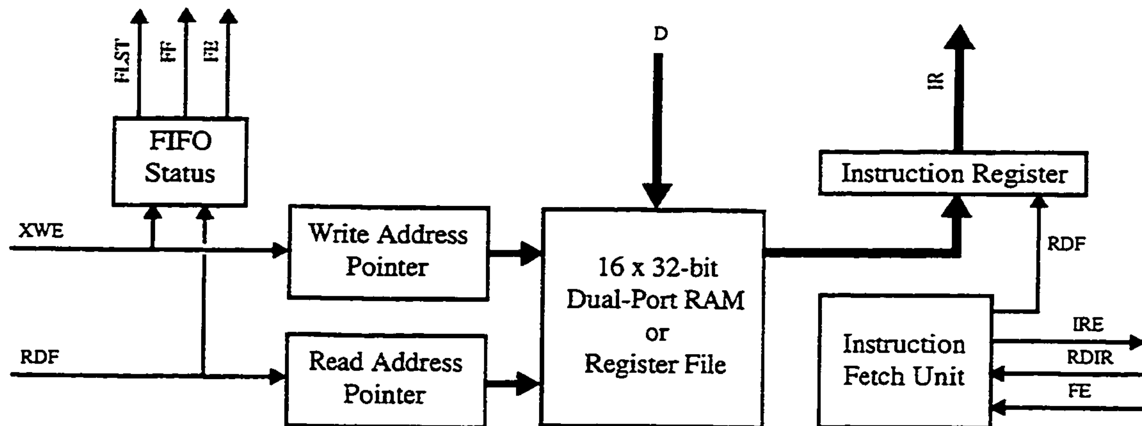

Figure 4.4 Instruction Queue Unit

4.3.1 Effect of IQU on Performance of PE Array Controller

The total time, T_{exe}, to execute an instruction on a CRAM-like three-layer instruction flow (Host Computer ⇒ Controller ⇒ PE Array) is given by

$$T_{exe} = T_{init} + T_{load} + T_{flow} + T_{pe} \tag{4.1}$$

where T_{init} is the host overhead of initializing the instruction and setting up its transfer to the CRAM controller, T_{load} is the time to transfer the instruction on the host bus, T_{flow} is

the time for the instruction to flow through the instruction path on the controller, and T_{pe} is the time to execute the instruction on the PE array. If the instruction has n microinstructions, and the cycle times for the CRAM system and the host bus are T_c and T_{bus}, respectively, then

$$T_{pe} = nT_c \tag{4.2}$$

$$T_{flow} = 2T_c \tag{4.3}$$

$$T_{load} = 2T_{bus} \tag{4.4}$$

Note that the instruction path on the CRAM controller is a three-stage pipeline (FIFO ⇒ IR ⇒ Pipeline Register) and hence an instruction takes two clock cycles to flow from the FIFO to the pipeline register. If the components of Equation 4.1 were always executed sequentially, then the total execution time for a sequence of N instructions would be

$$T_{exe} = N[T_{init} + 2T_{bus} + (n+2)T_c] \tag{4.5}$$

However, the instruction queue unit allows that some of these components are executed in parallel with each other. This improves the performance of the CRAM controller when compared to controllers that only have IR in the instruction path, such as the MIT Pixel Processor [14] and the LUCAS Associate Array Processor [31]. IQU affects T_{exe} differently for the cases $(T_{pe} + T_{flow}) \leq (T_{init} + T_{load})$ and $(T_{pe} + T_{flow}) > (T_{init} + T_{load})$.

For the case where $(T_{pe} + T_{flow}) \leq (T_{init} + T_{load})$, the performance of the CRAM controller is better than that of an IR-only controller because of two reasons. First, for an IR-only controller, only one instruction can be transferred at a time (after IR is empty). Therefore to transfer N instructions, the host computer has to set up the transfer N times. This increases T_{init} when compared to the CRAM controller where a transfer of more than one instruction is possible. Also, for an IR-only controller, the total instruction load time, T_{load}, is always equal to $2NT_{bus}$. On the other hand, for host buses that support burst transfers, the transfer of instructions from the host to the CRAM controller may be executed in a shorter time, $(N + 1)T_{init} \leq T_{load} \leq 2NT_{bus}$, depending on how many instructions are transferred during each burst. The second reason for the improved performance is due to the way data transfers to an external bus are executed on the host

computer. Once the host processor has set up a data transfer, the unit responsible for handling such transfers (an external bus interface unit, a load/store unit, or a dedicated I/O processor) can execute the data transfer in parallel with the execution of instructions in the other units of the processor. Since an instruction can be loaded onto the CRAM controller while the previous instruction is still executing, the initialization and setting up of the instruction on the host can be done in parallel (fully or partially) with either the execution of PE microinstructions (T_{pe}), the flow of an instruction in the controller (T_{flow}), or the transfer of an instruction on the host bus (T_{load}). This reduces T_{init} even further.

For the case where $(T_{pe} + T_{flow}) > (T_{init} + T_{load})$, a new instruction is loaded into the FIFO during the execution time of the previous instruction. This means that T_{init}, T_{load}, and T_{flow}, do not contribute to the total execution time of the instruction, and T_{exe} becomes equal to T_{pe}. This not only reduces the execution time to its minimum possible value, but also makes the execution time of CRAM instructions independent of the transfer characteristics of the host bus. This is very important for CRAM as a general-purpose system because it means that CRAM systems can be implemented for a variety of platforms, including slow host systems such as ISA-based computers and embedded systems that use slow microcontrollers.

Another parameter for measuring the performance of an SIMD controller is PE utilization. This is defined as the percentage of the total execution time for which the PEs are busy, i.e.

$$PE_{utilization} = \frac{T_{pe}}{T_{init} + T_{load} + T_{flow} + T_{pe}} \times 100 \quad \textbf{(4.6)}$$

For all cases, a controller approaches an ideal controller (100% PE utilization) if $T_{pe} >> (T_{init} + T_{load} + T_{flow})$. However, for the CRAM controller, this happens at a smaller value of T_{pe} because of the parallelism possible with IQU. In fact, at the point where $T_{pe} = (T_{init} + T_{load} + T_{flow})$, the utilization is 100% as opposed to 50%.

One way to characterize the performance of an array controller is to execute a sequence of instructions of varying number of microinstructions per instruction [14]. Figure 4.5 shows the simulation results based on a 50 ns CRAM system interfaced through the PCI bus to a 133 MHz Pentium PC. As noted before, not only is the performance

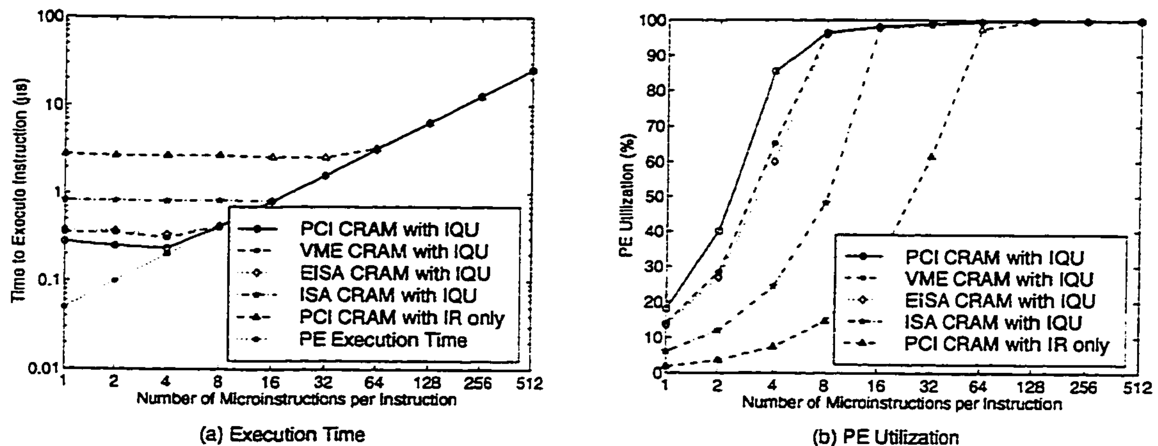

Figure 4.5 Effect of IQU on Controller Performance

better than for a controller with IR only, but it also becomes independent of the host bus at a smaller number of microinstructions per instruction. Another important thing to note from Figure 4.5 is the controller performance for short sequences, i.e. instructions with a small number of (less than 30) microinstructions. These instructions generally exhibit lower PE utilization because T_{pe} is much less than $(T_{init} + T_{load} + T_{flow})$. However, the IQU allows the performance of the CRAM controller to be as high as 10 times the performance of an IR-only controller. For application specific systems, especially pixel-processing, these instructions are dismissed as not typical because operands are generally 8-bit wide or bigger. But for a general-purpose system such as CRAM, short-sequence instructions are common, and are especially used in operand extension, updating of controller registers, and optimization of operations. They may also constitute the main operations in databases which have records with less than 8-bit precision. For example, a database of the ages of children at a nursery school, or the number of courses taken by a student per semester, is efficiently stored as 4-bit numbers on CRAM because these numbers are always less than 16, and 4-bit numbers use half the memory of, and execute twice as fast as 8-bit numbers. Therefore, short-sequence instructions may constitute a significant percentage of the total number of instructions if the application has either low-precision (or small-size) operands, or if the number of operations in the program is very small compared to register-initialization instructions. Table 4.1 gives examples of applications with a high percentage of short-sequence instructions. These applications are described in

detail in Chapter 7. Notice that for 4-bit database records, almost all instructions in the search operations have less than 30 microinstructions. Therefore, the performance improvement for short-sequence instructions in the CRAM controller is of vital importance for a general-purpose CRAM system.

Application	Total Number of Instructions	Number of Instructions with less than 30 Microinstructions	Percentage of Instructions with less than 30 Microinstructions (%)
Low-level image processing:			
Brightness adjustment	9	8	89
Spatial average filtering	329	266	81
Edge enhancement	77	60	78
Conversion to binary image	5	5	100
Multiple-threshold segmentation	72	72	100
Database searches (32-bit records):			
Equal-to search	9	8	89
Maximum search	7	5	71
Greater-than search	9	8	89
Between-limits search	13	10	77
Database searches (4-bit records):			
Equal-to search	9	9	100
Maximum search	7	6	86
Greater-than search	9	9	100
Between-limits search	13	13	100

Table 4.1 Percentage of Short-Sequence Instructions

4.3.2 Instruction Fetch Unit and FIFO

The fetch unit reads an instruction from the instruction FIFO into the 32-bit instruction register when the current instruction finishes executing. It is comprised of two simple one-hot encoded finite state machines (FSMs). The 3-state fetch FSM is synchronized to the FIFO read/write and is therefore clocked by the external bus clock. The 2-state read FSM connects the IQU to the instruction execution unit and runs at the internal clock of the CRAM controller.

The instruction FIFO is a 16 x 32-bit dual-port RAM or register file. This acts as a buffer for the instructions from the host computer to the instruction execution unit. Instruction FIFO status include FIFO Full (FF), FIFO Empty (FE) and FIFO Last Word (FLST) flags. FLST is used for buses, such as the PCI, that signal the termination of a burst transfer one cycle before the last one.

The penalties of a small FIFO are twofold. First, a small FIFO is more likely to be completely empty before the next group of instructions is loaded into it, especially if long-sequence instructions are mixed with short sequences. With an empty FIFO, one or all of the components T_{init}, T_{load}, and T_{flow} of Equation 4.1 contribute to, and hence increase the total execution time. The second penalty of a small FIFO is the increased number of bus transfer retries that the host processor performs in order to transfer CRAM instructions into the FIFO. This number increases for a small FIFO because the FIFO is full more often. While the number of transfer retries does not directly affect the total execution time of an application, it may affect the performance of the host processor if it is operating in a multitasking environment. Apart from increasing the size of the instruction FIFO, the number of transfer retries may be reduced by using an interrupt, instead of polling, to signal the FIFO-full condition

To determine an optimum size of the FIFO, a number of simulations were carried out using practical applications and other theoretical combinations of instructions. An example of the results is shown in Figure 4.6. These were obtained by running Vector Quantization (VQ), with a 64-word codebook and 2 x 2 vectors, on a 50 ns PCI-based CRAM system. (see Section 7.5.1 and Appendix D.3 for the definition, implementation, and CRAM C++ source code for VQ). VQ is reported here as an example because it

includes a variety of operations likely to be found in most applications. These operations include addition/subtraction, addition/subtraction with an immediate integer constant, operand extension, comparisons, conditional execution, updating of controller registers,

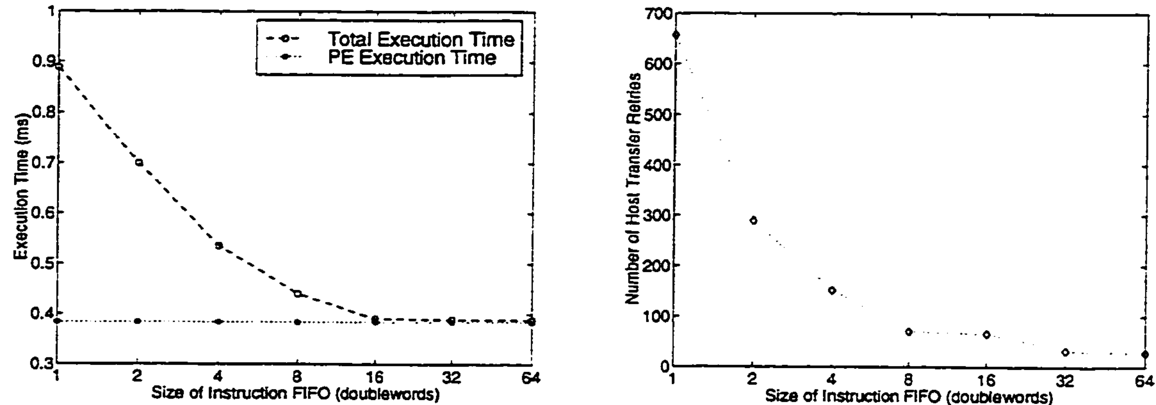

Figure 4.6 Effect of FIFO Size on Performance

and variable assignment. Based on these simulations, a FIFO size of 4 to 16 doublewords is adequate. It was also noted that increasing the size of the FIFO beyond 4 doublewords does not result in a proportional decrease in the execution time. However, a size of 16 doublewords was chosen for the Xilinx and ASIC implementation because of two reasons. First, for a Xilinx FPGA, a 16-word RAM occupies exactly the same number of CLBs as any other RAM with 15 words or less [40]. Secondly, the maximum number of continuous bus cycles allowed for the PCI bus is 16 [36]. Therefore a FIFO with more than 16 words does not have a significant transfer-time advantage, especially if the host processor has to re-arbitrate for the bus after completing the first 16-cycle burst transfer [36].

4.4 Instruction Execution Unit

The instruction execution unit (IXU) gets the instruction loaded in IR and outputs the lowest-level control, address and data signals to the CRAM. It is composed of a microprogram sequencer, an address unit, and parameter, status and command registers. The following sections describe the components of the execution unit and the formats of CRAM instructions.

4.4.1 Instruction Word

All CRAM instructions are 32-bit wide and have a uniform format (Figure 4.7). This RISC-like format reduces the overhead of instruction decoding and flow.

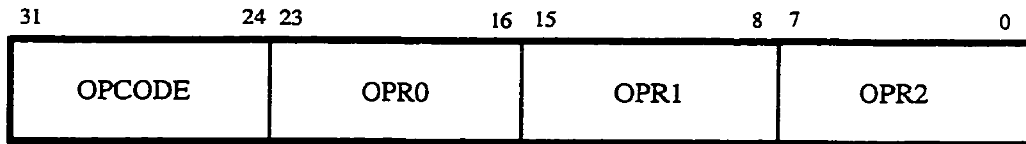

Figure 4.7 instruction Format

The 8-bit OPCODE represents the operation to be performed by the instruction, e.g. ADD. This field points to (is the address of) the beginning of the operation microinstructions in the control store. The 8-bit width means that the control store can be up to 256 words unless address mapping is used. While it might appear that this also limits the number of instructions to at most 256, this is not so because the control store is writable (Section 4.4.5) and hence any address can be re-mapped (even during run-time) to a new instruction by loading its microinstructions from the host computer. Since CRAM application programs will most likely be written in the CRAM C++ programming language (Section 6.2), this address-instruction re-mapping will rarely be necessary because the size of the control store (256 words) has been chosen to accommodate all the microroutines required to implement the language. Section 4.4.5 gives more details on the choice of the size of the control store.

There are three 8-bit operand fields in the instruction. These can either be addresses of operands in a CRAM operation, or they can be immediate values in instructions for loading CRAM controller registers. Immediate values for CRAM operations are not packed into the instruction. They use the write buffer and constant broadcast unit as outlined in Section 4.5.2. To address CRAM memory rows beyond 256, the operand addresses are extended using address extension registers (AX0-AX2). This is described under Address Unit in Section 4.4.6.

One major disadvantage of coding all instructions with a uniform format is that for instructions that require no or fewer operands, some bit fields carry no relevant information and hence waste bandwidth. For example, the instruction to reset a variable in

CRAM memory requires only one operand (the address of the operand). However, this instruction is coded as a 32-bit word, with 16 bits carrying no relevant information. One solution is to pack these unused bits with other information that may be used by this instruction (i.e. extending the functionality of the instruction) or other subsequent instructions. This, however, means extra decoding and a more complex execution unit. In our design, we have exploited the special architecture of the CRAM PEs and the multiplexed CRAM TTOP/Data bus by grouping microroutines that differ by COP and/or TTOP on only one microinstruction. Instead of loading all the microroutines of such instructions into the control store, only one representative routine is loaded and the differing COP and/or TTOP of the instruction is packed into the unused operand fields of the instruction. Two microinstruction bits indicate whether COP and/or TTOP is either coded into the microinstruction (internal COP/TTOP) or should be extracted from the instruction operands (external COP/TTOP). Using this approach has reduced the number of microinstructions required for implementing the CRAM Assembly Language (Section 6.4) from 462 to 197 (more than 50% reduction). This is the main reason why a control store of less than 256 words is feasible in our design. The only hardware overhead added by this are two 8-bit muxes to select between internal and external COP/TTOP. However, these muxes do not add to the critical path of the controller since they are after the pipeline register while the critical path is between IR and the pipeline register (Section 4.4.3). This microinstruction grouping is described in detail in Section 6.6.

4.4.2 Microinstruction Word

The 32-bit CRAM microinstruction word is shown in Figure 4.8. The following paragraphs describe the individual control bit fields. Details of the actual bit settings are described in Appendix A.1.

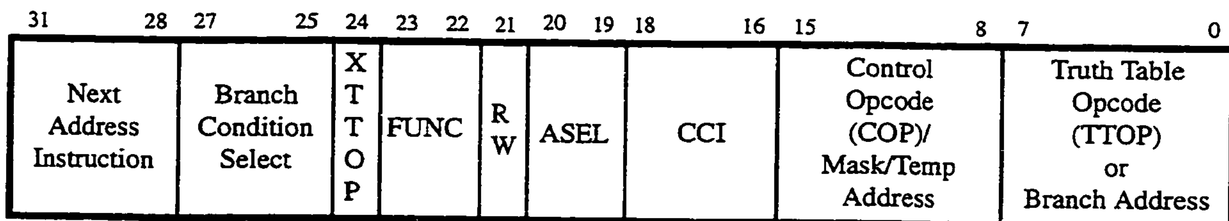

Figure 4.8 Microinstruction Word

The 4-bit Next-Address Instruction field selects the next address of the control store (i.e. the next microinstruction). The Branch Condition selects one of the eight conditions to be used in a conditional jump or loop instruction. These conditions include the status of the loop counter, the status of the CRAM global-OR and PE shift outputs, and whether the byte read from CRAM using a special read-bank instruction contains all zeroes. The last condition is used for examining the results of searches and comparisons as described in Section 4.4.3. EXTTOP selects TTOP to be the value of either operand 2 (OPR2) of the instruction word or bits [7:0] of the current microinstruction. FUNC, together with RW, is used to indicate a no-operation, a PE operation, or an internal or external CRAM memory access. Note that any controller instruction (CCI) may still be embedded in the microinstruction that has FUNC set to CRAM no-operation. ASEL selects the value to be put on the CRAM address bus. This can either be one of the three address registers, or it can be COP or the address of a system mask or temporary (scratch) location (TMPA). CRAM controller instructions (CCI) are used for loading parameter registers and incrementing address registers. These instructions are described in Table A.4 and Table A.5.

Control bits [15:8] carry either COP or TMPA. There are two reasons why these are put on the same position in the control word. Firstly, they are always used exclusively and hence can occupy the same control word bits without requiring any hardware to multiplex them. This reduces the size of the control word. Secondly, since COP is multiplexed with addresses on the CRAM chip address bus, coding COP as an address means that the address bus multiplexer has to select between four addresses only (AR0, AR1, AR2 and TMPA) without the need for a 5-to-1 multiplexer to cater for the COP input. This reduces the size and critical path of the address unit (Section 4.4.6).

TTOP is coded on the eight least significant bits of the microinstruction. To limit the size of the microinstruction, these bits can also be occupied by the branch-to address during jump (JMP, JMPS, JSR) instructions. This means that these instructions cannot be used in a microinstruction that executes a PE operation (they can however be used with a CRAM read/write operation). This restriction does not affect the performance of the system because, as mentioned in Section 4.4.3, jump instructions are used very rarely in simple microroutines like the ones for CRAM. Also, since a PE operation is usually

followed by a memory operation, it is more likely that the jump instruction will be after the memory instruction than after the PE operation. In this case, the jump can be put in the same microinstruction as the CRAM memory operation, thus avoiding the overhead of using an extra microinstruction just to implement the jump.

4.4.3 Microprogram Sequencer

The microgram sequencer is designed as a fairly standard architecture (Figure 4.9). The 8-bit opcode from the instruction register (IR[31:24]) points to the starting address of the microinstructions to be executed. The microprogram counter (uPC) provides the address of the next sequential microinstruction. The stack, the current address register (CAR) and the loop counter are used in subroutine calls, branches and looping.

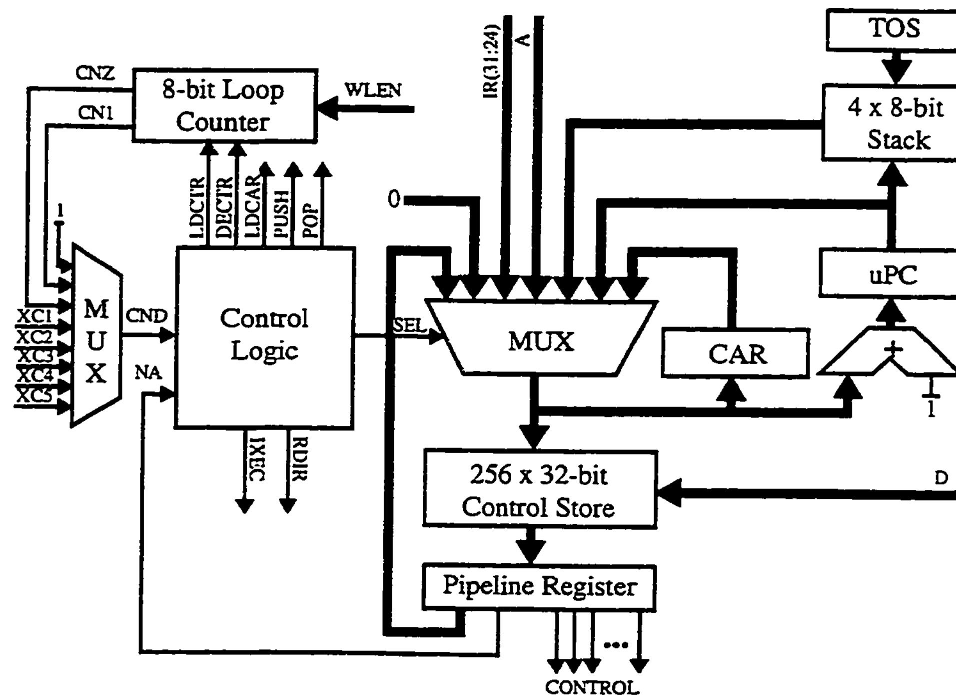

Figure 4.9 Microprogram Sequencer

The loop counter is loaded from the Word Length (WLEN) register (Section 4.6) at the beginning of every instruction. This represents the bit length of the operands in the current instruction. Even though the current CRAM C++ library consists of only CRAM integer objects (usually up to 32-bit values on many systems), the WLEN and loop counters are

designed as 8-bits for future expansion (such as the inclusion of floating-point CRAM objects, which will be 64-bits or more). It must be pointed out that the CRAM controller and the CRAM C++ library do not restrict the size of CRAM integers to 32 bits. An application program can use integers of any size up to 256 bits. Also, the 8-bit loop counter allows operations that involve all the rows of the CRAM memory to execute faster. For example, clearing the memory of the C64p1k CRAM chip (1 Kbits/PE) requires only 4 MCLR (memory clear) instructions to be issued by the host processor compared to 32 instructions if the loop counter was 5-bit wide (for up to 32-bit integers).

The control logic generates signals for selecting the next address of the control store (the next microinstruction), loading and decrementing the loop counter, and reading the next instruction from the instruction register (RDIR). The conditions that can be used in looping, branches and subroutines are shown in Table A.2. Conditions derived from within the sequencer include CNZ and CN1. The external conditions (XC1-XC5) are mostly from the CRAM chip (GOR, SLO, and SRO). DZ is primarily designed for scanning out the result of CRAM search or comparison operations. In this case, the boolean result value is scanned (read) a byte at a time until a byte containing a 1 is read in. This byte is saved in the DTR register and can be used, together with the bank address register (CBA), by the host processor to calculate the index of the first PE that yielded a true (1) value during a search or comparison.

4.4.4 A Simplified Microprogram Sequencer

Because CRAM is a general-purpose parallel-processing system, the control store is not intended to include microroutines of application-specific algorithms such as FFT, pixel smoothing, etc. Otherwise the control store would be prohibitively large to be on-chip. Instead, only microroutines of basic operations such as addition, subtraction, maximum/minimum search, etc. are loaded into the control store. One noticeable feature of these basic microroutines is that they do not use most of the sequencing features shown in Figure 4.9. Their dominant feature is the n-looping through the microinstructions depending on the number of bits of the operands. Subroutines and sophisticated branching are not as necessary. Even simple branching is used rarely since the testing of the most

commonly used CRAM status (Global-OR) is usually done on the CRAM chip and not on the controller (even though the status of the CRAM bus-tie is also passed to the controller). Conditions that are usually tested on the controller (such as DZ) usually occur as end-of-looping conditions and not as branch conditions. Therefore, a simplified microprogram sequencer, shown in Figure 4.10, was designed to reduce the area of the controller. In this design, both the subroutine and branching have been removed. The Loop from (LPR) register contains the address where looping is to start from. To avoid the overhead of including an extra instruction when the first microinstruction in a microroutine is the one to loop from, LPR is automatically loaded with the address of the first microinstruction when any instruction starts executing. It can also be loaded by the sequencer SETLP next address instruction.

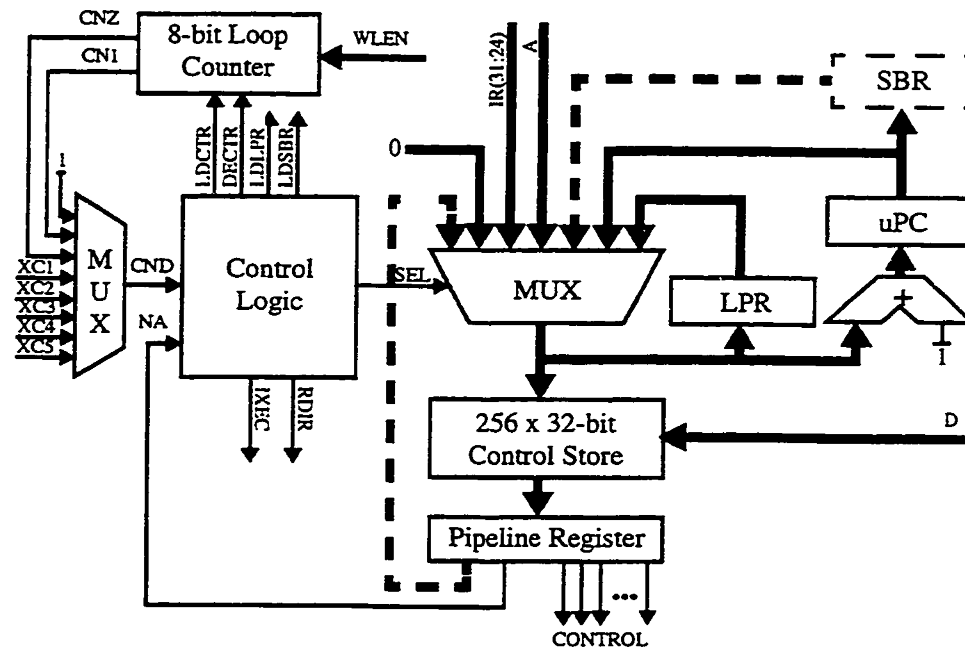

Figure 4.10 A Simplified Microprogram Sequencer

While the microroutines used to implement the CRAM C++ library do not use branching, it is our desire to retain the branching feature for future expansion. Also, the subroutine feature will not be entirely removed, but rather the multi-level subroutines will be replaced by simple single-level subroutine calls. This can be a very useful feature in implementing microroutines that are just the concatenation of the basic microroutines. If this feature is implemented, the stack is replaced by a subroutine return address (SBR)

register. In order to be able to use the existing basic microroutines in their usual stand-alone mode (which will be their most common use, anyway) as well as in the concatenation mode, subroutines will not be implemented with the usual return-from-subroutine (RTS) instruction. Rather, the calling microinstruction will set a 1-bit subroutine mode (SMR) register. Whenever a basic microroutine finishes executing, if SMR is set, the next microinstruction address is selected to be the one in SBR, otherwise the address is selected as before. These two features (branching and single-level subroutine calls) are not implemented in the prototype controller (hence are shown dotted in the simplified sequencer of Figure 4.10) because of two reasons. First, excluding these features increases the speed and reduces the area of the controller. Second, and more important, these features are not used by any of the microroutines used to implement the first revision of the CRAM C++ library. However, a controller incorporating these two features, and one using the architecture shown in Figure 4.9, has been designed, synthesized, and fully simulated (including timing simulation). The synthesizable VHDL design files can be used in future prototypes.

4.4.5 Control Store

The control store is writable from the external bus, permitting dynamic refilling of microinstructions. This is important because, unlike general-purpose processors, CRAM processing elements perform operations, such as addition, one bit at a time. This means that a very large number of such operations are possible in CRAM, thus necessitating an arbitrarily large control store. By making the control store writable, it can be made small and refilled when an application requiring different microroutines than those in the control store is to be run. A motivation for making the control store small is to reduce the cost of the controller, and hence the whole CRAM system. A smaller control store also means that it can easily be implemented on-chip, thus increasing the speed of the controller as well as reducing the cost and ease of implementation of the CRAM system (fewer discrete components on the PCB). Also, since microinstruction reading represents the heaviest data movement in the controller, putting this on-chip reduces the controller power dissipation.

The size of the control store is of vital importance to the effective design and performance of the CRAM controller. While a small control store results in a smaller area, for a reasonably large application, it can also degrade the performance of the CRAM system since microinstructions have to be continuously refilled. On the other hand, a large control store completely removes the overhead of microinstruction refilling, but adds considerably to the area of the controller and may even reduce the speed of the controller, especially if it is to be implemented off-chip. As a compromise between performance and area, it was decided to make the control store big enough to store all the microinstructions required to implement the CRAM assembly code, i.e. all the primitive operations in the CRAM C++ library. With this control store size, an application that uses standard C++ operators (+, -, +=, >, &&, etc.) does not have to refill microinstructions. Dynamic refilling or pre-loading of microinstructions is necessary only if the programmer decides to implement some application-specific routines as microcode. As mentioned in Section 4.4.1, an innovative approach has been used to group similar microinstructions. This has halved the number of microinstructions to be stored on the control store. More important, this number (197) is below 256. It is easier and cost-effective to implement on-chip synchronous RAMs of less than 256 words in both an FPGA (the Xilinx Memgen program can automatically generate RAMs of up to 256 words) and an ASIC process (some technologies, such as the Nortel 0.8 μm BiCMOS and TSMC 0.35 μm CMOS, include 128-word and 256-word synchronous RAM macros in their standard cell libraries).

From the above argument, the size of the control store can be set to 197 words or more. However, a size of 256 was chosen because of two reasons. First, the depth of RAM macros in ASIC processes and some FPGAs run in steps of 2^n ($n = 3, 4, 5, \ldots$). Common values are 16, 32, 64, 128 and 256. Secondly, the extra free space in the control store may be used for future microcode extension. It also allows pre-loading of application-specific microcode, thus reducing the penalty of refilling microinstructions during run-time. Figure 4.11 demonstrates the penalty of microinstruction refilling when the size of control store is set to less than the number of the required microinstructions (197). Again, this is based on a 64-word VQ and a 50 ns PCI CRAM system.

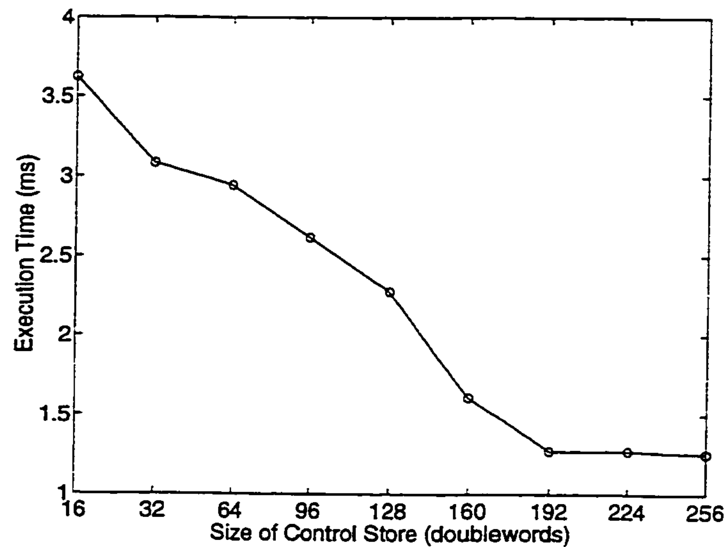

Figure 4.11 Effect of Control Store Size on Performance

In Figure 4.11, the main contributions to the increased execution time when the control store is small is the host processor execution time and the time to load microinstructions. The host processor execution time is increased because it now executes more conditional code, and also sets up more bus transfers. The time to load microinstructions is even more significant if CRAM is interfaced to a slow host bus such as the ISA bus. The other contributors to the increased execution time are the time to fill an empty instruction queue, and the time to load instructions into the FIFO. Before executing an instruction whose microinstructions are not in the control store, its microinstructions are first loaded into the controller. To make sure that active microinstructions are not overwritten, the controller status has to be read (through interrupts or polling) to make sure that all instructions in the IQU have finished executing before new microinstructions are loaded. Apart from increasing the host execution time, this also means that the instruction queue is always empty before a new instruction is loaded. This is the reason why the time to fill an empty instruction queue and the time to load instructions into the FIFO both go up as the control size decreases. It must be pointed out that as the size of the control store increases, more microroutines are loaded permanently into the control store depending on the frequency with which they are used in applications, and whether they belong to a group of instructions whose microroutines are already in the control store. For example, if the last

microroutine in the control store is for addition, then the microroutine for subtraction will be the next one to be loaded should the size of the control store be increased. The same for microroutines for comparisons (greater than, equal to, etc.), or logical operations (AND, OR, XOR, etc.). This is the reason why the curve in Figure 4.11 has a stair-case shape in some regions. The flat portions, for control size of less than 192, represent increments of control store size for which there is no corresponding improvement in performance simply because microroutines loaded in that memory range are not used by that particular program.

4.4.6 Address Unit (ADU)

The address unit shown in Figure 4.12 consists of three 8-bit address registers (AR0-AR2), three corresponding address extension registers (AX0-AX2), an 8-bit bank address register (CBA), and an 8-bit incrementer/decrementer. The address registers AR0, AR1 and AR2 (ARn) are loaded from the instruction register at the beginning of every instruction (i.e. when an instruction is read from IR by the instruction execution unit). For memory access instructions, the values in ARn are concatenated with the values of AXn and CBA to form the address of the memory. AXn:ARn forms the row address of the PE local memory, which is the address used during internal CRAM memory accesses. The bank address is used during memory accesses between the CRAM and the controller to select a group of 8 PEs (for 8-bit RAM) whose memory is to be accessed. That is, for external memory accesses, the address is CBA:AXn:ARn. AXn and CBA are preloaded from the host processor using specific controller load instructions (LDAXn and LDCBA). This allows packing up to three operand addresses in the 32-bit instruction (Figure 4.7), while using addresses which are more than 8 bits wide. 256 continuous rows can be accessed in a program without the need to update AXn. The CRAM C++ compiler can also optimize the allocation of variables in memory so that AXn need as little updating as possible. The incrementer/decrementer computes the next value of the currently selected ARn register if there is an INCA or DECA instruction in the current microinstruction.

In the first prototype, not all bits of AXn and CBA are implemented. AXn is implemented as a 2-bit register because the maximum memory per PE for the C64p1k and

C512p512 CRAM prototype chips is 1Kb (i.e. 10-bit address). CBA is implemented as a 6-bit register because the C512p512 CRAM chip has 512 PEs.

Since the address registers are always loaded with IR[23:0] at the beginning of an instruction, they also act as temporary registers for latching operands of any instruction, especially the controller load-register instructions. In other words, OPR0, OPR1 and OPR2 are just aliases of AR0, AR1 and AR2, respectively.

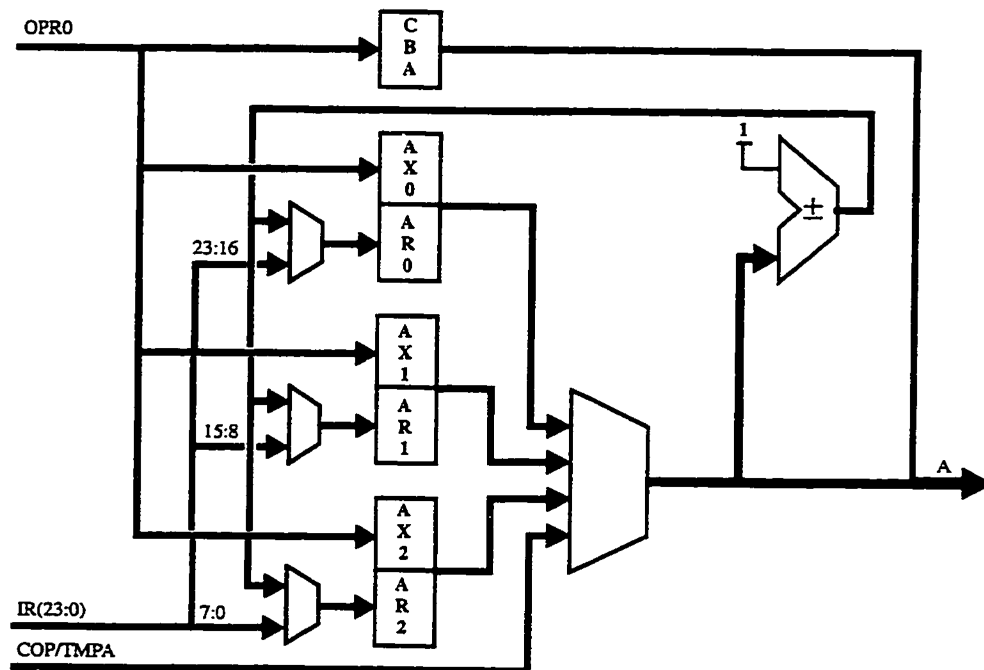

Figure 4.12 Address unit

4.5 Read/Write Buffers and Constant Broadcast Unit

Like CRAM instructions, data transfers from the host processor to the CRAM chip are buffered through the read/write buffers. In this way, data can be transferred quickly from the host processor to the controller using burst cycles, and then later transferred to CRAM. This feature also allows parallelism in that when the CRAM is executing instructions currently in the FIFO, the host can preload data into the write buffer, and then issue an instruction (which will be queued in the FIFO) to transfer this data from the buffer to CRAM. This reduces the time of initializing or reading CRAM variables.

The other reason for including data buffers in the controller architecture is that they simplify the timing of data transfer between the host and CRAM. If CRAM was connected

directly to the external bus, there would have been a need to synchronize the data transfers with the flow of instructions in the instruction execution unit. Another problem of connecting CRAM directly to the host bus is the danger of violating the bus electrical characteristics as more CRAM chips are added to the system. For example, the PCI specification states that each device (embedded or an add-in card) may place only one load on each shared PCI signal line [36]. By buffering all data transfers through the controller, a number of CRAM chips may be connected to the controller, and the loading problem becomes local to the controller. This is less of a problem because the controller is not connected to as many devices as the system bus.

The fourth reason for using the read/write buffers is that the write buffer has been ingeniously combined with CRAM data bus multiplexing circuitry to implement a novel constant broadcast unit. This is described in Section 4.5.2. Figure 4.13 shows the block diagram of the read/write buffers and the constant broadcast unit. RIBA and WIBA are the counters for the internal address of the read and write buffers, respectively. These point to the byte that is to be used in the current data transfer between the CRAM and the controller. For transfers to/from the external bus, the external bus address (A) is provided from the CRAM-host interface unit.

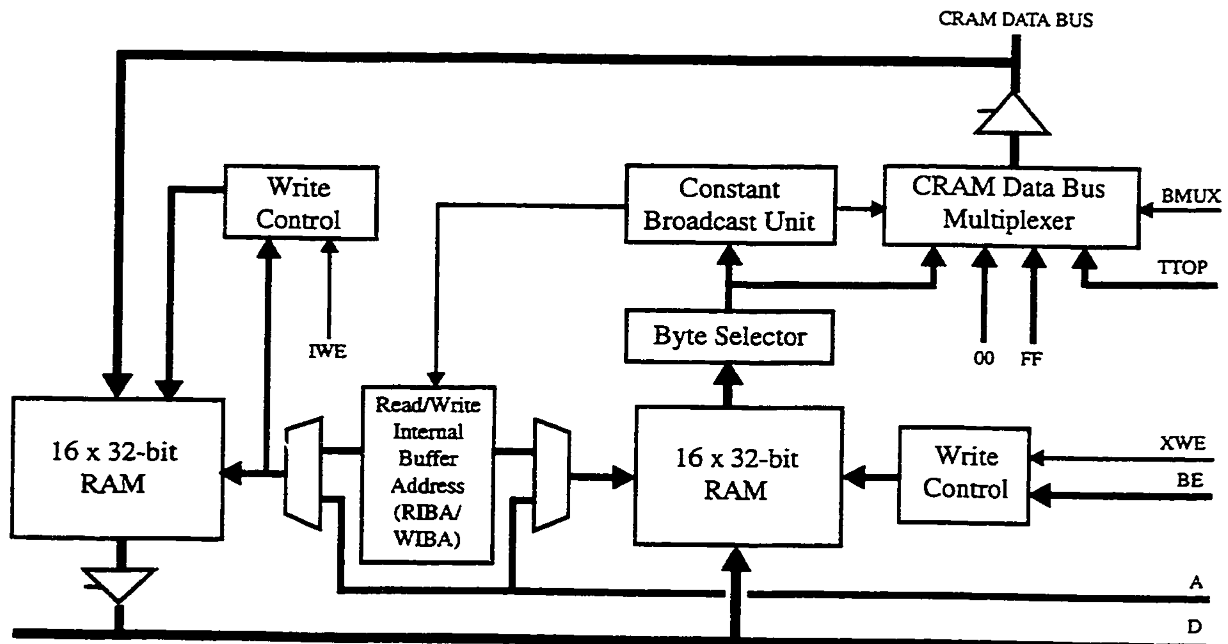

Figure 4.13 Read/Write Buffer and Constant Broadcast Unit

4.5.1 Effect of Read/Write Buffers on Variable Accesses

The read/write buffers are designed to match the sizes of the host and CRAM buses. To allow for scalability, the buffers are made of independent 8-bit memory blocks. A few such blocks are combined to match the width of the external bus so as to maximize data transferred in one cycle. Figure 4.14 shows how the buffers are organized for a CRAM system with an 8-bit CRAM data bus and a 32-bit PCI bus or a 16-bit ISA bus. Notice that because of this organization, the ISA buffers have fewer bytes when compared to the PCI. In other words, the size of the buffer is automatically scaled down to match the transfer capabilities of the host bus.

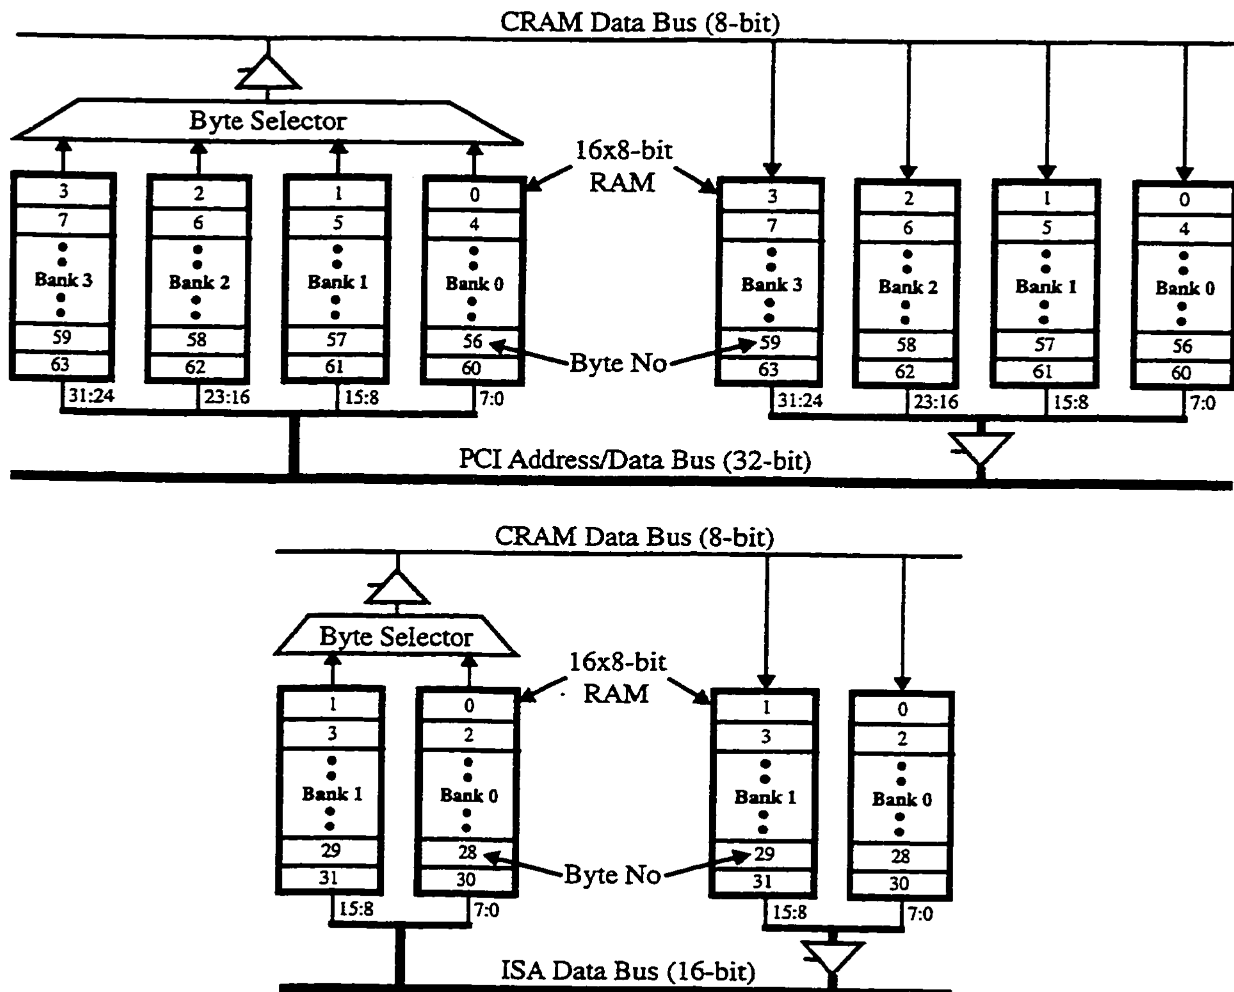

Figure 4.14 Data Buffers Organization

To transfer data for CRAM variables from the host, the data is first loaded into the write buffer. A `WRITE` instruction is then loaded into the instruction FIFO, after which it is executed, transferring the data from the buffer to CRAM. In order to reduce the total data transfer time, the loading of data and the `WRITE` instruction from the host to the controller must be done in parallel with the execution of the `WRITE` instruction. This is accomplished by dividing the buffer into two, and then executing the data transfers in chunks equal to half the size of the buffer. Assume a buffer size of B bytes, a CRAM system of cycle time T_c, and a host bus of cycle time T_{bus} (the following analysis is based on either an ISA bus or any 32-bit bus that supports burst transfers). Therefore, during each step, $B/2$ bytes are loaded into the buffer in time T_{xdata} given by

$$T_{xdata} = T_{host} + \begin{cases} \frac{B}{2} T_{bus}, & ISA \\ \left(1 + \frac{B}{8}\right) T_{bus}, & otherwise \end{cases} \tag{4.7}$$

where T_{host} is the average time to initialize and set up a data transfer on the host. The time to load the `WRITE` instruction (4 bytes) from the host to the CRAM controller instruction FIFO is

$$T_{xins} = T_{host} + \begin{cases} 4T_{bus}, & ISA \\ 2T_{bus}, & otherwise \end{cases} \tag{4.8}$$

Therefore, the total time to load the $B/2$ data bytes and the `WRITE` instruction word from the host to the CRAM controller is

$$T_{xload} = 2T_{host} + \begin{cases} \left(4 + \frac{B}{2}\right) T_{bus}, & ISA \\ \left(3 + \frac{B}{8}\right) T_{bus}, & otherwise \end{cases} \tag{4.9}$$

For CRAM with an 8-bit data bus, the execution time for the `WRITE` instruction is

$$T_{exe} = \left(2 + \frac{B}{2}\right) T_c \tag{4.10}$$

The $2T_c$ is for the instruction to flow through the instruction pipeline. If the loading from

the host is done in parallel with the execution of the `WRITE` instruction, then the total time to transfer N bytes of the CRAM variables from the host is given by

$$T_{tx} = T_{lat} + \frac{2N}{B}\begin{cases} T_{exe}, & T_{exe} \geq T_{xload} \\ T_{xload}, & T_{exe} < T_{xload} \end{cases} \qquad (4.11)$$

T_{lat} is the latency time for the parallelism, and is equal to the time to load the first $B/2$ data bytes, the `WRITE` instruction word, and three other instruction words for setting up parameter registers. Since no other operation on the CRAM chip can be executed in parallel with the transfer of data between the R/W buffers and CRAM, the minimum T_{tx} possible is when only T_{exe} contributes to the data transfer time, i.e. $T_{exe} \geq T_{xload}$. Using Equation 4.9 and Equation 4.10, the minimum size of the buffer at which this occurs can be found by using Equation 4.12.

$$B_{min} = \begin{cases} \dfrac{2(2T_{host} + 4T_{bus} - 2T_c)}{T_c - T_{bus}}, & ISA \\ \dfrac{8(2T_{host} + 3T_{bus} - 2T_c)}{4T_c - T_{bus}}, & otherwise \end{cases} \qquad (4.12)$$

Since T_{host} depend on the type of computer used as the CRAM host, the value of B_{min} can be evaluated either by simulating the data transfers on a particular host, or by approximating the average value of T_{host} (using a few simulated values) and then using Equation 4.12. The former is more accurate. Table 4.2 shows B_{min} for different CRAM systems simulated on a 133 MHz Pentium PC. For the ISA bus, the condition $T_{exe} \geq T_{xload}$ is never true for any of the tabulated CRAM cycle times because of its low data transfer rate.

	$T_c = 100$ ns	$T_c = 50$ ns	$T_c = 25$ ns
PCI	16	32	80
VME	16	48	N/A
EISA	16	56	N/A
ISA	N/A	N/A	N/A

Table 4.2 Minimum Buffer Size for $T_{exe} \geq T_{xload}$ during Variable Initialization

Figure 4.15 shows the variation of the time to load 64K bytes of a 256 x 256 8-bit image from the host to a 50 ns CRAM system. It must be pointed out that for buffer sizes of 4 bytes or less (2 bytes or less for ISA), there is no advantage in using parallelism because the number of bytes during each transfer is less than the size of the external data bus. Therefore, the loading of data into the write buffer, and the execution of the `WRITE` instruction are done sequentially. Also note that even though the reduction in the total load time in Figure 4.15 is more than 90%, we quote it at 80% because a practical system would not implement a single byte buffer (or data register) for an external bus that can transfer four bytes in one cycle. Therefore, a practical minimum buffer size is actually four bytes for EISA, PCI, and VME, and two bytes for the 16-bit ISA bus. These are the sizes used as datum for performance evaluation.

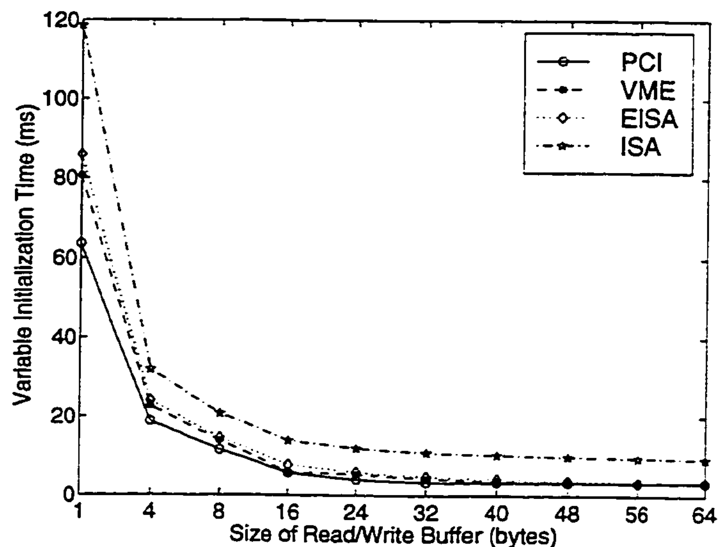

Figure 4.15 Effect of R/W Buffers on Variable Initialization

For buffer sizes of less than B_{min}, T_{xload} dominates the total load time. For this case, increasing the buffer size reduces the total load time because the host overhead is reduced (less looping, fewer data transfer initialization and setup, etc.), and for burst cycles, the actual transfer time on the external bus is also reduced. For buffer sizes of greater than B_{min}, the total load time is equal to T_{exe} (plus the latency), and hence becomes less dependent on the data transfer characteristics of the host bus. Again, this is important for

CRAM as a general-purpose system that can be used on different platforms.

In terms of choosing the size of the buffer, these simulations have shown that most of the reduction in the total load time occurs by increasing the buffer size to 32 bytes. Further increases results in correspondingly smaller performance improvements. However, because of the configuration of the buffers as explained earlier in this section (Figure 4.14), a size of 16 doublewords (64 bytes) for the 32-bit buses, and 16 words (32 bytes) for the 16-bit ISA bus, was chosen for the Xilinx FPGA implementation because RAMs with a depth of 16 or less occupy the same number of CLBs [40]. Also, for the 32-bit buses, this size is within the range of B_{min} for the cycle times of the prototype CRAM system (T_c = 100 ns), as well as that of a CRAM system likely to be implemented (T_c = 50 ns).

4.5.2 Constant Broadcast Unit

There are a number of ways to implement operations with scalar constants in bit-serial SIMD processors. The controller may produce the microroutines for each of the 2^n constants (for n-bit operands). But the amount of memory required to store all the 2^n microroutines might be prohibitively large unless n is very small. A more common method is to multiplex the microinstructions depending on the bit coming out of a select (constant) register [41]. This means that all microinstructions for constant operations must be in pairs, one for operation when the value of constant bit is '1', and the other for operation with '0'. If these pairs are laid side by side, the microinstruction word, and hence the control store size, is increased. If on the other hand, the microinstruction pairs are placed one after the other (which is the preferred approach), then the bit coming out of the select register is used as the LSB of the address to the control store. Apart from increasing the depth of the control store, this approach has several disadvantages. First, the use of the select register as one of the bits for the control store address adds extra multiplexing for the address. This may reduce the speed of the sequencer since control store address selection is usually part of the sequencer critical path. Second, code generation becomes more complicated because for operations with constants, microinstructions for operations with a '1' always have to be on an odd address, while those for '0' have to be on an even address. The third and, as far as CRAM is concerned, the most critical drawback of this

approach is that the constants' size is limited by the size of the select register, unless special programming techniques are employed. This size limitation is acceptable for systems designed to process fixed-size operands such as in pixel-processing (usually 8-bit operands). It would however limit the versatility of CRAM, in which the ability to vary the bit-length of operands (even during run time) is one of its major strengths.

A novel constant broadcast unit shown in Figure 4.16 is used for CRAM. It takes advantage of the existing write buffer and also exploits the fact that for CRAM, the PE truth table opcode (TTOP) is multiplexed with data on the CRAM data bus. The constant bit (KBIT) is selected by the write buffer internal address (WIBA) register and the constant bit pointer (KPT). WIBA, which is an already implemented feature of the write buffer, selects the byte from the write buffer using the byte selector (Figure 4.13). KPT, which is a 3-bit binary up counter, then selects one bit (KBIT) from this byte. If the currently executing microinstruction is for an operation with a constant (i.e. LDK is asserted), then LDK and KBIT tell the data bus multiplexer to select either 0x00 (if KBIT is '0') or 0xFF (if KBIT is '1'). These two values are the opcodes (TTOP) to reset or set a CRAM PE register. The register to be reset/set is specified the same way as in ordinary PE microinstructions, and is reflected in the value of COP. At the end of executing an LDK microinstruction, KPT is incremented to point to the next bit of the constant. If KPT was 7 ("111") before incrementing, it wraps over to zero and instead WIBA is incremented. This

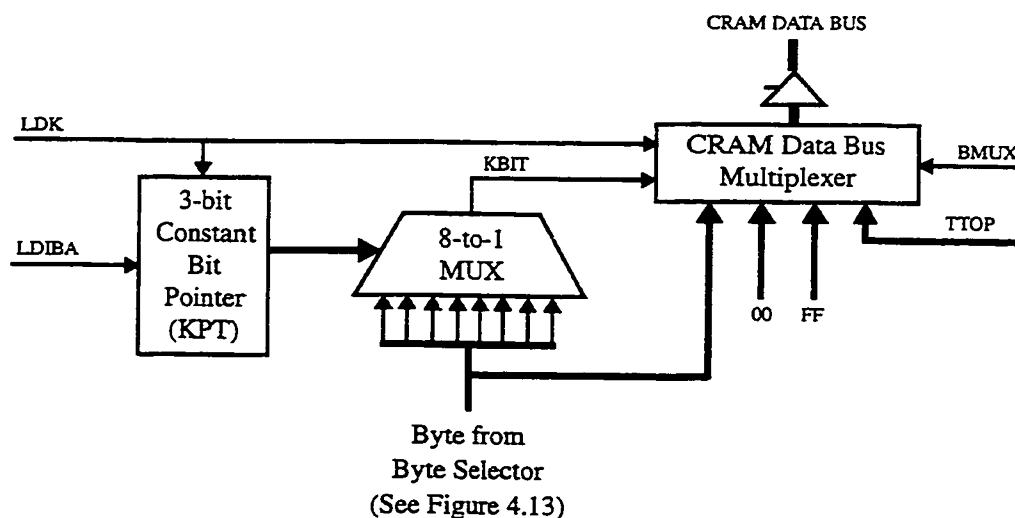

Figure 4.16 Constant Broadcast Unit

takes care of constants which are more than one byte in size or preloaded multiple constants. If constants are preloaded into the write buffer, you can point to any of these constants by loading WIBA with the corresponding buffer address. Loading WIBA also resets KPT to 0 in order to point to the first bit of the constant. If a microinstruction does not involve an operation with a constant, the data bus multiplexer will select either TTOP (microinstruction bits [7:0]) for PE operations or the data byte from the write buffer for memory writes from the controller to CRAM.

The extra hardware for the broadcast unit is the 3-bit counter (3 flip-flops) and an 8-to-1 multiplexer. The data bus multiplexer is also increased from a 2-to-1 to a 4-to-1 multiplexer. In a Xilinx FPGA implementation, the increased multiplexer does not add to the delay or area since a 2-to-1 and a 4-to-1 multiplexer both use a single CLB. In an ASIC implementation, the additional area is minimal, and the extra one-gate delay is of no effect since the critical path still remains in the sequencer.

Apart from the simplicity of this constant broadcast approach, there are other very important advantages. First, the constant size is now only limited by the size of the write buffer, which is large. Note that the size of the write buffer was made large not because we wanted to accommodate large constants, but rather because we wanted such a large buffer for efficient data transfers. The second advantage is that now a number of constants can be preloaded using the faster burst cycles, or constants can be loaded in parallel with the execution of the operate-immediate instruction or other instructions in the FIFO. This is similar to the use of the R/W buffers for data transfers as described in Section 4.5.1. The big difference is that operate-immediate instructions take a much longer time to use the data loaded in buffer. For example, an add-immediate (`ADDI`) instruction with a b-bit integer constant has $(5b + 1)$ microinstructions. Therefore, for 8-bit constants, an `ADDI` instruction takes 41 CRAM clock cycles to use one byte of the data in the buffer, whereas a `WRITE` instruction takes only 1 clock cycle to use the byte. Because of this, the value of B at which $T_{exe} \geq T_{xload}$ is much smaller for operate-immediate instructions. For 8-bit `ADDI`, the expression for B_{min} is given in Equation 4.13, and is simulated to be 8 bytes for all the system combinations tabulated in Table 4.2.

$$B_{min} = \begin{cases} \dfrac{4(T_{host} - T_c)}{41T_c - 5T_{bus}}, & ISA \\ \dfrac{16(T_{host} + T_{bus} - T_c)}{164T_c - 5T_{bus}}, & otherwise \end{cases} \quad (4.13)$$

Figure 4.17 shows the execution time of adding 256 8-bit integer constants to 8-bit CRAM variables for CRAM systems used in Figure 4.15. For buffer sizes greater than 8, the execution time of the ADDI instruction reduces by more than 35%. Also, this execution time becomes almost independent of the host bus, once again an important feature for the implementation of CRAM on different platforms.

The value of B_{min} for operate-immediate instructions is dependent on the type of instruction and the precision of the operands. Simpler operate-immediate instructions such as load-immediate (broadcast a constant value to all elements of a CRAM variable), or small-precision operands, result in fewer CRAM microinstructions. This increases B_{min}. However, since the fastest CRAM operate-immediate instruction (load-immediate) has $2b$ microinstructions (for b-bit operands), the execution time for operate-immediate instructions is always bigger than (typically more than 10 times) that of the WRITE instruction in Equation 4.10. Therefore, the buffer size set for efficient data transfers in Section 4.5.1 is automatically more than adequate for the constant broadcast unit.

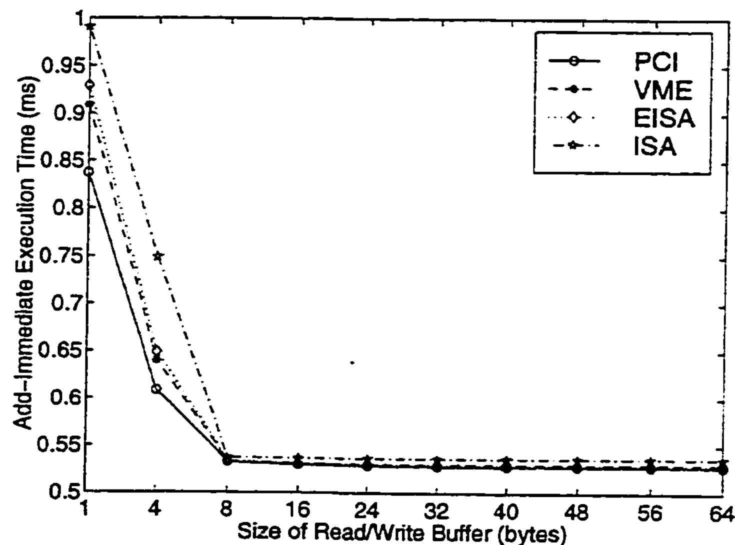

Figure 4.17 Effect of Buffer and Constant Unit on Operate-Immediate Instructions

4.6 User-Accessible Registers

There are two groups of registers that can be accessed by the programmer. The first group is memory-mapped. These include registers that are usually initialized at the beginning of program execution or at boot-up (such as command registers), and registers that are read to evaluate the status of the CRAM system. CRAM command registers specify the characteristics of the CRAM chip (such as type of PE), and the use of other controller resources such as the control store and interrupt lines. CRAM system status include global-OR and PE shift outputs, whether instructions are still executing or pending in the controller, and whether the controller issued an interrupt to the host computer. The second group of registers are those that usually change between instructions. These are used to set controller parameters such as address extension and word length. Most parameter registers are not memory-mapped and can only be initialized with specialized load-register instructions. Table 4.3 summarizes their characteristics. Details of all user-accessible CRAM controller registers can be found in Appendix A.2.

Figure 4.18 shows the memory map of the user-accessible registers. Registers outside the memory map can only be written using controller load-register instructions. This ensures that parameter registers are updated synchronously with the instructions that they are supposed to affect.

Register	Bits	R/W	Read/Write Method
Read Buffer Internal Address (RIBA)	6	W	LDIBA
Write Buffer Internal Address (WIBA)	6	W	LDIBA
Buffer Address Increment (BAI)	6	W	LDBAI
CRAM Bank Address (CBA)	6	R/W	LDCBA/Memory-mapped
Read-back Data (DTR)	8	R	Memory-mapped
Word Length (WLEN)	8	W	LDWLEN
Shift-Input selection Code (SIC)	2	W	LDSIC
Interrupt Number (INTNO)	1	W	SETINT
Address Extension Register 0 (AX0)	2	W	LDAX0
Address Extension Register 1 (AX1)	2	W	LDAX1
Address Extension Register 2 (AX2)	2	W	LDAX2

Table 4.3 CRAM Parameter Registers

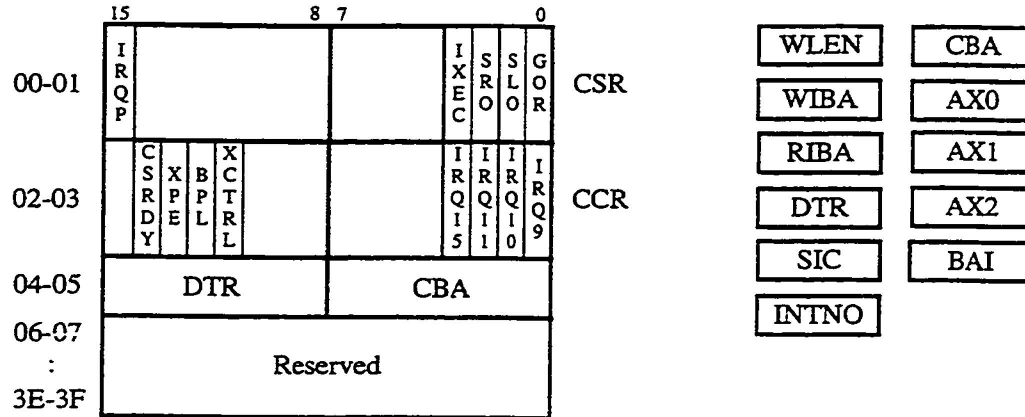

Figure 4.18 User-Accessible Registers

4.7 Memory Map

Apart from the parameter registers described in Section 4.6, all CRAM controller units that can be accessed by the programmer are memory-mapped. Since PCI configuration registers are qualified by the PCI device select (IDSEL) pin, they are not included in the CRAM controller memory map shown in Figure 4.19.

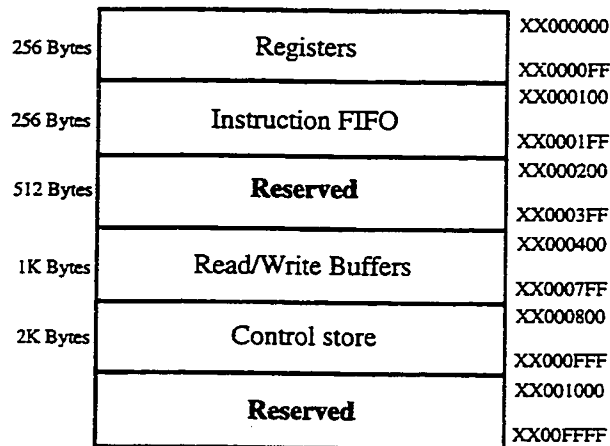

Figure 4.19 CRAM Controller Memory Map

The eight most significant bits of the address select the CRAM card. These are bits 31:24 for 32-bit buses (PCI, VME and EISA), and bits 23:16 for the 16-bit address of the ISA bus. The memory map has been allocated in such a way as to minimize the decoding of the controller units' addresses. For example, to decode the address of the control store, only

bit [11] is examined, for the read/write buffers, address bits [11:10] are compared to the value "01", and for the controller registers and instruction FIFO, their addresses are decoded by comparing address bits [11:8] to "0000" and "0001" respectively.

4.8 Summary

In this chapter, the architecture of the CRAM controller has been described. The controller has been designed with a general-purpose architecture to enhance the use of CRAM as a general-purpose parallel-processing system. A microprogrammed controller allows the optimization of application-specific tasks in software. The architecture has also been kept as simple as possible, and its area minimized, in order to reduce system costs, allow easy implementation of a CRAM system on a single size-constrained PC card, and facilitate future integration of the controller and the PE array on the same chip. In this regard, a simplified microprogram sequencer is used, and an innovative approach has been used to group microinstructions, halving the size of the control store. Various features have been used to enhance the performance of a rather relatively simple controller. A FIFO-based instruction queue improves the performance of short-sequence instructions by a factor as high as 10 times. It also allows the controller to approach an ideal controller (100% PE utilization) at a smaller number of microinstructions per instruction. This translates into increased performance for a wider range of applications. The use of read/write buffers for data transfers from the host to CRAM chips simplifies synchronization, eases electrical and physical loading of the host bus, and may reduce the time of loading data onto CRAM by as much as 80%. A novel constant broadcast unit based on the buffers improves the performance of operate-immediate instructions by more than 35%. It also reduces the size of the control store, and simplifies the use of variable-size constants. Finally, we have demonstrated that these three performance-enhancement features also minimize the effect of the host bus on the performance of a CRAM system [42]. This is very important for CRAM as a general-purpose system because it means that CRAM systems can be implemented for a variety of platforms, including slow host systems such as ISA-based computers and embedded systems that use slow microcontrollers.

Chapter 5

CRAM System Implementation

This chapter describes the implementation of a CRAM system. Section 5.1 describes the design flow, CAD tools, and technologies used in the design, verification, and implementation of two CRAM controller prototypes. Section 5.2 gives the implementation details of the PCI and ISA CRAM controllers in a Xilinx FPGA and TSMC 0.35 μm CMOS technology. This section also describes the testing of the prototypes. Section 5.3 describes the implementation of a CRAM system prototype using the C64p1k CRAM chip and the Xilinx ISA CRAM controller.

5.1 Controller Logic Design and Verification

5.1.1 Design Flow

Although a full-custom design approach results in higher performance and smaller die size [48], design automation and flexible design strategies are usually used in the design of processors and other complex digital systems in order to reduce the design time and to provide a design model that is highly adaptable to different implementation technologies and applications [49], [50], [51]. Therefore, the design flow used in the design of the CRAM controller is that of a typical automated synthesis design. The detailed steps are however specific to Synopsys, Xilinx XACT Step, and Cadence, the three CAD tools used in this case for logic design and synthesis, FPGA placement and routing, and ASIC layout design, respectively. Figure 5.1 shows the details of the design flow. The first step, which is the determination of the architecture from the design specifications is described in Chapter 4. The other steps, plus the tools used, are described in the following sections.

5.1.2 VHDL Behavioral Description and Simulation

VHDL has the advantage as a specification and synthesis language in that it can describe hardware at various levels of abstraction, from the architectural (behavioral) level down to the structural level [52], [53]. It also provides a generic design entry platform in that the design may be created in VHDL and the target implementation technology chosen later [54]. In this case, the synthesis tool can be used to map the design to any specific target technology, thus forgoing the task of a complete logic redesign should the initial intended target technology become obsolete or unavailable. Lastly, VHDL synthesis tools allow designers to follow a true top-down design methodology and also allow a quick comparison of trade-offs in implementing the same design in different technologies or styles [55]. In general, high-level synthesis allows designers to describe a design in a purely behavioral form, devoid of implementation details, and then synthesize the design using CAD tools. This achieves higher productivity gains since the design process is moved to higher abstraction levels, where designers can specify, model, verify, synthesize, simulate and debug designs in less time [52]. It is for these reasons, plus the general

advantages of automated synthesis, that VHDL synthesis tools were employed in the design of the CRAM controller. The choice of VHDL over Verilog as the description language for the design was justified by the author's level of knowledge and experience with VHDL when compared to Verilog, and by the fact that the Xilinx tools available at the time (XACT 5.2.1) only supported VHDL for post-place and route (timing) simulation.

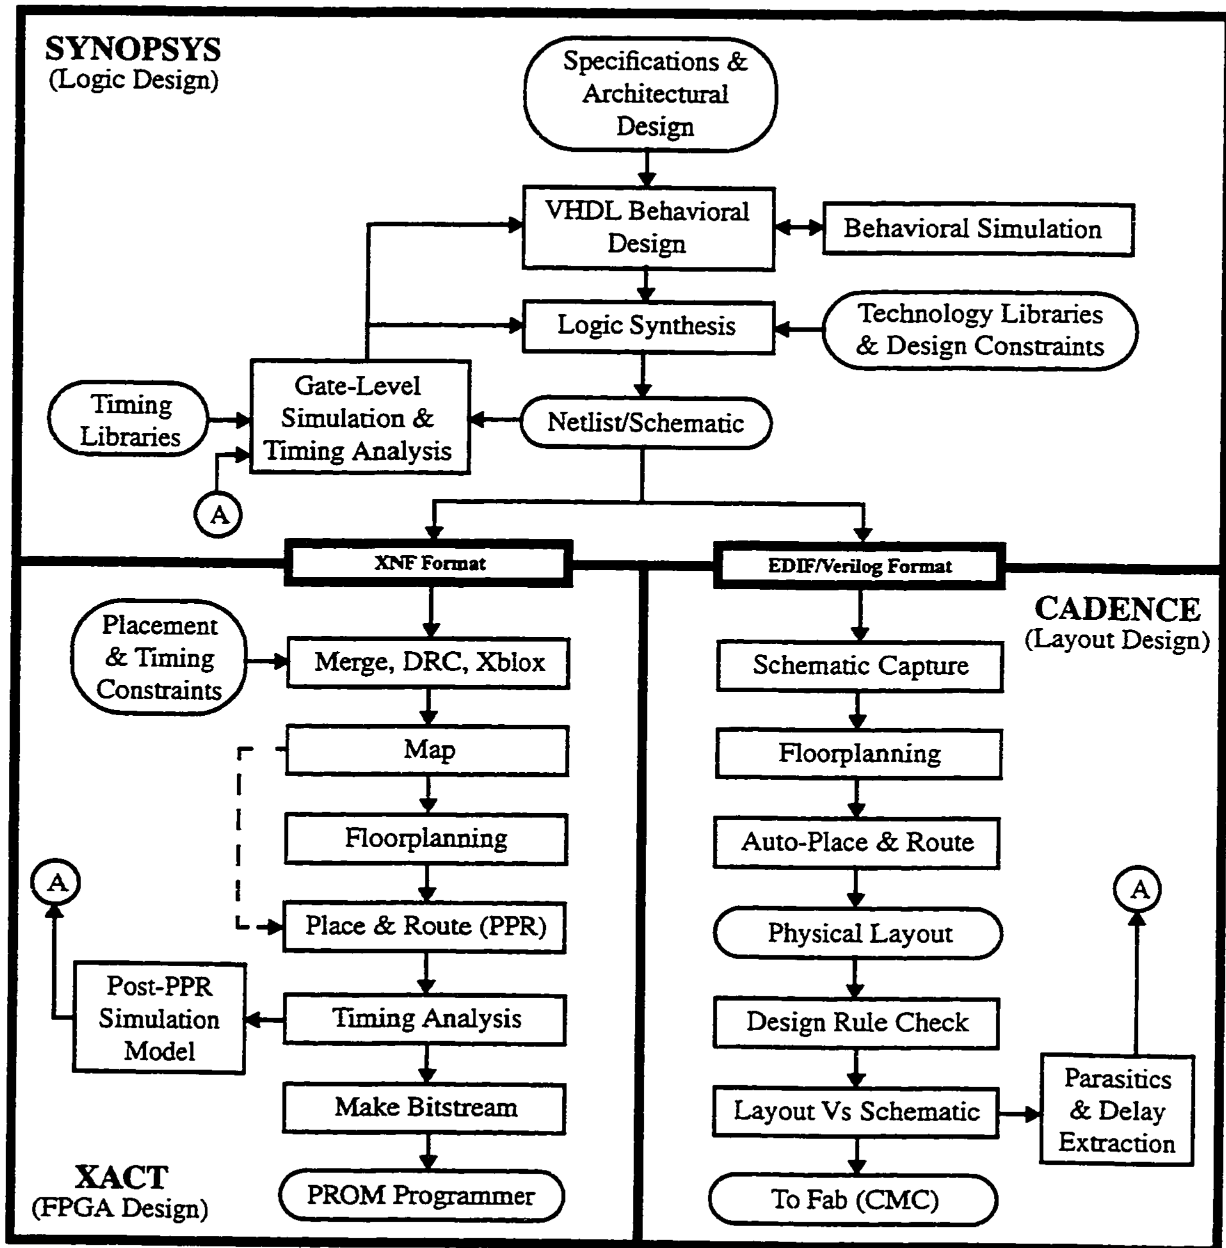

Figure 5.1 Design Flow

The VHDL behavioral description step involved a detailed specification of the controller through several iterative steps of partitioning, description and refinement. Variations in VHDL coding styles have significant effect on the logic synthesis optimization process. This is known as the problem of syntactic variance [52]. Generally, a structured, object-oriented style, that separates control and operations, is recommended for the writing of behavioral descriptions for synthesis of area-efficient hardware [56]. It is also important that the designer have very good knowledge of and experience with the VHDL language, as well as the translation, optimization and technology-mapping techniques used by a specific synthesis tool. In order to predict the results of the synthesis process, and hence effectively guide the tool in synthesizing more efficient and simpler hardware, one also needs to have a fairly good knowledge of conventional digital design techniques. All these factors have been extensively and effectively applied in this design.

After the behavior of each component has been described and simulated to verify its functionality, the components were connected up using structural VHDL, and then the whole controller was simulated to verify its functionality. Therefore problems relating to the functionality and interconnection of the controller functional blocks were detected and corrected at this stage. This structural design specification was introduced at this early stage, as opposed to using a flattened behavioral design, because the synthesis task becomes much easier and more efficient as the amount of structural detail increases [52].

5.1.3 Logic synthesis, Timing Analysis, and Gate-Level Simulation

High-level synthesis consists of converting the abstract VHDL behavioral description of the design into a register-transfer level (RTL) design implementing that behavior. Each component in the synthesized RTL netlist is then designed into the target implementation technology by means of logic synthesis and technology mapping tools. For this design, the target technologies include Xilinx XC4000E series FPGAs, Nortel 0.8 μm BiCMOS (BATMOS), HP 0.5 μm CMOS (CMOSIS5), and TSMC 0.35 μm CMOS.

After the synthesis, the Synopsys tools were used for static timing analysis and gate-level simulation. The timing analysis included investigating the design maximum/ minimum delay paths, critical paths, as well as setup and hold times. The netlist was then

saved as a VHDL gate-level netlist and simulated to verify the functionality of the synthesized circuit. Unlike behavioral simulation, the gate-level simulation also includes gate delays and hence gives approximate timing information of the design. However, this timing information does not include routing and parasitics delays, which may be more dominant than gate delays for FPGAs and small feature size technologies such as the 0.35 μm. Routing delays are simulated in post layout (or post-PPR) simulation.

5.1.4 Xilinx FPGA Design Flow

Once the Synopsys synthesis and simulation stages are complete, the design netlist is written into the Xilinx format (XNF) and exported to the XACT tools. The first stage in the Xilinx flow is to prepare the design for the succeeding tools. This includes converting the Synopsys XNF netlist to the 'real' Xilinx XNF, combining the XNF files to form one flat (non-hierarchical) design file, specifying timing and placement constraints (especially necessary in the PCI CRAM controller for which timing is very tight), performing design rule check (DRC), and using Xilinx-optimized blocks of logic (XBLOX). After this, the design is mapped and floorplanned. While floorplanning is seldom used in the Xilinx design flow, it yields very good results especially if timing is critical or CLB resources are limited. Therefore, for the PCI CRAM controller, floorplanning was used to reduce routing delays of critical signals in order to meet the PCI 7 ns setup time in an XC4013E FPGA. The placement constraints from the floorplan are then used by the PPR tool to place and route the design. Static timing analysis is performed using Xdelay. In order to perform a post-place-and-route VHDL simulation, a post-layout XNF file with detailed timing information is generated. This file is then used to generate the timing model of the design as a VHDL architecture and a corresponding SDF file, which are then used to simulate the controller. Once the design functionality and timing are verified through the post-layout VHDL simulation, the configuration bitstream is generated and formatted into an MCS-86 (Intel) PROM format. Finally, the controller configuration PROM is programmed.

5.1.5 ASIC Layout Design Flow

For the ASIC technologies, the design is imported form Synopsys to Cadence through an EDIF or Verilog netlist. Once in Cadence, an auto place and route view is created from the schematic, and the design is placed and routed using Cell Ensemble. After that, Design Rule Check (DRC) is performed and the layout is compared to the schematic (LVS). For BATMOS and CMOSIS5, there is no support for post-layout timing verification. That is, routing and parasitics delays are not backannotated into the original design. However, CMC is currently working on the supporting of this feature in the TSMC 0.35 μm CMOS technology. Otherwise, for now, LVS is the only final check on the design layout. This is not a major problem for a digital system running at slow speeds, like the CRAM controller is.

5.1.6 CRAM System VHDL Simulation Model

Figure 5.2 shows the model that is used in all VHDL simulations involving the CRAM controller or the CRAM system as a whole. This consists of the VHDL models of the CRAM controller, the CRAM chips, the host processor, and the host system buses, as well as C++ tools to create text-file models of CRAM microinstructions and CRAM/host code to be executed by the system.

When using this model to simulate the CRAM controller, the controller behavioral model is used for the behavioral simulation in Section 5.1.2, the Synopsys-generated gate-level VHDL model is used for the simulation in Section 5.1.3, and the post-layout VHDL model and SDF file are used for the simulation in Section 5.1.4. This use of the same testbench for simulating the controller at different design stages is advantageous because it ensures that the design is subjected to the same test stimuli. This not only reduces the design time, but also removes 'false simulation' errors that result if the design is mistakenly subjected to slightly different test inputs after it has been synthesized.

All the other components in the CRAM system simulation model use simulation-oriented VHDL behavioral models (i.e. some of the VHDL constructs used are unsynthesizable). The use of components behavioral models in a testbench, as opposed to using a more standard and common testbench that simply assigns values to the inputs of

the unit under test, allows the running of actual system machine code in testing the design functionality, thus saving time in generating the input test vectors and collecting the design test outputs [57], [58]. The VHDL simulation models of the CRAM chips (generic versions of C64p1k and C512p512) have been designed to exactly match the functionality of the actual designs. Even the bugs in the CRAM prototype chips have been incorporated in some of the models. This, when used with the synthesized model of the controller, makes the CRAM system VHDL model a very close representation of the CRAM system prototypes. Actually, the results of running machine code on this model yields the exact results and sequences as when the same code is run on the prototype system. This makes debugging of the prototype system and controller much easier and faster.

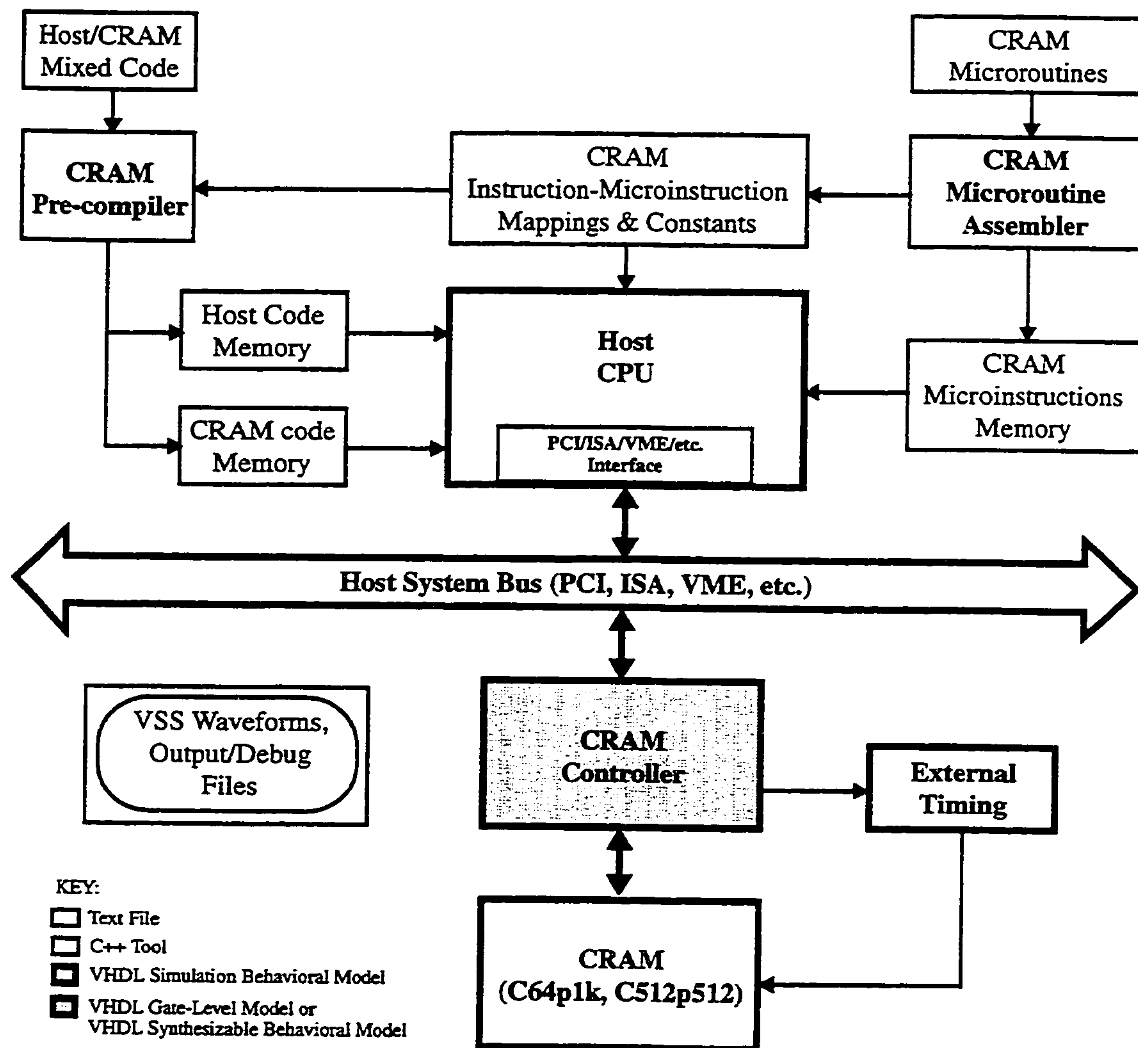

Figure 5.2 CRAM System VHDL Simulation Model

The simulation output include a display of the CRAM memory (through the Synopsys VSS waveform viewer) and output files showing results of executing the instructions on the host side of the system bus. Other output files show timing and debug information.

This CRAM system VHDL simulation environment is intended primarily for simulation during the hardware (controller and CRAM chips) design flow. It is also used for extracting more accurate CRAM controller timing information. Small algorithms and applications may also be run on the system. But since VHDL simulation, like all digital simulations, is slow, more lengthy applications or architectural exploration should be carried out using the C++ CRAM simulator described in Section 6.7.

5.2 CRAM Controller Implementation and Testing

5.2.1 Xilinx Implementation

Programmable devices are well suited for prototyping because a design can easily, cheaply, and quickly be implemented in-house without the need for the otherwise expensive and time-consuming IC fabrication process. CPLDs and FPGAs are therefore the backbones of most prototype digital systems because such projects require tighter schedules, output low volume products, and require low non-recurring engineering (NRE) expenses [59]. The choice of Xilinx FPGAs over other programmable devices was influenced by availability of CAD tools and experience with the devices in the department.

The target FPGA for the controller is a 576-CLB XC4013E FPGA. The specific part used in both the ISA and PCI CRAM controller prototypes is a 240-pin, speed-grade -2 part housed in a flat plastic quad package (XC4013EPQ240-2). Table 5.1 shows the FPGA device utilization for the different controller designs. All these designs are based on the simplified architecture described in Section 4.4.3. The ISA design has a control store size of 192 words. The PCI CRAM controller with a control size of 192 words does not fit in an XC4013E part. Therefore a PCI controller prototype with a reduced control store size of 128 words was implemented in the FPGA.

FPGA Resource	No. Available	ISA		PCI†	
		No. Used	% Used	No. Used	% Used
Occupied CLBs	576	564	97	576	100
Bonded I/O Pins	192	91	47	83	43
F and G Function Generators	1152	969	84	986	85
H Function Generators	576	262	45	233	40
CLB Flip Flops	1152	234	20	308	26
IOB Input Flip Flops	192	0	0	0	0
IOB Output Flip Flops	192	0	0	0	0
3-State Buffers	1248	118	9	235	18
3-State Half Longlines	96	40	41	78	81
Edge Decode Inputs	288	0	0	0	0
Edge Decode Half Longlines	32	0	0	0	0
CLB Fast Carry Logic	576	30	5	38	6

†Control Store Size = 128 Words

Table 5.1 CRAM Controller FPGA Device Utilization

The implemented ISA CRAM controller runs at 14 MHz, while the PCI CRAM controller (with a control size of 128 words) runs at 10 MHz. Table 5.2 shows the percentage FPGA utilization (in terms of CLBs and tristate buffers) of each functional block of the three CRAM controllers. Again this is for designs based on the simplified architecture of Figure 4.10 with a 192-word control store. Tristate buffer numbers are tabulated separately because in a Xilinx FPGA these logic gates are not implemented in a CLB. Notice that in each case the control store occupies more than 39% of the total area, and that the PCI interface unit uses almost six times more resources than the ISA interface.

Functional Block	CLB'S				TRISTATE BUFFERS			
	No Used		% of Total		No Used		% of Total	
	ISA	PCI	ISA	PCI	ISA	PCI	ISA	PCI
CRAM-Host Interface	17	99	3	15	36	137	24	50
Instruction Queue Unit	58	57	10	9	0	0	0	0
Microprogram Sequencer	43	42	8	6	40	40	27	14
Control Store	256	256	46	39	0	0	0	0
Buffers & Constant Unit	73	95	13	14	16	32	11	12
Registers, Address Unit, and Other Logic	108	109	20	17	56	65	38	24
Total	**555**	**658**	**100**	**100**	**148**	**274**	**100**	**100**

Table 5.2 CRAM Controller Functional Blocks FPGA Utilization

5.2.2 ASIC Simulation

Since this is the first work on the design and implementation of a CRAM system, there has been, and will continue to be a lot of iterations in arriving at an optimum architecture for the CRAM controller as well as the CRAM system software tools. For this reason, the implementation of a CRAM controller as an ASIC was not within the scope of this thesis. Otherwise, the main focus has been to refine the architecture by using programmable devices to build and test prototypes, and use simulation tools (developed in this work) to explore different architectural features. However, because of the universality of the design entry method (the same VHDL can be used for implementing the controller in an ASIC technology), the design of the controller using CMC-supported ASIC technologies has also been performed. The main objective of this exercise was to get a more accurate picture of the controller characteristics, especially speed and area. The technologies used are the Nortel 0.8 μm BiCMOS technology (BATMOS), Hewlett Packard 0.5 μm CMOS technology (CMOSIS5), and TSMC 0.35 μm CMOS technology (CMOSp35).

Table 5.3 shows the area of the controller using CMOSp35. The area is given in 2-input NAND gate equivalents so that it can be used to approximate the area of the controller in other technologies. For both the PCI and ISA, the control store occupies a substantial percentage of the total area of the CRAM controller. Another important thing to note is the area of the buffers. As shown in Figure 4.14, buffers are designed using small RAM blocks. In an ASIC technology, the area per bit of an SRAM core increases as the size of the core decreases. For example, a 256 x 32 SRAM core has 256 times the

Functional Block	PCI		ISA	
	†wand2 equiv. gates	% of Total Area	†wand2 equiv. gates	% of Total Area
CRAM-Host Interface	1184.2	9.6	105.0	0.9
Instruction Queue Unit	880.5	7.2	996.3	8.2
Microprogram Sequencer	663.3	5.4	665.5	5.5
Control Store	3898.7	31.7	6782.8	56.2
Buffers & Constant Unit	4458.5	36.3	2349.4	19.5
Registers, Address Unit, etc.	1201.4	9.8	1174.6	9.7
Total	**12281.6**	**100**	**12073.6**	**100**

†wand2 is a 2-input NAND gate.

Table 5.3 Area of CRAM Controller in TSMC 0.35 μm Technology

memory capacity of an 8 x 4 core, but its area in CMOSp35 (3893.7 equivalent gates) is only 34 times bigger. This makes the area of the buffers much bigger, especially since the buffers have to be designed using 8 x 4 SRAM cores because of the absence of 16 x 8 or 8 x 8 cores. This is also the reason why the ISA instruction queue unit is bigger than that of the PCI CRAM controller (two 16 x 16 cores have a total area of 562.4 gates, where as one 16 x 32 core has an area of 449.7 gates), and why the ISA control unit, which uses four 256 x 8 SRAM cores, has a bigger area than the 256 x 32 SRAM core used for the PCI.

The speed of the controller in CMOSp35 is well above the projected maximum CRAM speed of 50 MHz (20 ns cycle time). The cycle time is limited by the 11 ns (90 MHz) clock period of the SRAM cores. The maximum propagation delay in the other paths (without optimization) is 6.6 ns.

5.2.3 CRAM Controller Prototype Testing

Figure 5.3 shows the configuration used for testing CRAM controller prototypes. It consists of a CRAM controller board, a CRAM PCB, a Hewlett Parckard Logic Analysis System, a power supply, and a clock generator. The controller board is comprised of a Xilinx 240-pin XC4013E FPGA, an AMD27C512 512-bit EPROM to hold the FPGA configuration bitstream, and a MACH210A CPLD that controls the downloading of the controller configuration from the EPROM onto the FPGA. The CPLD implements a

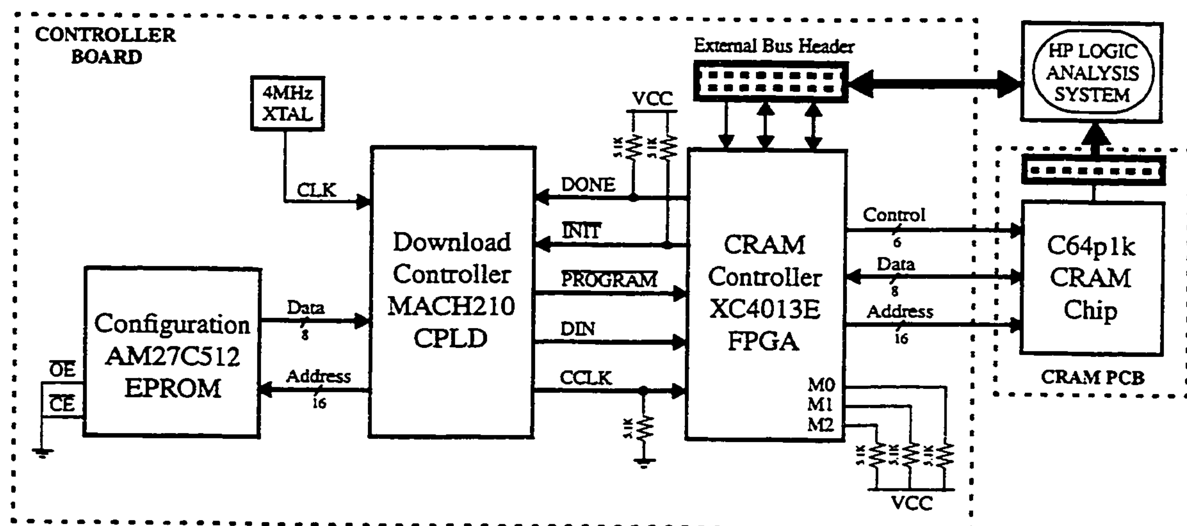

Figure 5.3 CRAM Controller Test Fixture

simple FSM that allows the use of standard byte-wide EPROMS instead of serial PROMS (which are more expensive and less widely available) during configuration. The CPLD logic was designed using the PALASM software. Details of Xilinx FPGA configuration can be found in the Xilinx devices data book [60]. The External Bus Header is a collection of test-point pins that are connected to the FPGA and represent the host bus signals. The number of these pins is big enough so that all the signals of the buses we are contemplating of using at some point (PCI, ISA, EISA, VME) can be accommodated. This makes the board a generic test fixture. All the main components of the controller board (FPGA, CPLD, and EPROM) are housed on sockets. The CRAM PCB consists of one socket for an 84-pin C64p1k CRAM chip and a few headers to allow tracing and driving of all the signals of the C64p1k. The first test fixture used a wire-wrapped controller board. But later, a CRAM system PCB comprising of the controller hardware and one socket for the C64p1k chip was fabricated. The layout of this latter board is shown in Appendix B. This is the board that is used in building the CRAM system prototype described in Section 5.3.

To test the controller, patterns of data are generated from the HP system to mimic specific host processor system bus cycles. This philosophy simplifies the testing because it mirrors exactly how the controller is used, connected, and tested in both the C++ and VHDL simulation models as well as the real CRAM system. Because of the absence of the C/C++ interface in this test fixture, CRAM instructions from the host processor (the HP pattern generator) are written in machine code. This machine code can be copied directly from the output of the CRAM Pre-compiler used in the VHDL simulation model (Figure 5.2), or from the debug output of the CRAM C++ Simulator. Both the host bus and CRAM bus signals are traced on the logic analyzer. Because of the simplicity of generating the test vectors and observing the test output, this test fixture has also been used to exhaustively test (memory and PE functionality) the C64p1k chips, something that was almost impossible to do with the chips tested in isolation.

5.3 ISA CRAM System Prototype

An ISA CRAM system, named CS64p1kISA, has been built. It comprises of one C64p1k CRAM chip, an ISA CRAM Controller FPGA, and the controller downloading CPLD and EPROM. It is made of two PCBs, the CRAM system PCB described in Section 5.2.3 (Appendix B) and a very small PCB that consists of an ISA male connector only. The ISA connector PCB is attached to the main PCB by screws, and its ISA signals are soldered to the main PCB External Bus Header using very thin wires. Figure 5.4 shows a photograph of the CS64p1kISA ISA PC card.

The system has been tested in a PC under both Linux and MSDOS. Because of problems with the software driver for the card, all routines tested on the prototype system were written in machine code (generated automatically using CRAM C++ simulator). Also, because of the small size of memory per PE and the small number of PEs, only small algorithms, especially basic arithmetic operations, were run on the system. The main objective of building the prototype was to demonstrate a working model of the CRAM concept rather than to highlight the performance advantages of CRAM. As mentioned earlier, for many applications, the performance advantage of CRAM becomes evident as the number of PEs increases. Therefore, the performance analysis of Chapter 7 is based on simulations of CRAM systems with a realistically large number of PEs.

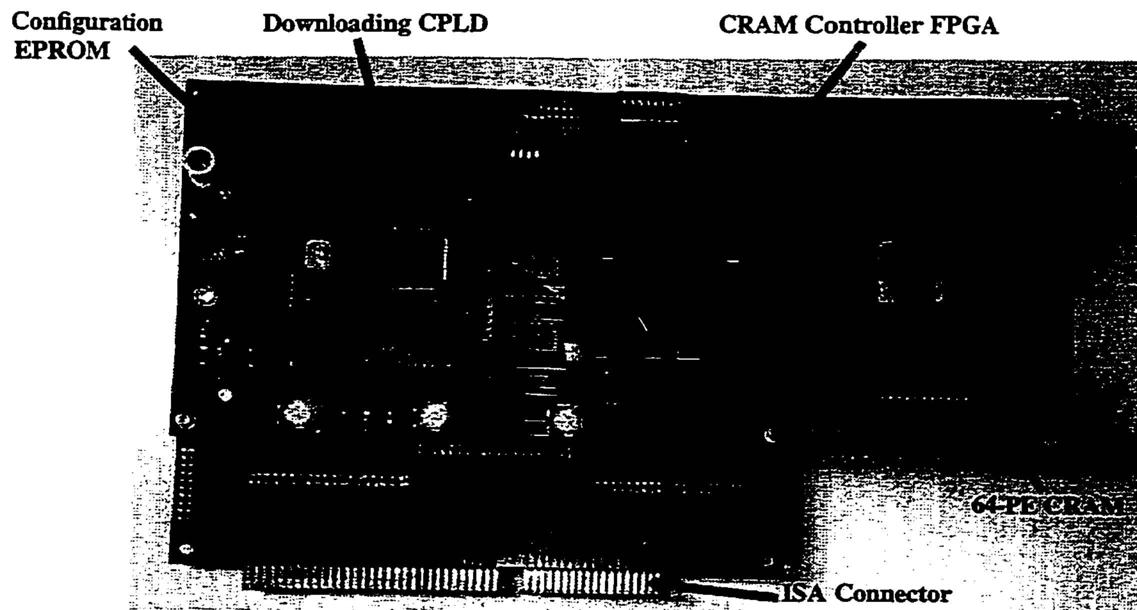

Figure 5.4 CRAM System ISA Card

5.4 Summary

A VHDL synthesis design flow was used in the design of the CRAM controller in order to reduce design time and provide a generic design that can easily be targeted for implementation in different technologies. A CRAM system VHDL simulation model has been designed to allow the running of actual system assembly or machine code when testing the functionality of the controller and CRAM chips. This reduces the time of generating input test vectors and collecting test output. The simulation model consists of VHDL models of the controller, the CRAM chips, a generic host processor, and host system buses, as well as a couple of C++ tools for processing text-file models of CRAM microinstructions and CRAM/host code.

Two controller prototypes have been implemented in a Xilinx XC4013EPQ240-2 FPGA. The ISA CRAM controller, with a 192-word control store, uses 564 CLBs and runs at 14 MHz. Because of area constraints, the PCI CRAM controller was implemented with a control store of 128 words. It uses all the 576 CLBs, and runs at 10 MHz. In TSMC 0.35 μm CMOS technology, the CRAM controller has an area of just over 12000 gates (2-input NAND gate equivalents) and runs at 90 MHz (limited by the cycle time of SRAM cores). To demonstrate a working model of the whole CRAM concept, a 64-PE ISA CRAM system prototype has been built and tested in a 133 MHz Pentium PC under both Linux and MSDOS.

Chapter 6

CRAM System Software Tools

This chapter describes the high-level software tools for application programming, software development, and system simulation. Section 6.2 describes the CRAM C++ Compiler, which is a CRAM C++ library that allows the use of the standard C++ language and standard C++ compilers when writing CRAM programs. Section 6.4 describes the CRAM assembly code, and Section 6.5 briefly describes a high-level tool for developing CRAM microcode. The CRAM C++ Simulator, a tool that is used to simulate the behavior of a CRAM system, is described in Section 6.7. This is used to analyze applications, software tools, and CRAM architectural features. Apart from software tools, this chapter also discusses two other software issues: data transposition (Section 6.3) and grouping of microroutines (Section 6.6). Data transposition is used to convert the format of data between the bit-serial CRAM and the bit-parallel host computer. Microroutine grouping is used to reduce the number of microinstructions in the control store so that a smaller microprogram memory can be used.

6.1 Introduction

The lowest level CRAM instructions are the control bits issued to the CRAM chip by the CRAM controller. These include the PE opcodes (COP and TTOP), as well as the control, data and address signals. These are stored as microinstructions in the CRAM controller control store, and are collectively referred to as the CRAM Machine Language.

Writing application programs using the CRAM machine code requires detailed knowledge of the architecture of the CRAM chip and the controller. Also, like any other machine code, it is both tedious and slow to program using these instructions. Therefore, high-level software tools have been developed that programmers and designers can use when writing CRAM applications or other software development tools. This also includes a C++ CRAM simulator that can be used by programmers for testing and timing CRAM applications. The simulator is also a vital tool for CRAM hardware and system designers in that it allows different architectural features to be explored before committing to the actual hardware implementation. Figure 6.1 shows the relationship and use of the CRAM software tools.

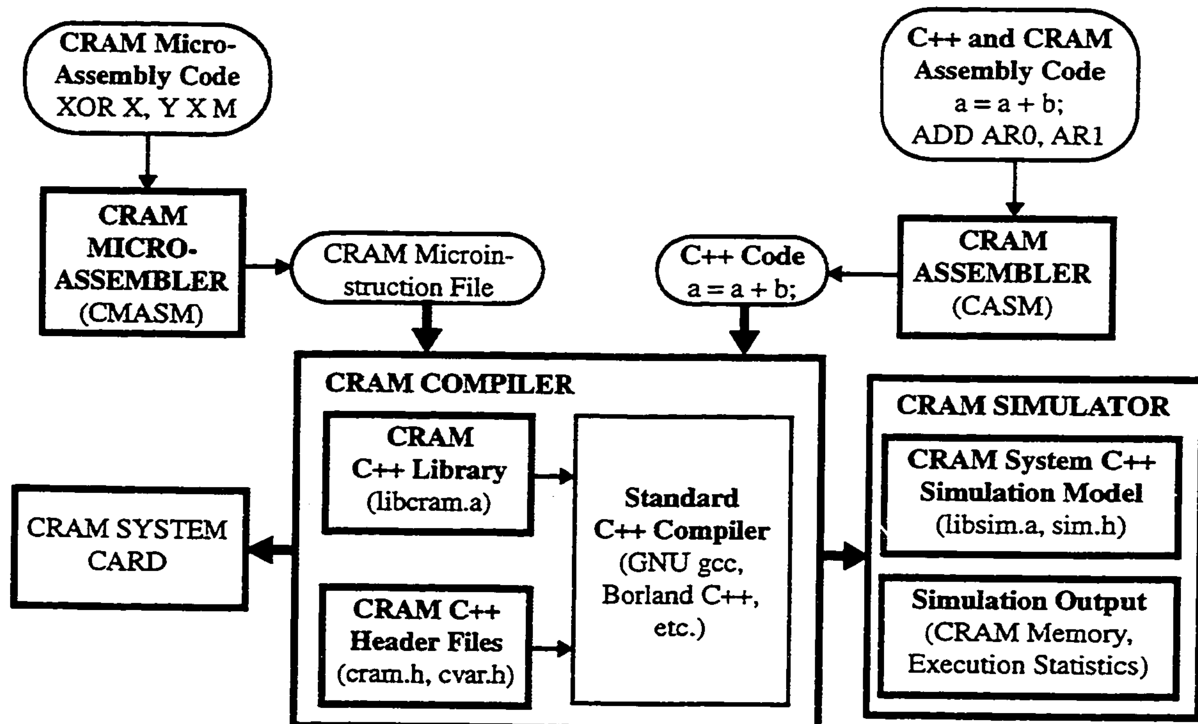

Figure 6.1 CRAM System Software Tools

CRAM applications can either be written in C++ code and compiled using a standard C++ compiler, or they can use both C++ and CRAM assembly code which is preprocessed through the CRAM Assembler before compilation. The CRAM Microassembler is a high-level tool for generating CRAM microcode. Note that the VHDL simulator described in Section 5.1.6 is used primarily for verification of the behavioral and synthesized VHDL controller design and is therefore not considered as a high-level tool for use by application programmers. It is therefore not included in Figure 6.1.

6.2 CRAM Compiler (CRAM C++ Library)

6.2.1 Using a C++ Compiler

One of the most negative attributes that SIMD machines are tagged with is that they are difficult to program. Because of this, designers of most well known SIMD machines have gone to great efforts to develop very elaborate software tools for their machines. There are three major approaches that can be followed. The first is to design the programming language for the machine from scratch. An example is the programming language for STARAN [43]. The second approach is to design a dialect of a common standard programming language such as C, Pascal or Fortran. This is a popular approach in most very high performance SIMD machines. Examples of such languages include the MasPar Fortran (MPF) and MasPar C (MPC) [44], the Connection Machine C*™ language [33], and the Terasys PIM data parallel bit C (dbc) [13]. The third approach is to use C++ as the programming language for the SIMD machine, with the addition of specialized libraries [4], [14].

The C++ programming language [45] allows creation of new classes of objects. These are actually new data types. The class constructor allows specialized initialization or operation when an object of a class is created. Such initialization might include allocating memory for the object. The class destructor is invoked when the object goes out of scope and is to be destroyed. This can be used in operations such as memory deallocation. The other important characteristic of C++ is operator overloading, which allows the use of standard C++ operators such as +, -, /, for any C++ class. By using operator overloading,

class objects can be used in expressions in exactly the same way as built-in data types. The third characteristic of C++, called inheritance, allows classes to be built from existing classes, thereby inheriting all the characteristics of the base class.

We have chosen to use C++ as the programming language for CRAM because we don't have to build a compiler, and hence the development cycle is short. It is also easier to upgrade since this only requires the addition of new classes or extension (adding new class functions, members, etc.) of existing classes. More important, C++ is easy to learn and use since many application programmers will already have used it before.

6.2.2 CRAM Classes (Data Types)

The CRAM C++ library contain classes that represent CRAM parallel integer variables. Because of the complexity of manipulating floating-point numbers in bit-serial format, and the fact that our current prototypes have very small memory per PE to consider using CRAM for floating-point operations, the design of CRAM C++ classes for parallel floating-point variables has been left for future work. Figure 6.2 shows the physical definitions and characteristics of CRAM data types. The following sections describe the CRAM classes.

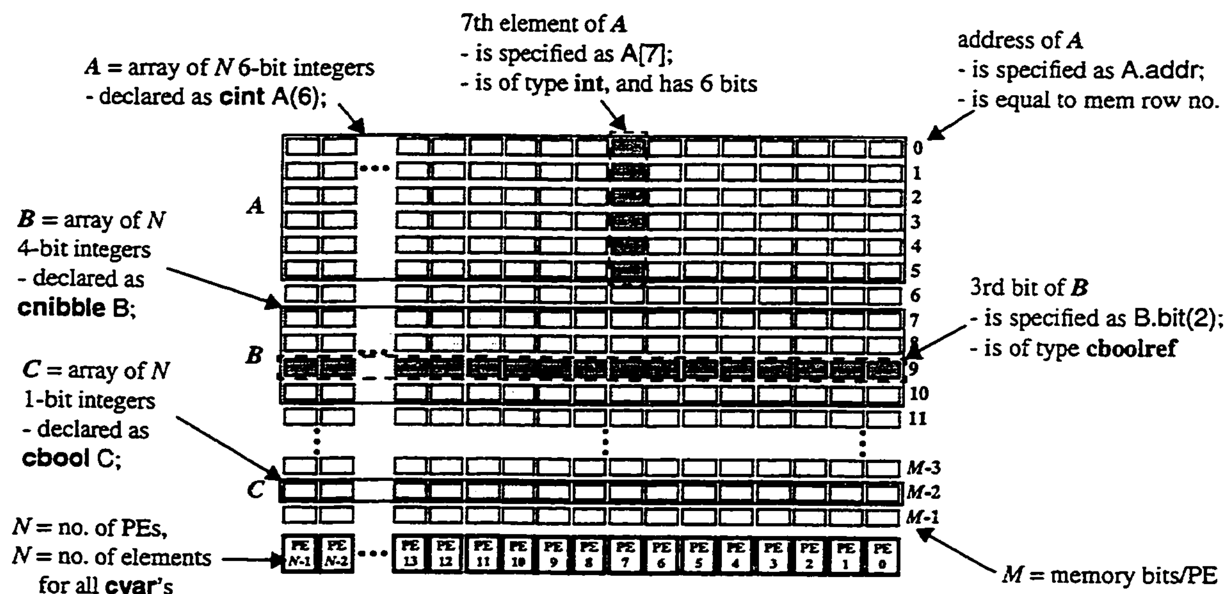

Figure 6.2 Definitions of CRAM Data Types

- **cvar:** The basic class from which all CRAM classes are derived is `cvar`. It represents an array of CRAM parallel integer variables. Its main member variables/functions include:
    - *bits* - The number of bits of the variable. This is limited by the size of memory/PE.
    - *addr* - The address of the variable in CRAM memory. This is assigned by the `cvar()` constructor when the variable is being created.
    - *type* - signed or unsigned integer.

  `cvar` also contains the functions for all the C++ overloaded operators and functions. These are described in Section 6.2.4 and Section 6.2.5.

- **cint** and **cuint:** These classes are descendants of `cvar` and represent signed and unsigned CRAM integer variables. `cint` and `cuint` only differ by *type* (`cint` is `CVAR_SIGNED` and `cuint` is `CVAR_UNSIGNED`). The size (*bits*) of the integer can be any positive number less than PE memory size. However, if *bits* is not specified, it defaults to `sizeof(int)` to mimic integer variables on the host computer system. This is usually 32 on most systems. The following shows the declaration syntax for `cint` and `cuint` variables.

  Declaration Syntax

```
cint a;                  // a is a 2-bit signed CRAM integer
cuint sum1(5);           // 5-bit unsigned
cuint error(INIT_0);     // 32-bit unsigned, initialize all its elements to zero
cint b(8, INIT_1);       // 8-bit signed interger, initialize all bits to 1 (0xFF)
```

- **cbool** and **cboolref:** `cbool` is a CRAM data-parallel boolean variable. Unlike on standard computer systems where boolean variables are usually implemented as 8-bit (`char`) variables, on CRAM, `cbool` is, as it should be, a 1-bit variable. Like `cint` and `cuint`, `cbool` is a descendant of `cvar`. It however has extra functions to handle logical operators used to combine relational operators (`||`, `&&`, and `!`). `cboolref` is a pointer to a bit of a `cvar` variable. It is not a physical object in CRAM memory, but it contains all the information about the address of a bit of `cvar`. This is an important class in optimizing operations involving `cbool` or 1-bit `cvar`, especially where there is no need to create a `cbool` variable. Otherwise, `cboolref` behaves exactly like `cbool`. The declaration syntax for CRAM boolean variables is shown below.

Declaration Syntax

```
cbool a;                  // a is a CRAM boolean variable
cbool a = 0;              // initialize a to 0 on declaration
cbool mask1(INIT_1);      // CRAM boolean, initialize to 1's
cuint d;                  // 32-bit CRAM unsigned integer
cboolef c(d.addr);        // c points to bit 0 of d
cboolef e = d.bit(3);     // e points to bit 3 (4th bit) of d
```

- **cchar, cshort, clong, cnibble:** These are fixed bit-length CRAM integer variables. They are also direct descendants of `cvar`. Like `cint`, they all have corresponding unsigned integer classes (`cuchar`, `cushort`, `culong`, and `cunibble`). The sizes (*bits*) of these classes are derived from the sizes of their corresponding C++ built-in data types. For example, the size of `cchar` is calculated from `CHAR_BIT*sizeof (unsigned char)`. On most machines, `cchar` would be 8-bit, `cshort` 16-bit, `cint` 32-bit, and `clong` would also be 32-bit. `cnibble` is a 4-bit CRAM integer variable. `cchar`, `cshort`, `clong` and `cnibble` have only one constructor with no parameters. In other words, you can not specify the number of vectors, nor can you initialize them on creation (you can always initialize them after they have been declared, without any increase in execution time). The first reason for having fixed-length classes is to allow programmers to declare variables in exactly the same way as is done for built-in C++ data types (no parameters in the constructor). The second reason, and the more important of the two, is that C++ does not allow declaring arrays if the class declaration (constructor) requires parameters. For example, the declaration `cuint a(8)[10]` is a syntax error, but `cuint a[10]` is correct. Therefore, where there is need to declare arrays, the programmer can use data types that do not require parameters. Since these classes cover only bit-lengths of 4, 8, 16, and 32, programmers can create any fixed bit-length CRAM integer class by using the declarations and functions of these classes. Declaration syntax for these variables is shown below:

Declaration Syntax

```
cchar a;                  // 8-bit signed cint
cushort diff = 4;         // 16-bit cuint, initialize all elements to 4
clong d[10];              // an array of 10 32-bit cint's
cnibble rec[8][12];       // a 2-dimensional array of 96 4-bit cint's
```

6.2.3 Memory Allocation

Any call to the `cvar` constructor invokes a CRAM memory allocation function. Since all classes of CRAM variables are descendants of `cvar`, the `cvar()` constructor is always called when objects of such classes are being created. In CRAM, variables are not required to be aligned to any particular address boundaries. There is one exception to this rule. A variable cannot be allocated memory that crosses over the 256*n* address boundary (n = 1, 2, 3,...). For example, a 6-bit `cint` variable cannot be allocated CRAM memory at CRAM row 254 because this will require its end address to be 259 (crossing over 256). This is so because the only part of the CRAM address that can be incremented or decremented by the controller to move through the bits of a CRAM variable is 8-bit. This is the part of the address that is stored in the AR0-AR1 registers (AXn registers are fixed).

When a CRAM variable goes out of scope, or when the `cvar` destructor (`~cvar()`) is explicitly called, the variable is destroyed and its memory is deallocated. Deallocated memory is returned to the CRAM free memory heap if the variable was at the end of the allocated memory. Otherwise it is returned to the fragmented garbage memory. A new variable is allocated memory in the fragmented garbage memory if it can fit in any memory fragment, else it is allocated in the continuous memory heap. If the memory is badly fragmented, i.e. so many small free unallocatable memory holes, the compiler will defragment the memory automatically. For the compiler to be able to do this, it keeps a list of pointers to all currently allocated CRAM variables. All memory allocation and deallocation issues are handled fully by the compiler.

6.2.4 Overloaded Operators

Most of the standard C++ operators have been overloaded for CRAM classes. These operators can be used with CRAM variables in exactly the same way as for standard C++ data types. Table 6.1 lists all operators that can be used with CRAM variables.

<table>
<tr><th>Operator
(@)</th><th>Description</th><th>Examples
(cint a, b; cuint c; cbool h; int i)</th></tr>
<tr><td>=</td><td>Assignment</td><td>a = b;</td></tr>
<tr><td>+
-
*
/
%
&
|
^</td><td>Addition
Subtraction
Multiplication
Division
Modulus
Bitwise AND
Bitwise OR
Bitwise XOR</td><td>a = a + b;
c = c % 6;
b = 7 & a;</td></tr>
<tr><td>+=
-=
*=
/=
%=
&=
|=
^=</td><td>Add and assign
Subtract and assign
Multiply and assign
Divide and assign
Modulus and assign
Bitwise AND and assign
Bitwise OR and assign
Bitwise XOR and assign</td><td>a += b;
c &= 0xFF;</td></tr>
<tr><td>++
--
-
~</td><td>Increment
Decrement
Negate
Complement</td><td>a ++;
b = --a;
a = -b;
c = ~a;</td></tr>
<tr><td><<
>></td><td>Shift left
Shift right</td><td>c = a << 5;
b = a >> i;</td></tr>
<tr><td><<=
>>=</td><td>Shift left and assign
Shift right and assign</td><td>c <<= 5;
b >>= 2;</td></tr>
<tr><td><
>
==
!=
<=
>=</td><td>Greater than
Less than
Equal to
Not equal to
Less than or equal to
Greater than or equal to</td><td>cif (a >= 6)
  a = b;
cend</td></tr>
<tr><td>||
&&
!</td><td>Logical OR
Logical AND
Logical NOT</td><td>cif ((b<c) && (!b))
  a = ++b;
cend</td></tr>
<tr><td>[]</td><td>Subscripting (ith element of cvar)</td><td>i = a[24];
b[3] = 6;</td></tr>
</table>

Table 6.1 CRAM C++ Operators

6.2.5 Library Functions

Table 6.2 lists functions provided in the CRAM C++ library for CRAM data types. These include functions for bit-slicing, bit-setting, minimum and maximum searches, and PE shifts. Note that PE shift functions (`PE.shiftl()` and `PE.shiftr()`) are used to shift the elements of a CRAM variable, whereas standard shift operators (<< and >>) operate on CRAM variables in exactly the same way as they do on standard integer variables, i.e. the bits of the elements are shifted left or right (equivalent to multiplying or dividing by 2^n).

Function(s)	Description	Examples (`cuchar c; cint a, b; cbool h; int y`)
`set()`, `ones()`	Sets all bits of `cvar` to '1'	`c.set();` // c[PE] = 0xFF
`reset()`, `clear()`, `zero()`	Sets all bits of `cvar` to '0'	`c.clear();` // c[PE] = 0x00; equivalent to c = 0
`bit()`	Returns *i*th bit (`cbool`) of `cvar`	`a.bit(5) = h;` // set bit 5 of a to the value of h
`from()`	Returns bit-slice (`cvar`) of `cvar`	`c = a.from(4,11);` // c is set to bits 11:4 of a
`search_bitset()`	Returns the PE number with bit of `cbool` set to '1'	`y = search_bitset(h);` // y is set to a value such that h[y] = '1'
`ismax()` `ismin()`	PE has max(min)imum element of `cvar`	`h = ismax(a);` // h[PE]=1 if a[PE] is the maximum element of a
`max()` `min()`	Returns max(min)imum of two `cvar`'s	`a = min(a, b);` // a[PE] = (a[PE] < b[PE]) ? a[PE] : b[PE]
`maxindex()` `minindex()`	Returns index of PE with max(min)imum element	`y = maxindex(a);` // yth PE contains maximum element of a
`maxele()` `minele()`	Returns the max(min)imum element of `cvar`	`y = minele(a);` // does int x=minindex(a), then reads a[x]
`abs()`	Returns absolute value of `cvar`	`a = abs(a);` // a = (a < 0) ? -a : a
`shiftr()` `shiftl()` `rotater()` `rotatel()`	Shift PE elements left/right (`rotatex` connects edge PEs (PE_0 and PE_{N-1}] when shifting)	`PE.shiftr(a, b, 4, 1);` // a[PE] = b[PE+4], fill edge PEs with '1' `PE.rotatel(a,b);` // a[PE] = b[PE-1], a[0] = a[N-1]

Table 6.2 CRAM Library Functions

6.2.6 Data-Parallel Conditional Statements

- Data-parallel conditional execution if supported by the `cif` statement proposed by Elliott [4]. This is similar to the `ifarray` construct used in STARAN [43]. The syntax of the cif statement is shown below:

Syntax

```
cif (cbool expression)
{
  ctstatements ...
[} celse {
  cfstatements ...
}]
cend
```

Example

```
cif (a > b)
  c += a;
celse
  c += b;
cend
```

cbool expression is any expression that returns a `cbool` object. The resulting `cbool` value has a 1 where the PE evaluated the expression to be true, and a 0 where the expression evaluates to false. This `cbool` value is then written to the PEs write enable (WE) registers to enable or disable the PEs from writing to their own memory. One key difference between `cif` and the C++ `if` constructs is that in the `cif`, all the PEs will execute the statements in both the `cif` (*ctstatements*) and `celse` (*cfstatements*) branches regardless of whether the *cbool expression* evaluates to false or true for a particular PE. The conditionality of the `cif` construct lies in the fact that the PE will only write the results of the executed statements to its memory if the condition evaluated to be true. Therefore, a useful statement inside a `cif/celse` will be one that assigns its result to a `cvar` (variable in CRAM memory), otherwise the statement will be executed unconditionally. Also, unlike the C++ `if`, `cif` cannot be used for say code execution speedup since all statements inside a `cif` or `celse` are always executed.

In order to support nested `cif`'s, a stack of `cbool` variables is maintained, with the top of the stack containing the current value of the write enable registers. This stack is automatically managed by the constructor and destructor of the `write_en` CRAM class. When a `cif` is executed, the result of the `cbool` expression is ANDed with the top of the stack (WE values), and the result written back to the WE registers and the top of the stack. The statements inside the `cif` are then executed. After this, if `celse` is encountered, and the `cbool` value next to the top of the stack is true, then the contents of the WE registers

and the top of the stack are inverted. After this, the statements inside the `celse` are executed. Notice that `celse` is optional. Every `cif`, or `cif/celse` pair, must be terminated with a `cend`. When `cend` is executed, the top of the stack is popped out and the write enable registers are restored to their values before the conditional execution.

`cif`, `celse` and `cend` are all defined as preprocessor macros. The curly brackets surrounding the statements inside a `cif` or `celse` are optional, but it is advisable to include them if there is more than one statement inside a conditional branch.

6.2.7 Operations with Scalar Constants

Operations that have a `cvar` and an integer value as operands (except << and >>) use the CRAM controller constant broadcast unit (Section 4.5.2). Instead of loading the integer into a temporary `cvar` in CRAM memory and doing the operation on two `cvar` variables, the integer is loaded into the controller write buffer and the operate-immediate instruction (such as `ADDI`, `SUBI`, etc.) is executed. Thus, the load-constant instruction (`LDK`), which has 2*bits microinstructions, is avoided, and no extra CRAM memory is required to store the integer constant. Also, the constants can be pre-loaded into the write buffer and operations performed directly with the constants in the buffer. The advantages of this are outlined in Section 4.5.2.

The CRAM C++ compiler performs some optimizations for operations with certain integer constants. This optimization is done by substituting the operation with an operation which results in fewer microinstructions. Table 6.3 shows all optimized operations. The compiler also optimizes multiplication and division by 2^n ($n = 0, 1, 2, \ldots$) by replacing them with shift operations (<< and >>, respectively). However, to reduce the overhead of testing against all the possible constants, this optimization is only done for the constants 2, 4, 8, 16, 32 and 64. The programmer can perform the optimizations for the rest of the constants either at compile time (if they are known in advance) or at runtime (by testing against all such constants of interest).

Intended Operation	Optimized Operation	Microinstructions Saved(n=bits)
`a = b + 0 ⇒ ADDI a, b, #0)` `a = b - 0 ⇒ SUBI a, b, #0`	`a = b ⇒ MOV a, b`	$3n+1$

`a = b | 0 ⇒ ORI a, b, #0`
`a = b ^ 0 ⇒ XORI a, b, #0`
`a = b & 1 ⇒ ANDI a, b, #1`
	`a = b ⇒ MOV a, b`	$2n$

`a = b & 0 ⇒ ANDI a, b, #0`	`a = 0 ⇒ CLR a`	$4n-1$
`a = b | -1 ⇒ ORI a, b, #-1`	`a = 1 ⇒ SET a`	$4n-1$
`a = b ^ -1 ⇒ XORI a, b, #-1`	a = $\overline{b}$ ⇒ NOT b	$2n$

`a = b | 1 ⇒ ORI a, b, #1`	`a(0) = 1 ⇒ SET a(0)`
	`a(n-1:1) = b(n-1:1) ⇒`
	`MOV a(n-1:1), b(n-1:1)`
		$2n+1$

`a = b & 1 ⇒ ANDI a, b, #1`	`a(0) = b(0) ⇒ MOV a(0), b(0)` `a(n-1:1) = 0 ⇒ CLR a(n-1:1)`	$4n-3$

`a = b ^ 1 ⇒ ORI a, b, #1`	a(0) = $\overline{b(0)}$ ⇒ NOT a(0), b(0)
	`a(n-1:1) = b(n-1:1) ⇒`
	`MOV a(n-1:1), b(n-1:1)`
		$2n$

`a = b + 1 ⇒ ADDI a, b, #1` `a = b - (-1) ⇒ SUBI a, b, #-1`	`a = b + 1 ⇒ INC a, b`	$2n-1$
`a = b - 1 ⇒ SUBI a, b, #1` `a = b + (-1) ⇒ ADDI a, b, #-1`	`a = b -1 ⇒ DEC a, b`	$2n-1$
`a = 0 - b ⇒ SUBI a, #0, b`	`a = -b ⇒ NEG a, b`	$2n-1$
`a = -1 - b ⇒ SUBI a, #-1, b`	a = -1 - b ⇒ †MINUS a, b	$2n-1$
`a = 0 ⇒ MVI a, #0`	`a = 0 ⇒ CLR a`	‡$n-1$
`a = -1 ⇒ MVI a, #-1`	`a = -1 ⇒ SET a`	‡$n-1$
`a = 1 ⇒ MVI a, #1`	`a(0) = 1 ⇒ SET a(0)` `a(n-1:1) = 0 ⇒ CLR a(n-1:1)`	‡$n-1$
`a = b * 0 ⇒ MLTI a, b, #0`	`a = 0 ⇒ CLR a`	‡$6n^2+3n-1$
`a = b * 1 ⇒ MLTI a, b, #1`	`a = b ⇒ MOV a, b`	‡$6n^2+n$
#a = b * 2^k ⇒ MLTI a, b, #2^k	`a = b << k ⇒ SL a, b, #k`	‡$6n^2-1$
`a = b / 1 ⇒ DIVI a, b, #1`	`a = b ⇒ MOV a, b`	‡$16n^2+36n+2$
#a = b / 2^k ⇒ DIVI a, b, #2^k	`a = >> k ⇒ SR a, b, #k`	‡$16n^2+35n+1$

†Special type of NEG with Y initially set to 1.
‡Plus the time for loading the constant from the host into the controller write buffer.
#k = 1, 2, 3, 4, 5, 6.

Table 6.3 Optimizations for Operations with Constants

6.2.8 Operand Extension

Operand extension is required if source operands have different number of bits, or if the number of bits of one operand is less than that of the destination operand. In CRAM, rather than physically extending the operands by creating temporary variables, each overloaded C++ operator has a CRAM instruction that is executed if a specific operand extension is required. This increases the speed of execution and also removes the need for using extra CRAM memory. For example, consider the execution sequence of an expression `c = a + b`, where `a` is 8-bit, `b` is 4-bit, and `c` is 16-bit. Firstly, the ADD instruction is executed for the first 4 bits. Then `ADD1` is executed for the next 4 bits to extend `b` to the size of `a`, and finally `ADD0` is executed to extend `a` and `b` to the size of `c`. In other words, `ADD1` only has to ripple carries, and `ADD0` is just sign extension. These instructions give a saving of $2n$ and $4n$ cycles, respectively, when executed for n bits instead of using the full ADD instruction. CRAM operand extension instructions have very few microinstructions and hence have been implemented for all the main operators.

6.2.9 Support Classes

Apart from the CRAM variable classes described in the previous sections, the CRAM C++ library also has classes for objects of the CRAM controller, CRAM PEs, and the host processor. These are briefly described below.

- **Controller Objects:** The CRAM controller class contains all information about the controller that is necessary to the CRAM C++ compiler. This includes the memory map of the units, the size (in bytes) of the instruction FIFO, the read/write buffers, and the control store, and information about controller registers and other units. This information is provided in form of controller sub-classes and instantiation of their objects in the controller class object. This approach is advantageous because it mimics the hardware architectural composition of the controller. Therefore specific functions can be implemented for the class of an individual controller unit or register, in order to model its behavior and keep track of its status during instruction execution. This makes it easier for the compiler to perform code execution optimization that is dependent on the architecture of the controller. Controller sub-classes include classes for the

instruction FIFO, read/write buffers, control store, command/status registers and their individual bits, and all user-accessible parameter registers. Typically, a class of a controller register includes members to store the register address and value. Member functions include those to implement register assignment (e.g. `WLEN = 7`), and register increment/decrement (e.g. `AR0++`). Common member functions for classes of the other units (FIFO, buffers, control store) are functions for reading and writing data. As mentioned earlier, these sub-classes and their functions allow the use of standard C++ constructs when assigning, incrementing, decrementing, reading and writing controller registers and units. Also, by having controller hardware units as objects in the C++ library enables the compiler to have an intelligent and updated information about the status of the relevant units. This is used in code optimization, for example by avoiding unnecessary updating of a register when the register already contains the value that it is supposed to be updated with.

- **Processing Element (PE) Objects:** The CRAM PE class contains information about, and functions to manipulate PE components. This includes the instantiations of CRAM registers (X, Y, M, W) and their groupings (such as XY, XWY), and functions for register assignment and PE shift operations. Like the controller classes, PE classes simplify the use of PE objects in the compiler or application source code. For example, to set all the PE X registers to zero, instead of issuing a native CRAM instruction, one may code it in normal C++ syntax using `PE.X = 0`. A number of PE assignment operators have been included to implement common PE register operations such as initialization (`PE.W = 1`), register-to-register copy (`PE.X = PE.Y`), and cvar-to-PE-register copy (`PE.W = a.bit(8)`).

- **Host Computer Objects:** Host computer classes include the software transposing buffers (Section 6.3) and CRAM instruction types. The instruction types are used for specific instruction characterization during execution. Examples include a one `cvar` source operand instruction (`one_addr_instr`), two `cvar` source operands instruction (`two_addr_instr`), and one `cvar` and one integer constant instruction (`addr_cont_instr`).

6.2.10 CRAM System Objects and Their Initialization

When an application compiled with the CRAM C++ library is run, a number of CRAM system objects are created and initialized. The compiler initializes these objects to the initial values of the real hardware units they represent.These objects include:

- **PE** - This represents the components of CRAM processing elements described in Section 6.2.9. Therefore, in the application code, all references to CRAM PE registers should be qualified with PE, e.g. `PE.X` to denote all PE X registers.
- **controller** - This object denotes the CRAM controller, and hence must qualify all reference to controller resources, e.g. `controller.WLEN`.
- **host** - Represents the host processor. The software transposing buffers (Section 6.3) are defined in the host object (`host.read_buffer` and `host.write_buffer`).
- **ALL_ZEROES_MASK** - This is of type `cbool` and is initialized to a value of all zeroes. It is located at CRAM memory row 0.
- **ALL_ONES_MASK** - This is of type `cbool` and is initialized to a value of all ones. It is located at CRAM memory row 1.
- **MASK_01** - This is of type `cbool` and is initialized to the value 0x01. It is located at CRAM memory row 2.

6.3 Corner-Turning: Host Access of CRAM Variables

Since CRAM operates on data in bit-serial format while the host computer is a bit-parallel system, data must be converted from bit-serial to bit-parallel and vice-versa when it is transferred between the two systems. This is referred to as either corner-turning, data transposing, or format conversion. Typically, data is transposed using multidimensional access memory (MDA) and a flip or shuffle network like the ones used in STARAN [46], [47], or the MIT image processor [14]. These methods require a lot of hardware, especially if the system is general-purpose and its operands are not fixed to a particular size (bits). For example, even for such an application-specific processor as the MIT pixel-parallel image-processing system [14], whose data path is designed for 8-bit gray scale pixels, sixteen 64Kb x 4 SRAMs, plus shuffler, address and other glue logic, are used to transpose the data between the host and the PE array [41]

Because of the high hardware cost and complexity of hardware corner-turning, data on a CRAM system is transposed in software. Traditionally, software data transposition is done on the host computer. This is slower than a hardware transposer, and may severely degrade performance especially if the number of elements to be transposed in a parallel variable is large. But with current increasing processor speeds, host-based data transposition offers a good alternative for low-cost system implementation. An example of a logic-in-memory system that uses software data transposing is the Terasys PIM workstation, which transposes data using its Sparc-2 host processor [13].

To reduce the overhead of transposing a large number of elements of a CRAM variable, a parallel array-based corner-turning approach that exploits the large degree of parallelism of the PE array and the 1-D inter-PE communication network has been developed. A 6-bit Buffer Address Increment register has been included in the CRAM controller to allow faster corner-turning of multi-byte data. The following sections describe these two corner-turning approaches.

6.3.1 Host-Based Data Transposition

Host-based data transposition is used if the number of elements to be accessed by the host is small. This is supported in the CRAM C++ library through the subscripting operator ([ ]). This operator has been overloaded for `cvar` in order to support constructs similar to those used by C++ in accessing elements of arrays for built-in data types. Any element of `cvar` can be addressed using the `[ ]` operator. For example, to initialize the 5th element of a `cint` variable `x` with an integer value of 6, the expression `x[5] = 6` is used. Similarly, the *i*th element of `x` can be assigned to an integer variable `y` using the expression `y = x[i]`. Figure 6.3 shows how integer values are written to `cvar` elements using this method. The data is first corner-turned through the host software write buffer (using shifting and masking), then written to the CRAM controller write buffer, and finally transferred to the CRAM memory using a `WRITE` instruction. If CRAM is a byte-wide memory, for the values to be transferred using the `WRITE` instruction, eight consecutive `cvar` elements must be loaded at the same time. If cvar elements are assigned in a non-consecutive manner (e.g. the expression `a[0] = 2` is followed by `a[3] = 7`), or if fewer than eight consecutive elements are assigned to, or if neither the first nor the last of the eight elements is on the $8n^{th}$ PE (n=0, 1, 2, ...), the values are written to the PEs memory as scalar constants. This is much slower because for an element to be initialized

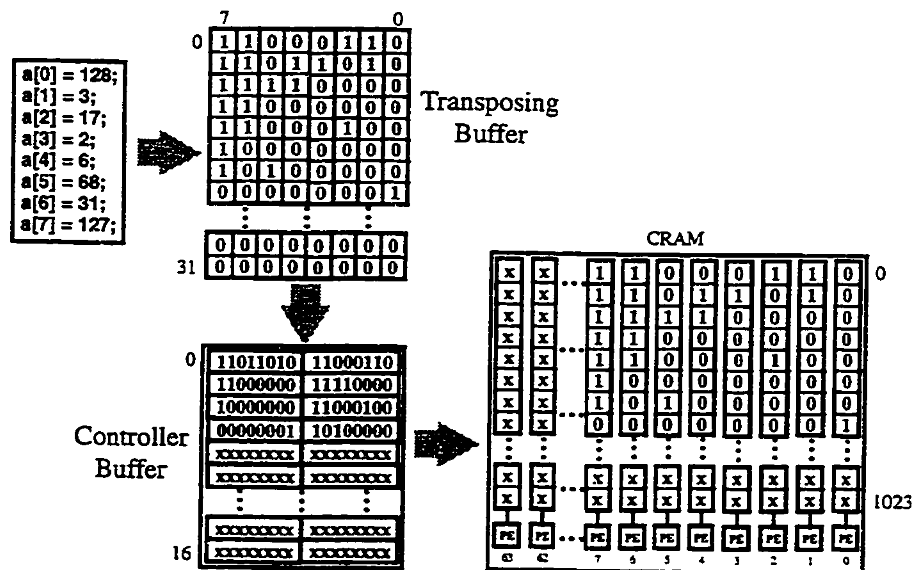

Figure 6.3 Host Corner-Turning

as a constant, the other PEs are first masked off by writing 0's to their write enable registers. After that, the constant is loaded into the controller write buffer, and finally loaded to all `cvar` elements using a `LDK` (load-constant) instruction. The WE registers are then restored to their previous values. If more than one element is to be loaded, this process is repeated for each element. Therefore to speed up the initialization of `cvar` elements, an effort must be made to initialize them consecutively. This can be achieved simply by initializing the elements using a `for` loop, with the initial value of the index starting at the $8n^{th}$ PE. To complement this, a `cvar load_elements()` function is provided which takes as arguments a pointer to the integer array, the size of the elements in the integer array, the `cvar` variable, and the number of `cvar` elements to be initialized. In both of these methods, the CRAM compiler automatically uses the transpose-WRITE method unless the number of `cvar` elements to be initialized is less than eight.

All elements of a `cvar` can be loaded with a uniform value by making the assignment without the `[]` operator. For example, for a `cint` variable `y`, all its elements can be loaded with 7 by using the expression `y = 7`. This uses a load-constant instruction with all WE registers enabled (by default, WE registers are always enabled). This is faster than the transpose-WRITE method since it requires only one load (to load the value into the controller write buffer) and one CRAM instruction (`LDK`).

Reading the value of a `cvar` element is almost the reverse of the write process described above. If CRAM is a byte-wide memory, eight `cvar` elements, including the one to be read, are read from CRAM to the controller read buffer using the `READ` instruction. These values are then transferred to the host read buffer, where the required value is transposed and returned to the assignment statement. Since eight elements were read, any subsequent assignment to the elements whose values are in the host read buffer will use these values without the need to read them from the CRAM chip. This is done until the elements in the read buffer are marked 'dirty', which happens when the `cvar` variable is used as a destination operand in an instruction.

6.3.2 Parallel Array-Based Data Transposition

To transpose the elements of an n-bit CRAM variable on the PE array itself, the data is loaded onto the CRAM array in the normal host (parallel) format, with n elements spread across n PEs. The data is then transposed in groups of n PEs by rotating the n x n bits. Typically, the minimum value of n is 8 since the byte is the smallest storage unit on most standard computers. Figure 6.4 illustrates array-based data transposition. For clarity, only

```
---------------------------------------
-- TXUP: Transpose from top          --
--   AR0 - Addr of LSB of source data --
--   AR1 - Addr of LSB of tx'd data   --
-- TXDN: Transpose from down         --
--   AR0 - Addr of MSB of source data --
--   AR1 - Addr of MSB of tx'd data   --
---------------------------------------
TXUP:
CREAD AR1;
LOAD RY, M;
CWRITE AR0, INCA, SETLOOP;
LOAD RY, Y;
CWRITE AR0, INCA, LOOPN, IF LPCNZ;
LOAD LX, X;
LOAD W, X, END;
--
TXDN:
CREAD AR1;
LOAD LX, M;
CWRITE AR0, DECA, SETLOOP;
LOAD LX, X;
CWRITE AR0, DECA, LOOPN, IF LPCNZ;
LOAD RY, Y;
LOAD W, Y, END;

/*------------------------------------*/
/* cram_transpose()                   */
/* CRAM-based parallel transposition  */
/* of 8-bit integers;                 */
/*------------------------------------*/
void cram_transpose (address& src,
                     address& dst)
{
  int i;

  // load the mask into regs X and W
  execute_instr (LDWX, MASK_01.addr, 0);

  // transpose from up: both addresses
  // increase after each iteration
  for (i=7; i>0; i--)
  {
    execute_instr (TXUP, dst, src, i);
    dst ++; src ++;
  }

  // load the mask from X regsiter to Y
  execute_instr (TXY, 0);

  // transpose from down
  for (i=7; i>0; i--)
  {
    execute_instr (TXDN, dst, src, i);
    dst --; src --;
  }
}
```

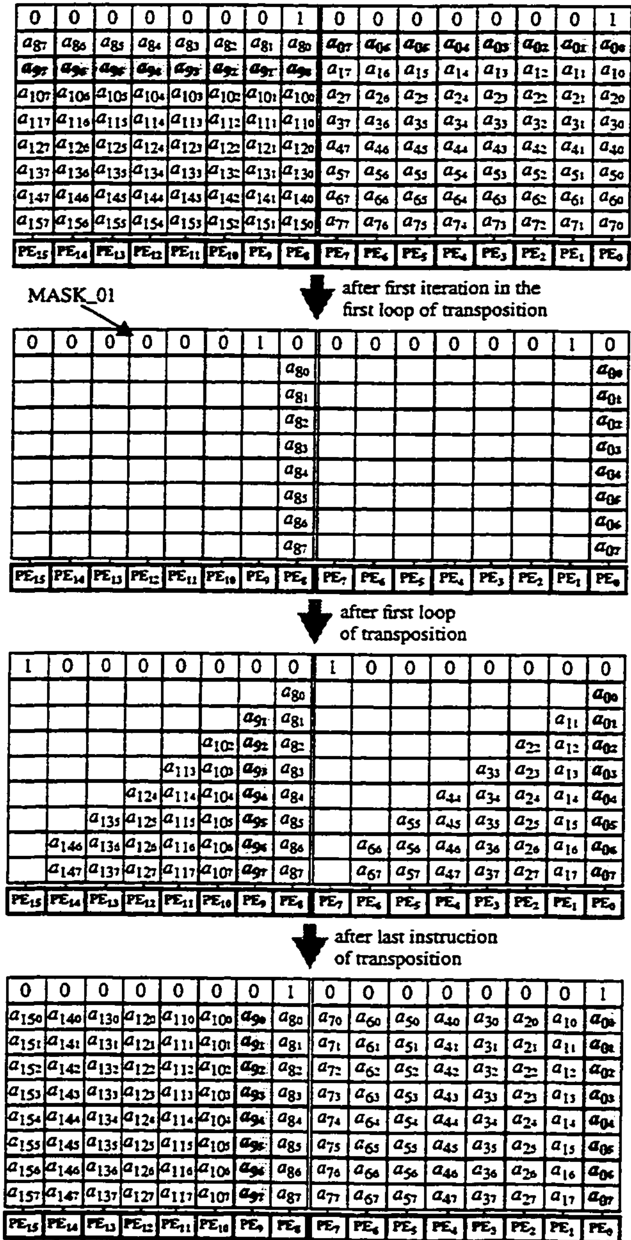

Figure 6.4 Parallel Array-Based Data Transposition

the first 16 elements of an 8-bit CRAM variable are shown. The exact steps are described in the actual transpose microroutines (`TXUP` and `TXDN`) and the C++ function (`cram_transpose()`).

If i is the number of the PE in the n-PE group, and j is the memory row of the PE, then data transposition is performed by moving the bits in the n x n window such that

$$a'_{ij} = a_{ji},\ 0 \le i, j < n \quad (6.1)$$

where a' is the transposed version of a. This is a symmetrical transformation and hence to transform data from CRAM format to host format, exactly the same operation is performed.

The `TXUP` and `TXDN` microroutines each executes $(5 + 2i)$ microinstructions, where i is the loop count variable in the transpose loops of the C++ `cram_transpose()` function. Each of these loops is executed $(n-1)$ times (n is the number of bits of the CRAM variable). Also note that before the transpose microroutine is executed, the WLEN register is first set to i. This makes the number of microinstructions executed for each transpose microroutine call equal to $(6 + 2i)$. Therefore, the total number of microinstructions executed in each of the transpose loops in `cram_transpose()` is equal to

$$M_{tx} = 6(n-1) + 2\sum_{i=1}^{n-1} i \quad (6.2)$$

For a CRAM system of cycle time T_c, the total time to transpose data, T_x, is equal to the time to execute the $2M_{tx}$ transposition microinstructions and six initialization microinstructions. The initialization instructions are used to update address extensions for `MASK_01`, `src` and `dst` operands, plus two microinstructions for `LDWX`, and one for `TXY`. Therefore, T_x is given by

$$T_x = \left\{2\left[6(n-1) + 2\sum_{i=1}^{n-1} i\right] + 6\right\}T_c \quad (6.3)$$

Simplifying Equation 6.3 gives

$$T_x = 2(n^2 + 5n - 3)T_c \quad (6.4)$$

Like all array-based PE operations, this time is constant for all numbers of elements equal

to or less than the number of PEs. Therefore, as the number of PE increases, the equivalent time to transpose an element decreases because of the parallelism in the PE array. For an 8-bit CRAM variable, the number of microinstructions executed in Equation 6.4 is equal to 202, and the transpose time for a 50 ns CRAM system is 10.1 µs. Figure 6.5 shows the performance of transposing a varying number of elements of an 8-bit CRAM variable using either the PE array of a 50 ns 64K-PE PCI CRAM system or a 133 MHz Pentium host PC. It is evident from the figure that host-corner turning is beneficial only if a very small number of elements (less than 8) is to be accessed.

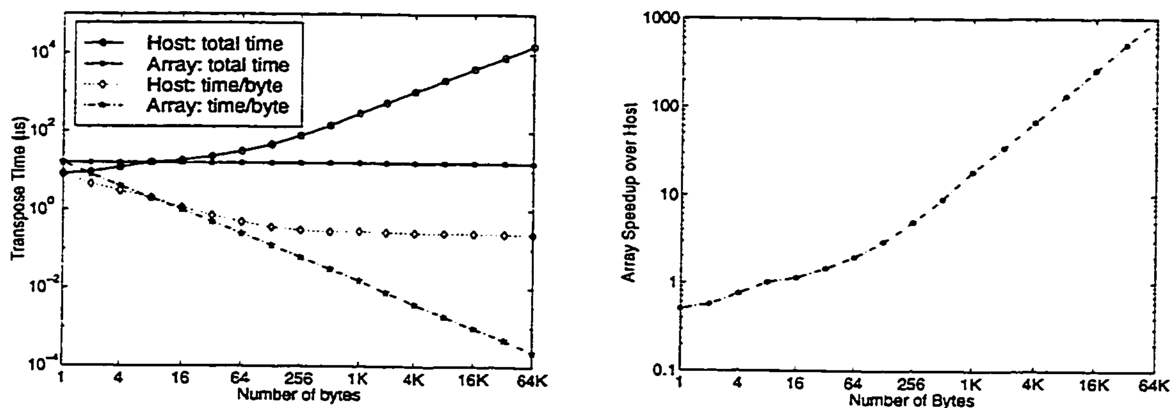

Figure 6.5 Array-Based vs. Host-Based Transposition

Transposing N-byte Variables

The transpose time for an n-bit variable has a complexity of $O(n^2)$. Therefore the transpose time for 16-bit and 32-bit variables would be almost 4 times and 16 times, respectively, the transpose time of 8-bit variables. More precisely, using Equation 6.4, the transpose time for a 16-bit and 32-bit variable is 3.3 times and 11.7 times, respectively, the transpose time of an 8-bit CRAM variable.

To reduce the time of transposing multi-byte variables, a procedure that has a complexity of $O(n)$ has been developed. This involves loading corresponding bytes of 8 consecutive elements in 8 consecutive memory rows, and then transposing them as 8-bit variables. Figure 6.6 illustrates this for a 16-bit variable. There is no extra overhead for this data arrangement because it is supported in the CRAM controller by simply loading the Buffer Address Increment register with N prior to issuing the `WRITE` instruction.

Using this approach, the transpose time for an N-byte CRAM variable is only N times that of a 1-byte (8-bit) CRAM variable.

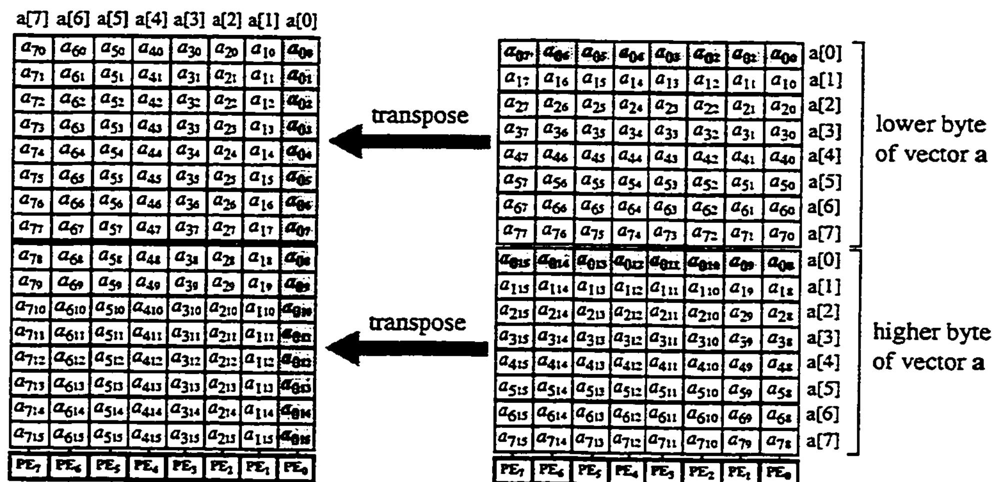

Figure 6.6 Transposing N-Byte Variables

6.3.3 Effect on Performance

The performance of data transposition is usually measured relative to the I/O overhead of transferring data between the bit-serial system and its host. This is so because in most cases the need to transpose data arises because data on the host has to be used by the bit-serial system and vice versa. An ideal data transposer adds nothing to the I/O overhead.

Figure 6.7 shows the data transpose time as a percentage of the total I/O overhead (data transposing + data transfer). Again, this is for a 50 ns CRAM system interfaced through a PCI bus to a 133 MHz Pentium PC. If the number of bytes to be transferred and transposed is small, both corner-turning approaches constitutes a very high percentage (more than 80%) of the total I/O overhead. But as the number of bytes increases, the percentage contribution of array-based transposition decreases, while that of host transposition remains almost constant. In other words, data transfer time increases in both cases, data transposition time when using the parallelism of the PE array remains constant, while the time of transposing data on the host also increases with the number of bytes. Typically, a CRAM system would have more than 4K PEs. This means that parallel CRAM variables would have 4K elements or more. For this case, data transposition using the PE array contributes less than 10% to the total I/O overhead. This number is as low as

0.5% for a 64K-PE PCI CRAM system. Note that the percentage contribution of data transposition is even smaller for slower buses because of the higher data transfer time. Also, when calculated as a percentage of the total execution time of an application (I/O overhead + code execution time), the effect of data transposition drops even further.

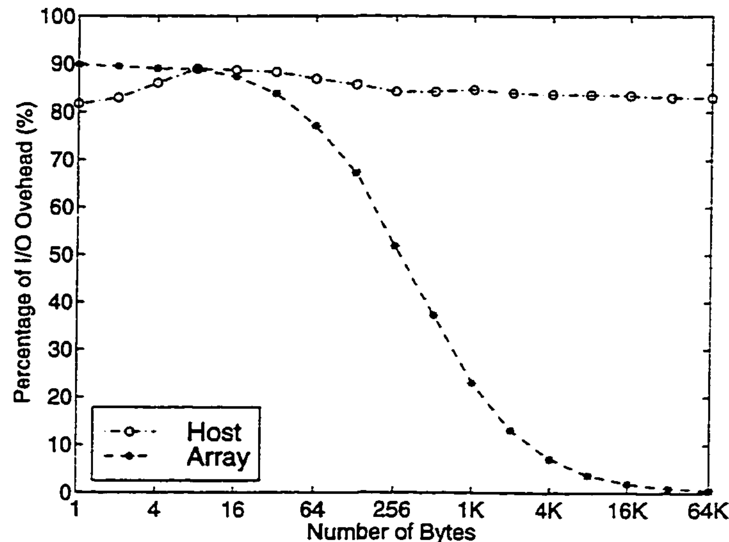

Figure 6.7 Data Transpose Time as Percentage of Total I/O Overhead

6.4 CRAM Assembler

6.4.1 Introduction

For applications that require explicit control of the raw hardware, applications can be written directly using the CRAM Controller Assembly Language (also known simply as the CRAM Assembly Language). This assembly code can either be mixed with the CRAM C++ code and tagged with the keywords CASM and END CASM, or it can written as separate code. If mixed with C++ code, the CRAM assembler (CASM) is used as a pre-processor before compiling the mixed code with a C++ compiler. In this case, CASM simply converts the assembly code into a C++ function to transfer the instruction to the controller instruction FIFO.

As with most assembly languages, it is very unlikely that a typical application programmer will ever use CRAM assembly code, except in situations where assembly

code offers considerable speed-up when compared to C++ code. Otherwise, CASM is a vital software development tool and has been used in, among other things, developing the CRAM C++ library and the CRAM C++ simulator, and generating input files for the controller VHDL simulations. The following section briefly describes the CRAM assembly language.

6.4.2 CRAM Assembly Language

A complete listing of the CRAM Assembly Language is found in Appendix C.1. This includes the syntax and usage examples of all the instructions in the language. The following sections give a very brief description of the instructions:

- **CRAM Memory Variable Instructions:** These are instructions that operate on CRAM memory variables (`cvar`). They include instructions to perform arithmetic, logical, relational, and search operations, as well as to shift, load, copy, set, and clear CRAM variables. Some of these instructions can have an immediate value as one of its source operands. A few of the logical operations are provided with what are called 'boolean instructions' in order to avoid the overhead of setting up the word-length (WLEN) register for operations involving boolean CRAM variables (see Appendix C.1 for details). A few examples of `cvar` instructions are shown below.

<u>**Examples**</u>

```
ADD #24, #16, #8;     // mem[24] = mem[16] + mem[8]
INC #24, #16;         // mem[24] = mem[16] + 1
MAX #27, #16;         // mem[17] = (is maximum element of mem[16])
MCLR #18;             // mem[18] = 0
MOV AR0, AR1;         // mem[AR0] = mem[AR1]
```

- **CRAM PE Register Instructions:** These are instructions that operate on PE registers. They include instructions to set or clear registers, and instructions to load PE registers from either other PE registers, or from `cbool` or `cboolref` variables.

<u>**Examples**</u>

```
CLR X Y;              // all X and Y PE registers are set to 0
LDPE X, !#16;         // load all X registers with inverse of mem[16]
TXPE X, Y;            // load X registers with contents of Y registers
```

- **CRAM Controller Instructions:** CRAM controller instructions include instructions to load, increment and decrement controller registers, as well as instructions for reading and writing data between CRAM and the controller. There is also a conditional read-bank instruction (RDBNK) used to scan a `cbool` or `cboolref` variable to check if there is a 1 at any PE position. This is done by first reading the byte of the variable corresponding to the first eight PEs. If this byte contains a 1 on any of its eight bits, the byte is saved in the DTR register and the RDBNK instruction terminates. Otherwise the bank address register (CBA) is incremented to point to the byte corresponding to the next eight PEs, and the cycle is repeated. This is repeated until either a 1 is found or all the bits of the variable have been scanned.The results of the scan (CBA and DTR) can be used by the host processor to compute the index of a PE that yielded a true value to a search or comparison.

Examples

```
LDAX0 #5;          // AX0 = 5
INCA AR1;          // AR1 = AR1 + 1
WRITE AR0;         // mem[AR0] = write_buffer[WIBA]
RDBNK AR0;         // loop (DTR = mem[AR0]) until mem[AR0] != 0
```

6.4.3 Using `cvar` Addresses and C++ Integer Variables

One special feature of the CRAM assembly language is that when it is mixed in CRAM C++ code, one can use the address of a C++ defined `cvar` variable where ever it is legal to use `#AR0`, `#AR1`, and `#AR2`. The address of a `cvar` is specified as *`var`*`.addr`, where *`var`* is the `cvar` variable. Any valid expression of `addr` can be used in the assembly code, e.g. *`var`*`.addr+2` for the address pointing to bit 2 (the third bit) of *`var`*. Similarly, any integer defined in the C++ code can be used where ever an immediate value or a `cvar` address is allowed in the assembly code. This allows easy optimization of part of the CRAM C++ code using CRAM assembly code.

Example

```
cint c; cuint b(8), a(8);                  // C++ defined cvar's
int x = 3;                                 // C++ defined integer variable
CASM
  LDWLEN #x                                // WLEN = x (i.e. WLEN = 3 or 4-bit word-length)
  AND c.addr+8, a.addr+4, b.addr+4;        // c[11:8] = a[7:4] & b[7:4]
END CASM
```

6.5 CRAM Microcode Assembler

The CRAM Microcode Assembler (CMASM), simply known as the CRAM Microassembler, is used as a high-level development tool when generating microinstruction routines. Instead of writing microinstructions in their native binary format, CRAM Microassembly Language is used. Typically, this is high-level language for low-level PE/RAM operations (e.g. `AND X, Y M X;` i.e. let each PE do an AND operation of the contents of registers Y M, and X, and write the result into X), and CRAM controller microinstruction execution control (e.g. `SETLOOP...LOOPN IF LPCNZ;` i.e. execute these microinstructions until the loop counter is zero).

Microinstruction routines for all operators and functions in the CRAM C++ library and the CRAM Assembly Language have already been generated (using CMASM) and comes together with the CRAM software library. Still there might be need to extend the C++ library or CASM, especially if a certain application-specific function occurs so frequently in application source codes as to warrant being implemented as a microroutine in the controller control store. In this case, the programmer can use CMASM to generate the microinstructions and append them to the existing ones. Otherwise, like CASM, the microassembler is mainly used by the CRAM designers in system development and analysis work. Details of the Microassembly Language can be found in Appendix C.2.

6.6 Grouping of Microroutines

As described in Chapter 3, the ALU result of a CRAM PE is derived from the truth table of its three input registers. Therefore, source operand registers are specified in terms of an 8-bit truth-table value, and this value is unique for each operation. This, together with the specification of destination operands, makes the number of microinstructions required for an operation much bigger when compared to standard processors. For example, with four possible destination registers (W, X, Y) and six possible source registers (X, Y, M, !X, !Y, !M), an instruction to move the contents (or inverse) of a register to another requires 34 unique 14-bit values of COP-TTOP combinations (6-bit COP and 8-bit TTOP).

To reduce the number of microinstructions in the control store, microroutines of all

instructions that differ by TTOP and/or COP in only one corresponding microinstruction have been grouped together. For each group, only one representative microroutine is loaded into the control store. The different values of TTOP and/or COP for the specific instructions in the group are coded in the unused operand fields of the instruction issued from the host. This has three advantages when compared to using bits in the instruction to point to specific microroutines or microinstruction subroutines. First, the size of the control store is reduced since we don't have to store the other microroutines or microinstruction subroutines. Second, the level of address decoding and multiplexing for the microprogram sequencer is reduced. This reduces both the area and critical path of the sequencer. Third, this approach further justifies the use of a uniform RISC-like instruction format for the macroinstructions because the wastage of bandwidth due to unused fields of the instruction word is now smaller. As mentioned earlier, a uniform instruction format results in simpler and faster instruction decoding and flow. Figure 6.8 illustrates microroutine grouping for instructions that set, clear, or transfer data between PE registers.

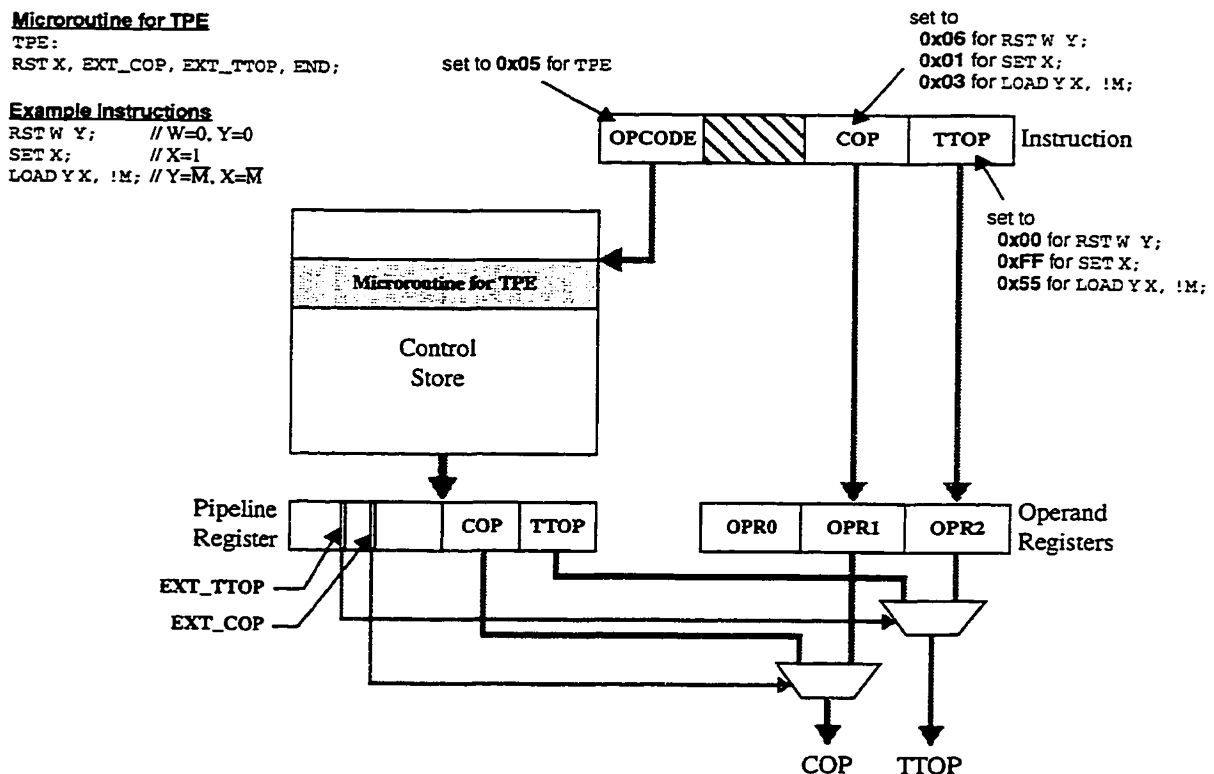

Figure 6.8 Grouping of Microroutines

Microroutines with more than one microinstruction are also grouped and executed in a similar way, with external TTOP and/or COP being selected only for the one specific microinstruction. Note that microroutine grouping is possible only if the instruction has less than three operands. Table 6.4 shows all instruction groups and those where microroutine grouping has been applied. Using this approach has reduced the number of microinstructions for basic operations from 462 to 197, i.e. a more than 57% reduction in the required size of the control store.

Instruction Group	No. of instructions	Total no. of microinstructions	†Grouped instructions	Microinstructions in control store
No operation	1	1	-	1
Controller-to-CRAM access	4	4	-	4
Controller instructions	9	9	1x9	1
PE reg-to-reg transfer, set, reset	48	48	1x48	1
Load PE-reg from memory	14	28	1x14	2
Load mem from PE-reg or constant, Extend result of logical operate	20	40	2x8, 1x10 1x9	6
Memory-to-memory move, Extend operand of logical operate	8	24	1x5 1x2	9
Add/sub, logical operate two `cvar`'s	7	39	-	39
Add/sub `cvar` & constant	3	18	-	18
Logical operate `cvar` & constant	3	12	1x3	4
Increment, decrement, negate, Extend operand of add/sub	16	81	2x3, 1x2 2x2	18
Extend constant of add/sub	6	27	3x2‡	15
Extend constant of logical operate	3	9	1x3	3
Extend result of arithmetic operations	12	47	2x6‡	10
Compare two `cvar`'s	2	8	1x2	4
Compare `cvar` and constant	3	9	1x3	3
Extend operand of compare	6	12	2x3	4
maximum/minimum search	2	18	1x2	9
PE Shifts	4	18	-	18
Data-parallel conditions	2	14	-	14
Data transpose	2	14	-	14
Total	**175**	**462**	-	**197**

†Example: 2x3 means 2 groups of 3 instructions each.
‡Partial or irregular grouping.

Table 6.4 Instruction Groups

6.7 CRAM C++ Simulator

6.7.1 Introduction

The CRAM C++ simulator is a tool that is used to simulate the behavior of a CRAM system. This can be used by both hardware and software designers. Hardware designers can use the simulator to explore different architectures of CRAM, its controller, and overall system design before committing to the actual hardware implementation. This would improve both the quality and performance of the design, while reducing the number of architectural errors that are only apparent after a design has been simulated or implemented. CRAM software tools designers can use the simulator to test and debug their tools, and applications developers can test their programs and extract more accurate timing information. The fact that the simulator offers a generic number of processing elements (PEs) makes it an ideal tool for performance analysis. While small prototype systems can easily be implemented at this stage of CRAM research work, the potential advantages of CRAM can only be demonstrated relevantly on systems with a large number of PEs. The simulator allows this to be easily realized.

For the simulator to be an accurate architecture exploration tool as well as yield more accurate program execution and timing behavior, the simulation model of the CRAM PEs and the CRAM controller is based on the hardware description of the actual implementations rather than on just their behavior. The disadvantage of this is that it makes the design of the simulator more complex and the simulator runs slower than that based purely on the behavior of the components. However, such a simulator gives out a behavior that is closer to that obtained by the design hardware description (VHDL/ Verilog) used in the actual design and implementation phases. This reduces the number of changes required when moving from an architecture obtained through simulation to that which is to be implemented using a hardware description language. It also yields more accurate timing information than that obtained in a behavioral simulator because in the latter a number of processes are usually bundled up in a single simplistic behavior.

The CRAM simulator is not an independent tool. Rather, it is designed to be a library that, when needed, can be linked with an application together with the CRAM C++

library. In this way, the execution of a program on the host, up to the point where the host transfers data onto the CRAM card or simulation model, is not a simulation but the real thing. This has some advantages. First, the complexity of the simulator is minimized since very few host features need be incorporated into the simulator. In particular, only the simulation of the CRAM card driver, and the timing information of the host system buses are required. All other aspects of program compilation and execution are already taken care of by the CRAM compiler. The second advantage is that even when you are using the simulator instead of an actual CRAM card, the timing information and program execution as related to the host system (except for transfers over the system bus) is exactly the same. This is especially important during the design and fine-tuning of the CRAM C++ library and other software tools.

6.7.2 Simulation Model

Figure 6.9 shows the CRAM C++ simulation model. When code executes on the host system, the simulator is invoked only when the host transfers data to the CRAM system. A single line in the CRAM compiler checks if the execution is on a simulator or on an actual CRAM card. If on a simulator, the host transfers data to the CRAM system by calling the simulator CRAM driver, otherwise the real driver is used. The simulation models of the CRAM system and the host environment are described in Section 6.7.3 to Section 6.7.5, and the different outputs of the simulator are described in Section 6.7.6 to Section 6.7.8.

6.7.3 Host Objects

The CRAM driver is the only component of the host system that is incorporated into the simulator. It is used as the link between the actual host system and the simulation model of the CRAM system. It simulates both the driver as well as different system buses (PCI, ISA, etc.). A link to the host system is through a function call to the driver (as would be done on a real system), while a link to the CRAM controller simulation model is through a function call that emulates the sequence of bus signals required to complete the data transfer on a specific system bus. Currently, only the ISA bus protocol is implemented in detail. The other buses are simulated by specifying the generic parameters

of the bus model such as bus width, bus speed, and whether the bus supports burst transfers.

Like all other models in the simulator, the host simulation model is implemented as a C++ class. Its class constructor is used, among other things, to model what happens when the CRAM driver is being initialized, such as the loading of microinstructions into the CRAM control store. Most timing parameters are also initialized at this time.

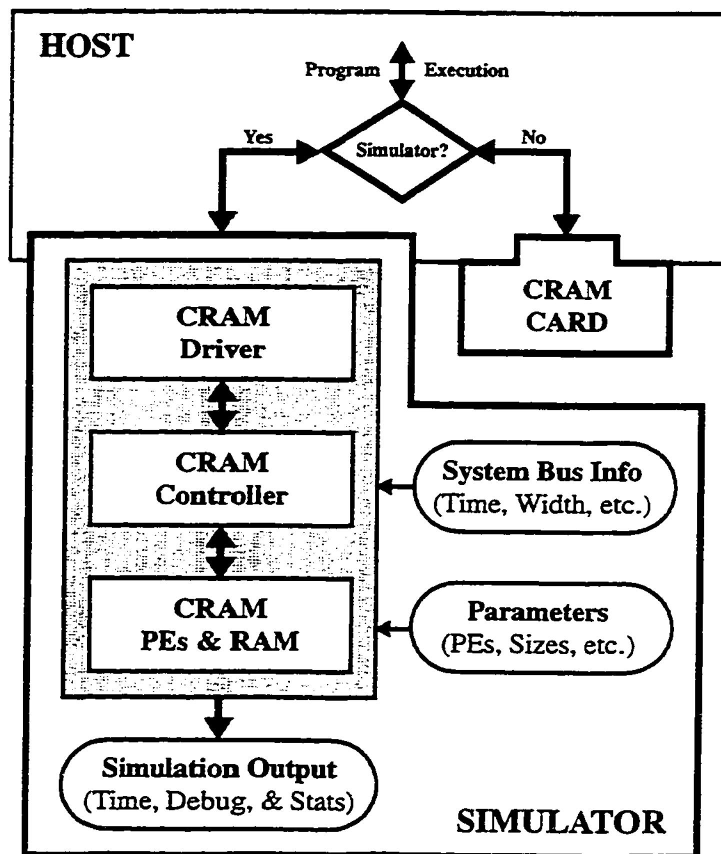

Figure 6.9 CRAM C++ Simulation Model

6.7.4 CRAM Controller

The simulation model of the CRAM controller is a hardware description (using C++) of the controller core components as well as the CRAM-host interface. Components are described as classes, with functions used to describe specific component features such as external and internal access and instruction execution. The flow of an instruction through the instruction execution unit (IXU) is described the same way as in the hardware, i.e.

instruction read from FIFO to IR, then executed in the microprogram sequencer using microinstructions, PE or RAM functionality decoded, address and data multiplexed with TTOP and COP and driven onto the CRAM bus, and finally the CRAM control signals (OPS, MCK, RW, etc.) asserted. All the other components are also modeled with such detailed hardware description. Such modeling is very easy to do if the real hardware was designed using a hardware description language such as VHDL or Verilog. This detailed modeling yields more accurate simulation and timing.

Like in the host simulation model, only the CRAM-ISA interface is described in detail. The other interfaces are simulated by specifying generic parameters of the bus model used in the host simulation model. Other generics in the controller model include the sizes (depths) of the instruction FIFO, the read/write buffers and the control store, as well as the size of the CRAM data and address buses.

6.7.5 CRAM PEs and RAM

The simulation model of the CRAM chips include the low-level hardware description of the PE architecture and the CRAM external interface. Other features of CRAM that relate to its computational nature (such as the bus-tie and shift operation) are also modeled based on how they connect and function in hardware. On the other hand, the memory component of CRAM is simply modeled as arrays of PE local memory since the hardware design of the memory is not a principal focus of the CRAM project. A generic number of PEs and memory bits/PE is used to simulate CRAMs of different sizes.

6.7.6 Timing Information

This measures the duration of time during specific stages of program execution from the host processor, via the system bus, to the CRAM controller, through the controller instruction execution unit, to the CRAM chip. This information can be used by both hardware and software designers. For example, the hardware designer can decide to adjust the size of the instruction FIFO or change the instruction pipeline scheme by studying the time that is spent filling an empty instruction queue or the time that the host processor has to wait before another CRAM instruction can be transferred to the CRAM controller.

CRAM software tools designers can evaluate the times that the host processor spends executing a specific CRAM C++ library routine and then decide whether to optimize or remove the routine. Applications designers can assess the time that an application takes to run on a CRAM system. Currently, the major timing parameters reported by the simulator include:

- **Host Execution Time:** The time that the host processor is executing some instructions while there are no instructions in the CRAM instruction queue, i.e. no parallel host/ CRAM activity.

- **Instruction Load Time:** The time taken by the host processor to load instructions into an empty CRAM instruction queue. This can be used, among other things, to determine the effect that the system bus transfer rate or the size of the instruction FIFO has on the overall program execution time.

- **Microinstruction Load Time:** The time that CRAM microinstructions are being loaded into the control store while the CRAM instruction queue is empty. This time is used to study the effect of the control store size as well as the make-up, size and selection of CRAM microroutines. Microinstruction load time does not include the time to load the initial microinstructions since this happens at boot-up during CRAM driver initialization. Therefore for a system with a control store big enough to require no microinstruction refilling, this parameter is always zero.

- **Data Access Time:** The time that the CRAM buffers are being accessed by the host processor while the CRAM instruction queue is empty. This time is used to study the effect of the size and configuration of the CRAM data buffers, as well as how an application data transfer loads and styles affect how fast it runs on a CRAM system.

- **Register Access Time:** The time that the CRAM memory-mapped registers are being accessed by the host processor while the CRAM instruction queue is empty.

- **Pipeline Fill Time:** The time to fill an empty CRAM instruction queue (FIFO⇒IR⇒Pipeline Reg) when an instruction begins executing from the FIFO. This is useful in studying, among other things, the effects of the FIFO size, the instruction

pipeline scheme and depth, and the inability of the host processor and the CRAM library to keep the CRAM instruction queue always full.

- **Microinstruction Execution Time:** The time taken to execute all CRAM microinstructions of routines invoked in an application. This is the major component of the total execution time of a program, and is calculated by multiplying the total number of microinstructions executed with the controller cycle time.
- **Total Execution Time:** This is the total time taken to run a program on a CRAM system. It is simply the sum of all the timing parameters described above.

All the timing parameters above accumulate as the program executes, but each can be individually reset at any point in the program or library routine in order to allow timing of specific features of a program or CRAM library routine. Note that most timing parameters are evaluated only when there is no concurrent execution of instructions in the CRAM controller because this is the only time that they contribute to the total execution time of a program. It must also be mentioned that there may be other uses of these parameters other than the ones mentioned here.

6.7.7 Debugging Information

The simulator outputs various information that can be used to debug an application, CRAM library routines, microinstruction routines, or even the CRAM software tools themselves. The main debug information includes:

- **CRAM Memory Image:** This gives the state, in binary format, of the local memory of each PE. The image is displayed with each PE and its local memory forming a single column of the memory array. The memory image can be displayed at any point in the application code or CRAM library routine, and the number of memory rows to be displayed can be specified. This, together with the fact that memory rows allocated to CRAM variables are grouped and separated from others, allows that a malfunctioning microroutine, library routine or part of an application code be debugged by simply looking at how they actually update the CRAM variables they are supposed to affect.

- **Microinstruction Execution:** As each microinstruction is being executed, its 32 bits, plus the resulting status of the CRAM address, data, and control buses are displayed. This may be used to debug microroutines or the simulation model of the controller.

- **Assembly Code Listing:** This lists the CRAM assembly code as the program executes. This feature is actually incorporated into the CRAM compiler and is activated using a #define statement. It may be used to debug CRAM C++ library routines.

- **Bus Transfers:** This displays the address and data of transfers between the host and the CRAM system.

6.7.8 Program Execution Statistics

Apart from timing and debugging information, a few statistics are reported by the simulator during program execution. These include:

- **Number of Specific Instructions:** Specific assembly code instructions may be counted. Typically, these are instructions used to update CRAM controller registers, such as LDWLEN and LDAXn. Such information may be used to optimize microroutines, the architecture of the controller, and the formats of both macroinstructions and microinstructions.

- **FIFO-Full Transfer Failure:** This is the number of times that the host data transfer to the CRAM instruction FIFO fails because the FIFO is full. In this case the host has to retry the transfer at a later time. This can affect the performance of the system especially if the host processor is working in a multitasking environment. The number of such failures depends on so many things, including the size of the FIFO, the depth of the instruction pipeline, and the cycle time of the controller in relation to the speed of both the host processor and the system bus.

- **Number of Microinstructions:** This is the number of CRAM microinstructions executed in a program. It is used, among other things, to calculate the main component of a program total execution time.

6.8 Summary

Writing application programs using low-level CRAM machine code requires detailed knowledge of the architecture of the CRAM chip and its controller. Therefore, to assist application programmers, software tools developers, and hardware or system designers, four high-level software tools that can be used to write and analyze CRAM applications and software tools, as well as analyze the different architectural features of a CRAM system have been developed. The CRAM C++ Compiler is a library of classes for CRAM parallel variables and controller objects that can be used to write CRAM programs using the standard C++ language. Such programs are compiled using a standard C++ compiler such as GNU C++ (gcc) or Borland Turbo C++. The CRAM Assembler can be used to write applications in CRAM assembly code, especially in cases where the C++ code does not yield the required speed or code size. The CRAM Microcode Assembler is a high-level tool for generating CRAM microinstructions. The CRAM C++ Simulator is used to simulate the behavior of a CRAM system. This tool outputs several timing and debug information, as well as other program execution statistics that can be used by both hardware and software designers.

This chapter has also described two other software issues: data transposition and microroutine grouping. Because of the high hardware cost and complexity of hardware corner-turning, data on CRAM is transposed in software. In this regard, a parallel array-based corner-turning approach that is several hundreds times faster than host-based transposition, and contributes less than 10% of the total I/O overhead for CRAM systems with more than 4K PEs has been developed. This removes the need for a hardware data transposer, thus reducing the area and complexity of a CRAM system. It also minimizes the effect of the host on CRAM performance, once again making it easier to implement CRAM systems on different platforms.

Finally, microroutines of similar instructions have been grouped together using unused fields of the instruction word. This has reduced the required size of the microprogram memory by more than 50%. This small size (less than 256 32-bit words) makes an on-chip control store feasible, even in standard ASIC technologies and FPGAs, and hence reduces the number of components on a CRAM system PCB.

Chapter 7

Applications and Performance Analysis

This chapter looks at CRAM applications and performance. The main objective is to show the suitability and performance of CRAM for practical applications of different characteristics. First, the characteristics of applications suitable for CRAM are highlighted, and the methodology used in performance evaluation is described. Section 7.2 then analyzes the performance of basic operations (addition, multiplication, etc.). After this, a few selected practical applications are described in terms of their definition, algorithm, implementation, and performance on both CRAM and conventional uniprocessor systems. The fields from which applications are selected include low-level image processing (Section 7.3), database applications (Section 7.4), and image and video compression (Section 7.5). The performance impact of the CRAM controller and other CRAM architectural features is analyzed in Section 7.6. The chapter concludes by comparing CRAM with other SIMD and logic-in-memory systems in terms of performance, hardware, and complexity.

7.1 Analysis Methodology and Parameters

Applications most suited for CRAM, like most massively parallel SIMD machines, are those that have fine grain parallelism and regular communication patterns. Such applications can be found in numerous fields including image processing, database operations, video and image compression, digital signal processing, computer-aided design, graphics, and numerical analysis [6], [7], [4]. The CRAM-implementation of a few of such applications is described in this chapter. The main objective is to highlight the use of CRAM for practical applications, as well as analyze the performance of the CRAM controller and the whole CRAM system based on practical applications. Three criteria have been used to choose the applications described here. First, different applications have been selected from different fields to show the general-purpose nature of CRAM. Second, these applications are of varying complexity and computation models in order to analyze the performance impact of different features of the CRAM system, such as inter-processor communication, global-OR usage, loading of scalars from the host, control and host overhead, degree of parallelism, and data types. Lastly, it is not within the scope of this thesis to find, develop, and test algorithms for all (or as many) applications that are suitable for CRAM. As mentioned earlier, the goal here is mainly to use some practical applications to analyze the CRAM hardware and software developed in this thesis, as well as to show that CRAM can indeed be used for practical applications. For this reason, the author has developed CRAM algorithms for only a few applications, and these are the ones described in this chapter. Appendix D lists the CRAM and uniprocessor C++ code for these applications. Other applications not described here, but whose CRAM algorithms have been developed by other people, include fault simulation, data mining, satisfiabilty problem, and FIR filters [4], as well as discrete cosine transform, adaptable and scalable vector quantization, and other MPEG-2 algorithms [65], [66], [67], [68].

Because of the small sizes and limited number of the prototypes implemented so far, the performance analysis work is carried out using the CRAM C++ Simulator. This not only allows the use of bigger CRAMs but also makes it possible to study how the performance of applications change as the size and other parameters of the CRAM system are varied. As described in Section 6.7, the CRAM C++ Simulator is designed to be a very

close approximation of the real CRAM system. Therefore the accuracy of the results obtained in these simulations is very close to the actual numbers. This has been verified by running some of the smaller algorithms, especially arithmetic, logic, and search operations, on the 64-PE ISA CRAM system prototype. The CRAM simulation model used for performance analysis is a 50 ns, 32 MBytes, 64K PE, PCI-based CRAM card on a 133 MHz Pentium PC. This is compared to the performance of running the applications on two uniprocessor systems: a PC driven by a 133 MHz Intel Pentium processor, with 32 MBytes of RAM (the same system serving as the host for the CRAM system), and a SUN Sparc Ultra workstation driven by a 167 MHz Sparc processor, with 64 MBytes of RAM.

7.2 Basic Arithmetic, Logic and Memory Operations

A basic way of assessing the performance of CRAM is to look at the execution of basic computational and memory-access operations. Table 7.1 shows the number of CRAM microinstructions required to perform these operations. It also shows the number of 32-bit and 8-bit operations performed per second by a 20 MHz, 64K PE CRAM.

Operation	Notation	Microinstructions per n-bit operation	8-bit Giga-Operations per second	32-bit Giga-Operations per second
Addition	C = A + B	$6n + 1$	26.75	6.79
Multiplication	C = A * B	$8n^2 + 16n + 2$	2.04	0.151
Division	C = A / B	$18n^2 + 53n + 2$	0.83	0.065
Logical AND	C = A & B	$5n$	32.77	8.192
Memory Clear	C = 0	$n + 1$	145.64	39.72
Load Immediate	C = #k	$2n$	81.92	20.48
Copy	C = A	$3n$	54.61	13.65

Table 7.1 CRAM Basic Operations

Table 7.2 compares the execution times of these operations on the CRAM and the two uniprocessor systems. This is for 64K *long* (32-bit) and *char* (8-bit) integer arrays. On CRAM, the execution time is proportional to the number of microinstructions in the

operations. But as shown in Table 7.2, these differences between operations do not translate into proportional differences in their execution times on the uniprocessor systems. This is reason why instructions with a smaller number of CRAM microinstructions yield higher speedups.

Application	**Workstation** Execution time (ms)		**PC** Execution time (ms)		**CRAM** Execution time (μs)		**CRAM** Speedup over PC	
	32-bit	8-bit	32-bit	8-bit	32-bit	8-bit	32-bit	8-bit
Addition	10.76	7.25	12.07	6.43	9.65	2.45	1251	2624
Multiplication	15.19	14.6	15.57	10.71	435.3	32.1	36	334
Division	27.28	27.21	33.01	17.03	1006.5	78.9	33	216
Logical AND	11.74	7.34	12.07	6.43	8.0	2.0	1509	3215
Memory Clear	3.89	2.99	4.62	3.86	1.65	0.45	2800	8578
Load Immediate	3.89	2.99	4.62	3.86	3.2	0.8	1444	4825
Copy	2.8	2.73	8.85	5.12	4.8	1.2	1844	4267

Table 7.2 Performance of Basic Operations

The performance of SIMD machines with bit-serial processors stems from the massive number of processors rather than from the computational power of the individual processors. It is therefore important to assess how the performance of CRAM varies with the number of PEs. Figure 7.1 shows this variation for the basic operations. There are two things to note from the figure. First, for larger numbers of PEs, the speedup is proportional to the number of PEs. But for smaller numbers of PEs, this variation is not linear because the execution time of the operation itself becomes much smaller when compared to the

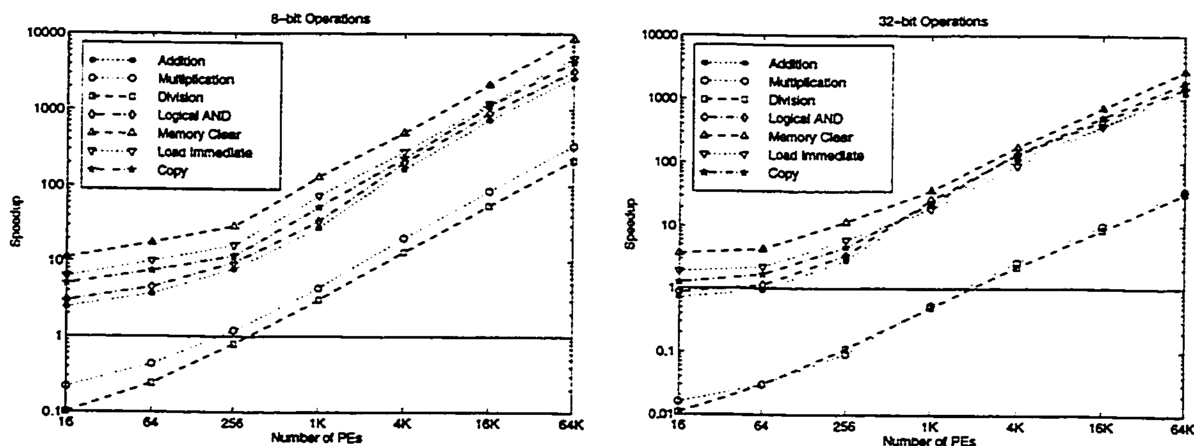

Figure 7.1 Effect of Degree of Parallelism on Basic Operations

other overheads of executing the operation. The second point is that for operations with bigger (bit-size) operands or higher degrees of complexity, not only does the speedup decrease, but the threshold at which CRAM performs better than the uniprocessor occurs at a larger number of PEs.

As shown earlier, the execution time of most CRAM operations is proportional to the bit size of their operands. On uniprocessor systems that are based on fixed-size (byte, word, double-word, and quad-word) operands, the execution time of most basic operations do not show substantial variation with operand size. For example, the execution time of 32-bit addition on the PC is less than twice the execution time of 8-bit addition, even though the operands are 4 times bigger. Because of these small differences in performance, it is common for programmers to use operand sizes that are bigger and more common (such as integers) than the required precision. But as shown in Table 7.1, the difference in execution times between CRAM operations with different operand sizes is very substantial. It is therefore more important to choose the right operand precision when writing programs for CRAM because there is a significant difference in the speedups over the PC if the operation is executed with different operand sizes. This is illustrated in Figure 7.2 for the basic operations on 64K arrays. Note that on the PC, 2-bit, 4-bit, and 8-bit operations all have the same execution times, thus yielding even higher CRAM speedups for operations on bit sizes less than 8.

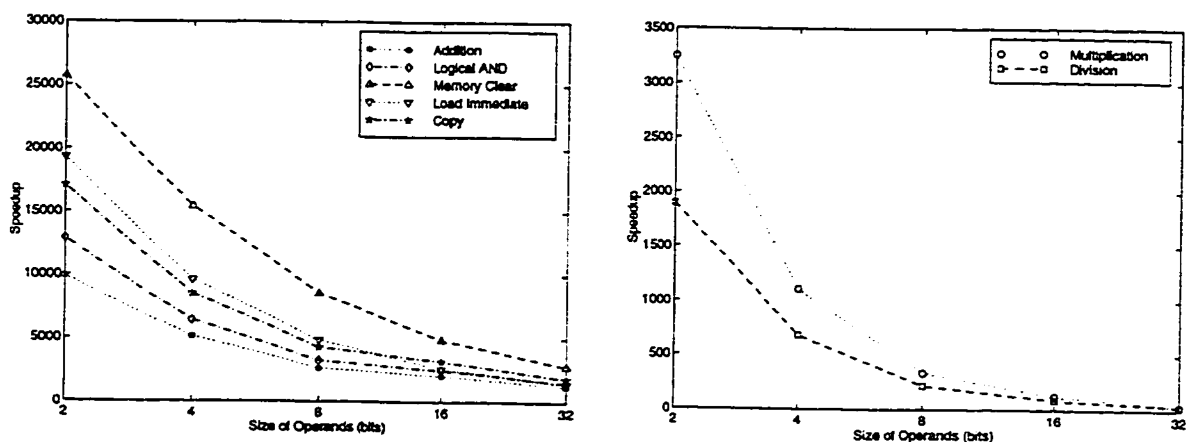

Figure 7.2 Effect of Operand Size on Basic Operations

7.3 Low-Level Image Processing

Low-level image processing involves sensing and pre-processing of an image before it is subjected to medium-level and high-level image processing operations such as segmentation, description, recognition and interpretation. Typical low-level operations include image enhancement, filtering to remove noise or improve image fidelity, edge detection, and threshold-based segmentation.

Low-level image processing is particularly suited for CRAM for two main reasons. First, most of the operations are massively parallel. The same computational operation is applied to numerous (usually more than 1000) pixel data. Therefore, each CRAM PE can be assigned to one or a few pixels and the computation done in parallel. Second, most low-level image processing operations are computed over small pixel neighborhoods [69]. These operations are either pointwise, i.e. the result at a pixel is a function of the initial data at that pixel alone, or they are local to a 3 x 3 or smaller pixel neighborhood. This means that data can be transferred from nearby PEs. This is important because of the limited inter-processor connectivity on CRAM.

The following sections describe the parallel algorithms and CRAM implementation of some low-level image processing operations. These are based on sequential algorithms adapted from those described in the literature such as [69], [70], and [71]. The simulations are based on a 256 x 256 grey-scale (8-bit) image and a 64K-PE CRAM, i.e. one pixel per PE. The timing results are for the case in which CRAM is used as the main memory. Otherwise, if the image is to be loaded onto CRAM from the host computer, processed, and then read back to the host, data I/O overhead has to be added to the total timing. This includes the time to write and read the data, and the time to transpose it. For a 256 x 256 8-bit image (64Kbytes), these times are 1.64 ms (2 x 0.82ms), and 31.58 μs (2 x 15.79 μs) respectively. It must be noted that usually several low-level image processing operations are applied to the image before it is used (displayed or further processed) elsewhere [11], [69]. For example, image enhancement operations such as contrast enhancement are usually followed by filtering. In this case, the I/O overhead is shared by several operations.

7.3.1 Brightness Adjustment

Brightness adjustment is used to brighten up dark images or darken images that are too bright. This is done by adding (or subtracting) a value to (from) each pixel in order to shift the pixel intensity by the specified number of grey levels. Algorithm 7.1 shows the brightness adjustment algorithm for both a uniprocessor and CRAM.

```
/* on uniprocessor */
for pixel=0 to N-1 pixels
  add/subtract a value to/from the pixel;
  if pixel is greater than grey-level limit
    saturate the pixel;
end for
```

```
/* on CRAM */
in parallel
  add/subtract a value to/from the pixel;
  if pixel is greater than grey-level limit
    saturate the pixel;
end in parallel
```

Algorithm 7.1 Brightness Adjustment

To adjust the brightness of a 256 x 256 grey-scale image took an average time of 9.12 μs on CRAM, 2.72 ms on the workstation, and 7.44 ms on the PC. This represents a CRAM speedup of 298 and 816 over the workstation and the PC, respectively.

7.3.2 Spatial Average Low-Pass Filtering

Apart from poor contrast, images may contain random pixels that have values higher or lower than what they should be. One way to reduce this type of noise is to replace the value of each pixel with the weighted average of its neighborhood pixels, i.e.

$$g'(x,y) = \sum_{i=-m}^{m} \sum_{j=-n}^{n} w(i,j) \cdot g(x-i, y-j) \tag{7.1}$$

where g is the noisy image, g' is the filtered image, and $w(i,j)$ are normalized filter weights in an m x n pixel neighborhood. A common and simplest class of spatial averaging filters has equal weights and is operated in a 3 x 3 pixel neighborhood, giving

$$g'(x,y) = \frac{1}{9} \sum_{i=-1}^{1} \sum_{j=-1}^{1} g(x-i, y-j) \tag{7.2}$$

Its algorithms are described in Algorithm 7.2.

```
/* on uniprocessor */
for pixel=0 to N-1 pixels
   find average of 3 x 3 neighborhood pixels;
   replace pixel value with this average;
end for
```

```
/* on CRAM */
in parallel
   find average of 3 x 3 neighborhood pixels;
   replace pixel value with this average;
end in parallel
```

Algorithm 7.2 Weighted Average of Neighborhood Pixels

Most SIMD pixel processing machines have inter-processor connectivity that makes data access from the nearest neighbors implicit in the architecture. Examples of such hardware inter-PE communications include the 2-D (North, East, West, and South neighbors) connectivity used in MIT Pixel Processor [14], the square or "X" (nearest 8 neighbors) connectivity of BLITZEN [72], and other more complicated networks such as the hypercube and the global router used in the Connection Machine [73] and MasPar MP-1 [29]. For a 3 x 3 neighborhood, data can be accessed from a neighbor PE in at most two cycles for the 2-D inter-PE connectivity, and one cycle for the square connectivity. Other networks can transfer data between any two PEs in a few clock cycles (e.g. 16 clock cycles on the global router of MasPar MP-1). On CRAM, only a linear (1-D left/right) inter-PE communication is implemented in order to reduce the size of the PEs. This, however, results in large communication penalty when data has to be accessed from neighboring PEs. Algorithm 7.3 and Figure 7.3 illustrate the method that has been used to implement the access and computation of pixel data in 3 x 3 or smaller neighborhoods. For clarity, the figure shows the case for an 8 x 8 image.

```
do twice, first for shift = shift left, and then for shift = shift right
   shift image pixels 1 position into temp variable;
   use temp to access pixel p(x-1, y) or p(x+1, y);
   shift temp W-2 positions into temp;
   use temp to access pixel p(x+1, y-1) or p(x-1, y+1);
   shift temp 1 position into temp variable;
   use temp to access pixel p(x, y-1) or p(x, y+1);
   shift temp 1 position into temp;
   use temp to access pixel p(x-1, y-1) or p(x+1, y+1);
end do
```

Algorithm 7.3 Accessing Pixels in a 3 x 3 Neighborhood on CRAM

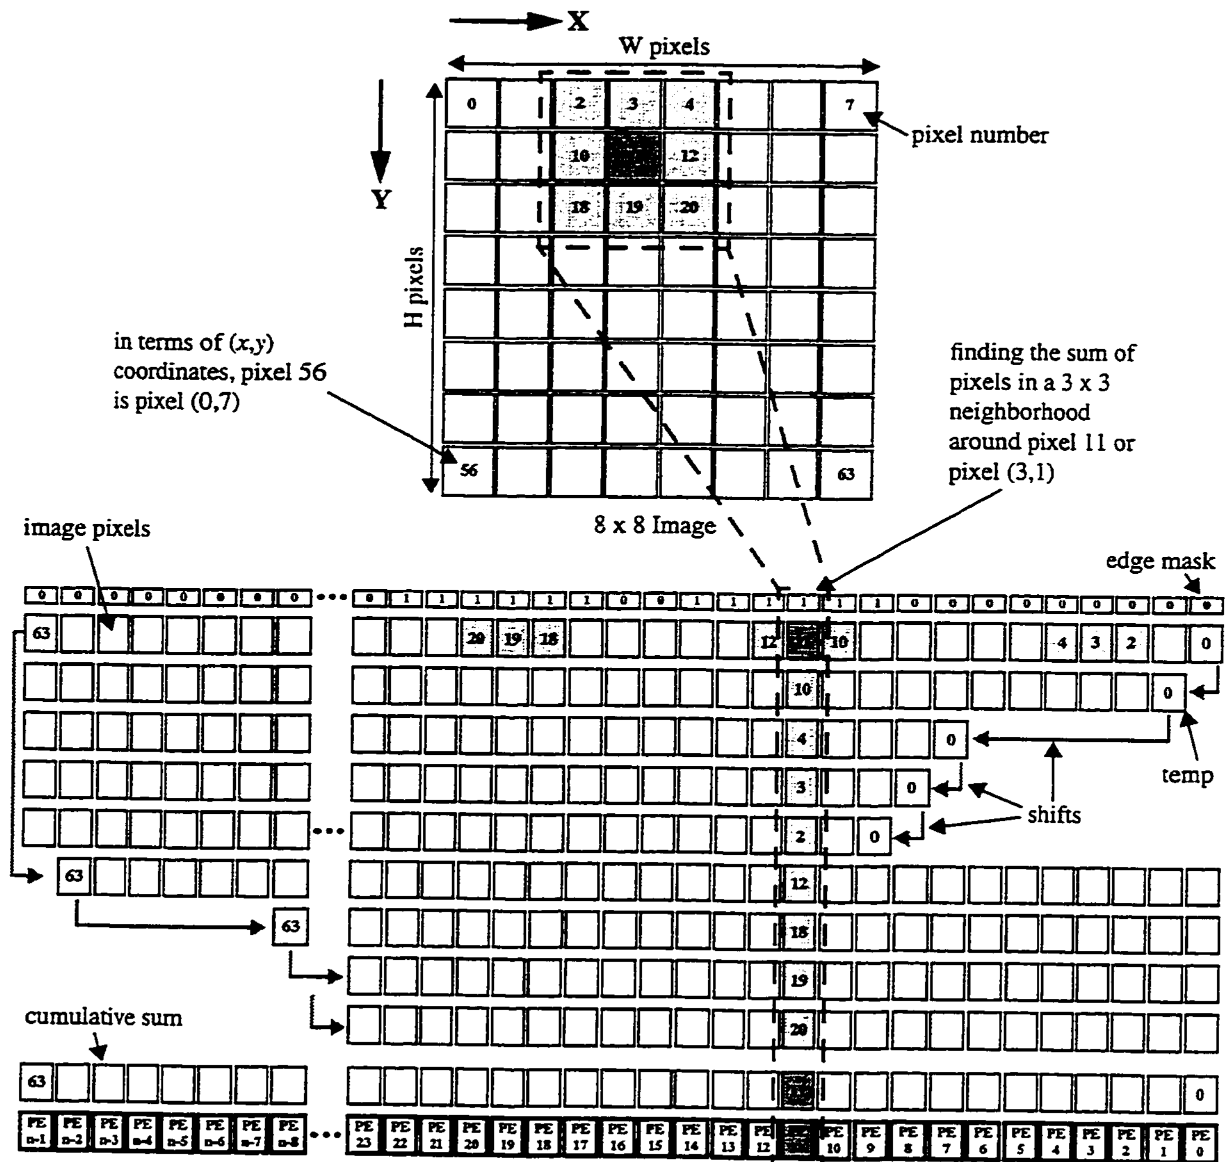

Figure 7.3 Computing in a 3 x 3 Pixel Neighborhood

The image is laid out across the PEs, one row at a time. For an image of width W pixels, the furthest pixel of the neighborhood will be W+1 PEs away. Instead of shifting the pixels independently, cumulative shifting is used. Therefore, to access a pixel to the left and the three pixels above, a total of W+1 shifts are required. The same is true for the other four pixels (one to the right, and three below). On the current CRAM system, shifting a b-bit variable by p PE positions takes $b(4+p)$ microinstructions. The case where ($p = 1$) has been optimized and uses only $4b$ microinstructions. Therefore the total time for the shift operations in Algorithm 7.3 is given by

$$T_{shift} = 2bT_c(14 + W) \quad (7.3)$$

where T_c is the CRAM cycle time. The edge mask in Figure 7.3 is used to disable the processing of border pixels. This mask can either be pre-loaded like the other masks, or its loading time (about 0.1 ms for a 64K image) can be added to the total execution time.

Despite the inter-PE communication bottleneck, the weighted average filtering of a 256 x 256 8-bit image exhibits a CRAM speedup of 136 over the PC, and 82 over the workstation. The respective execution times are 0.39 ms, 53.13 ms, and 31.78 ms. Of particular interest is that for CRAM the total shift time (T_{shift}) for the average filtering is equal to 0.216 ms and represents about 55% of the total execution time and 56% of the total CRAM microinstructions executed. This means that a CRAM with at least 2-D inter-PE connectivity would have about twice the speedup reported here. Also note that because the shift overhead represents only about half of the total execution time, there is no advantage in re-ordering the image to reduce W since this would require that the averaging routine be executed more than once, and/or that the borders of the partitioned image be further processed. In fact, the arrangement of the image pixels shown in Figure 7.3 is advantageous because it retains the format that is used for the pointwise operations.

7.3.3 Edge Detection and Enhancement

Edge detection and enhancement are important operations in image segmentation and object recognition [69], [71]. The Prewitt and the Sobel are two of the most commonly used 3 x 3 neighborhood edge-detection operators. These compute both the magnitude and direction of the edge. However, if only the presence of the edge and not its direction is of interest, the direction-invariant Laplacian edge detector (Figure 7.4) is usually used.

0	1	0
1	-4	1
0	1	0

Figure 7.4 The Laplacian Edge Detection Operator

Algorithm 7.4 [70] enhances the edges in an image by subtracting the Laplacian of a

pixel from the pixel itself. This is a high-pass filter implemented by simply subtracting the low-pass filter output from its input [71]. Apart from extracting edges, this algorithm is also used for sharpening images that are slightly out of focus or taken with a jittery camera. On CRAM, the pixels in the 5-pixel neighborhood are accessed using the method described in Algorithm 7.3 and Figure 7.3, with the slight variation that the last two 1-position shifts are eliminated, and the top and bottom pixels are accessed with *W*-1 shifts. The total execution times for a 256 x 256 grey-scale image are 0.25 ms on CRAM, 31.4 ms on the PC, and 14.83 ms on the workstation. These translate into speedups of 126 and 59. The pixel-access shifts represents 85% of the total execution time, and 87% of the total CRAM microinstructions executed.

```
/* on uniprocessor */
for pixel=0 to N-1 pixels
  find the Laplacian of the pixel;
  find difference between pixel and its Laplacian;
  find absolute value of the difference;
  if absolute difference is greater than 255
    saturate it to 255;
  replace pixel value with this difference;
end for
```

```
/* on CRAM */
in parallel
  find the Laplacian of the pixel;
  find difference between pixel and its Laplacian;
  find absolute value of the difference;
  if absolute difference is greater than 255
    saturate it to 255;
  replace pixel value with this difference;
end in parallel
```

Algorithm 7.4 Edge Enhancement

7.3.4 Segmentation

Segmentation is used to partition an image into a useful set of objects and the background. Each image pixel is classified into a class of pixels based on some property of the pixel. Two examples of segmentation based on thresholding are conversion to binary image and multiple thresholding.

(a) Conversion to Binary Image

One of the simplest approach to segment a grey-scale image g is to threshold it at a value T to produce a binary image g'. This threshold operation is defined as

$$g' = \begin{cases} 1, g \geq T \\ 0, otherwise \end{cases} \tag{7.4}$$

A suitable value of T can be obtained from the histogram of the image. Binary images are used in object recognition or image compression operations such as motion estimation [68], [74]. To convert a 256 x 256 8-bit image to a binary image (Algorithm 7.5) took 9.25 µs on CRAM, 5.23 ms on the PC, and 2.31 ms on the workstation. This represents speedups of 565 and 250.

```
/* on uniprocessor */
for pixel=0 to N-1 pixels
  if pixel is greater than or equal to threshold value
    set its value to 1;
  else
    set its value to 0;
end for
```

```
/* on CRAM */
in parallel
  if pixel is greater than or equal to threshold value
    set its value to 1;
  else
    set its value to 0;
end in parallel
```

Algorithm 7.5 Conversion to Binary Image

(b) Multiple Thresholding

If there are M objects with distinct grey-levels, the image may be segmented by multiple thresholds, i.e.

$$g' = \begin{cases} M, & g \geq T_{M-1} \\ M-1, & T_{M-2} \leq g < T_{M-1} \\ \ldots, & \ldots \\ 1, & T_0 \leq g < T_1 \\ 0, & otherwise \end{cases} \quad (7.5)$$

where T_k is the threshold for object k. The desired objects can then be thresholded out of image g'. Multiple thresholding a 256 x 256 grey-scale image with 16 objects exhibits a CRAM speedup of 513 on the PC, and 209 on the workstation. The execution times are 0.27 ms on CRAM, 138.41 ms on the PC, and 56.54 ms on the workstation.

7.4 Database Applications

Many database operations are suited for parallel processing [6]. However, because of the limited inter-PE communication and the slow speed of the individual bit-serial PEs, the type of operations particularly suitable for CRAM are those that are pointwise and have an

SIMD-implementation complexity of $O(1)$. Examples of such operations are basic searches and multi-field record matching. In this section, algorithms are presented along with the simulation results of running these operations on a database of 64K (65536) records of 32-bit (long integer) data. The algorithm for the Least Means Squared Match used for searching multi-field database records is adapted from the one described by Elliott [4].

7.4.1 Basic Searches

Examples of basic searches include:

- *Equivalence search* - search for records equal to or not equal to the search key.
- *Extreme search* - search for maximum or minimum record.
- *Threshold search* - search for records greater than or less than the search key.
- *Between-limits search* - search for records between two search keys.

These four search operations were simulated on a randomly-generated database of 64K 32-bit records and a 64K-PE CRAM. In both the uniprocessor and CRAM algorithms (Algorithm 7.6), all records matching the search criterion are replaced by a new value. Table 7.3 shows the simulation results.

```
/* on uniprocessor */
for record=0 to N-1 records
  if record matches search criterion
    replace record with new value;
end for
```

```
/* on CRAM */
in parallel
  if record matches search criterion
    replace record with new value;
end in parallel
```

Algorithm 7.6 Basic Searches

Basic Search Operation	Execution Time (ms)			CRAM Speedup Over	
	PC	Workstation	CRAM	PC	Workstation
Equal-to	5.96	2.94	0.0138	432	213
Maximum	6.85	3.38	0.0150	457	225
Greater-than	7.06	3.54	0.0128	552	277
Between-limits	8.73	4.67	0.0196	445	238

Table 7.3 Performance of Basic Searches

7.4.2 Least Mean Squared (LMS) Match

LMS is used in searches where the data records (and the search key) contain multiple fields or criteria. While an index may be used for each field of the records, the nearest match for the combination of criteria may not be the best match for any single criteria. One way to combine all the criteria in a record, in order to find the best match with a search key is to find the sum of the squared differences (SSD) between the records fields and the corresponding fields in the search key. The best match is the record(s) for which SSD is a minimum. Given a database of N records, each with K fields, $R^0 = \{r_0^0, r_1^0, \ldots, r_{K-1}^0\}$, $R^1 = \{r_0^1, r_1^1, \ldots, r_{K-1}^1\}, \ldots$, $R^{N-1} = \{r_0^{N-1}, r_1^{N-1}, \ldots, r_{K-1}^{N-1}\}$, and an K-field search key $S = \{s_0, s_1, \ldots, s_{K-1}\}$, the SSD is calculated as

$$SSD(S,R^i) = \sum_{j=0}^{K-1} (s_j - r_j^i)^2, \quad 0 \le i < N \tag{7.6}$$

Again, the LMS algorithm (Algorithm 7.7) replaces all matching records with a new value. Execution times for a database of 64K records, each with four 16-bit data, are 0.97 ms for CRAM, 81.7 ms for the PC, and 68.96 ms for the workstation. This represents a speedup of 84 over the PC, and 71 over the workstation.

```
/* on uniprocessor */
for record=0 to N-1 records
  reset difference to zero;
  for record field=0 to K-1 fields
    accumulate difference between record field
    and corresponding field of the search key;
  end for
  if difference is the minimum so far
    store the difference and the record number;
  end if
end for
for all records with minimum difference
  update record with new value;
end for
```

```
/* on CRAM */
// records (all fields) stored one per PE
in parallel
  reset difference to zero;
end in parallel
for record field=0 to K-1 fields
  in parallel
    accumulate difference between record field
    and corresponding field of the search key;
  end in parallel
end for
in parallel
  if difference is the minimum
    update record with new value
  end if
end in parallel
```

Algorithm 7.7 Least Means Squared (LMS) Match

7.5 Image and Video Compression

Image and video compression is a process of yielding a compact digital representation of images and video streams in order to minimize their storage and transmission requirements. This takes advantage of the high degree of redundancy present in these signals. Most image and video compression techniques involve a high degree of uniform computations on a large number of independent groups of pixels. This makes them highly suitable for implementation on CRAM. This section, presents algorithms and CRAM implementation of Vector Quantization and MPEG-2 Motion Estimation, two of the most commonly used image and video compression schemes [74].

7.5.1 Vector Quantization

Most multimedia applications are CD-ROM based and hence require only real-time decoding of the compressed video clips [74]. Unfortunately, most low-end machines cannot perform the decompression process in real time if the images are coded using transform-based coders such as JPEG and MPEG. Vector Quantization (VQ) is an alternative compression scheme that processes the image directly in the spatial domain. Its encoding processing is much more complex than transform-based coders, but it allows for very fast decoding using simple table lookups. Its simple decoder is the main reason why VQ is the method of choice for image and video coding on computing platforms with limited resources such as PCs and portable electronic notepads [74].

To compress an image using VQ, the input image is first divided into blocks of N x N pixels. These blocks are called M-dimensional vectors (where $M = N^2$). Each vector is then compared to each codeword of a previously generated K-word codebook to determine the codeword that best matches the vector (codebooks are generated using a collection of representative images). The index of the best matching codeword is transmitted or stored instead of the N^2 pixels, thus achieving a compression ratio of N^2:1. At the receiving (decoding) end, the image is reconstructed by using this index to point into a simple table lookup of an identical codebook. Common VQ distortion measures used in the matching process are mean square error (MSE) and mean absolute error (MAE). Given an N x N

input vector $V = \{v_0, v_1, \ldots, v_{M-1}\}$ and a set of K M-dimensional codebook vectors

$B^0 = \{b_0^0, b_1^0, \ldots, b_{M-1}^0\}$, $B^1 = \{b_0^1, b_1^1, \ldots, b_{M-1}^1\}$,....,

$B^{K-1} = \{b_0^{K-1}, b_1^{K-1}, \ldots, b_{M-1}^{K-1}\}$, MSE and MAE are calculated as

$$MSE(V,B^i) = \sum_{j=0}^{M-1} (v_j - b_j^i)^2, \quad 0 \le i < K \tag{7.7}$$

$$MAE(V,B^i) = \sum_{j=0}^{M-1} \left|v_j - b_j^i\right|, \quad 0 \le i < K \tag{7.8}$$

Image vectors are typically 2 x 2 or 4 x 4 pixels, and common codebook sizes range from 64 to 1024 words.

Algorithm 7.8 shows the VQ encoding algorithm for an image of I x J pixels. Its implementation on CRAM for a 512 x 512 8-bit image and a block size of 2 x 2 pixels is shown in Figure 7.5. Execution time for a 256-word codebook is 12.72 ms on CRAM, 74.95 s on the PC, and 38.22 s on the workstation. This represents speedups of 5892 and 3005.

```
/* on uniprocessor */
for image vector=0 to (I•J)/M - 1) vectors
  for codeword=0 to K-1 codewords
    reset distortion to zero;
    for vector pixel=0 to M-1 pixels
      accumulate distortion between vector pixel
      and corresponding pixel of the codeword;
    end for
    if distortion is the minimum so far
      store distortion and assign codeword index
      as the vector code in the coded image;
    end if
  end for
end for
```

```
/* on CRAM */
// image vectors are stored on PE's local memory
for codeword=0 to K-1 codewords
  in parallel
    reset distortion to zero;
  end in parallel
  for vector pixel=0 to M-1 pixels
    in parallel
      accumulate distortion between vector pixel
      and corresponding pixel of the codeword;
    end in parallel
  end for
  in parallel
    if distortion is the minimum so far
      store distortion and assign codeword index
      as the vector code in the coded image;
    end if
  end in parallel
end for
```

Algorithm 7.8 Vector Quantization

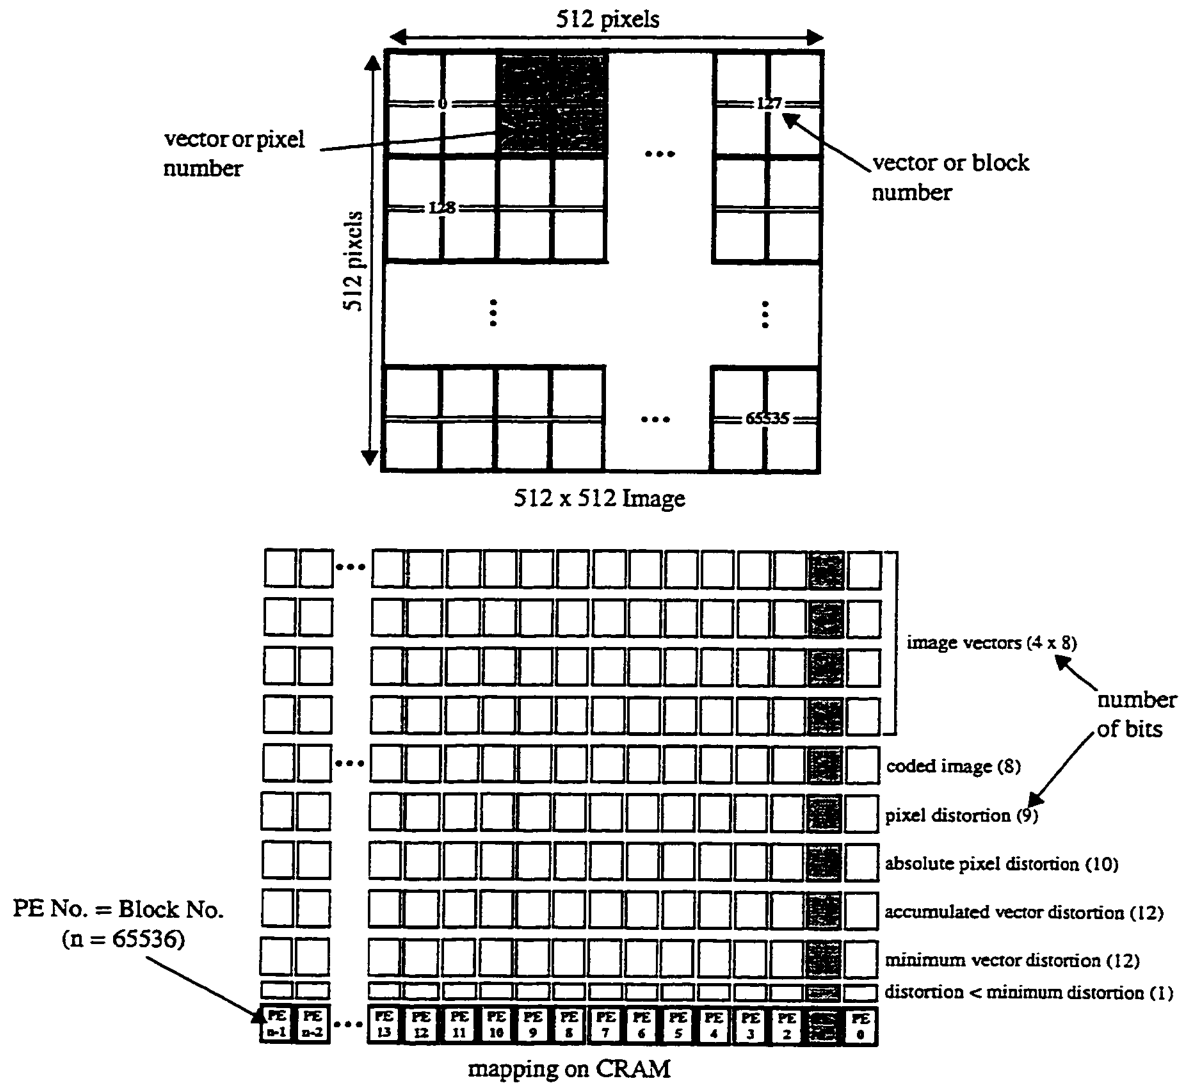

Figure 7.5 CRAM Implementation of Vector Quantization

7.5.2 MPEG-2 Motion Estimation

One of the most compute-intensive operations in the MPEG-2 video compression scheme is the motion estimation process [74]. If we define image frames at times t and $t-n$ as the current and reference frames respectively, then the objective of motion estimation is to determine the motion vector of an $N \times M$ block between the two frames. These vectors are used in forming compressed frames (P-frames and B-frames) that are transmitted between the minimally-compressed reference I-frames.

The motion vector of a block C at location (x,y) in the current frame is found by searching the region $[-p,p]$ around the block in the reference frame (Figure 7.6). There are $(2p+1)^2$ search locations in this region. The best matching block R is defined as a block at location $(x+i,\ y+j)$ in the reference frame for which the mean absolute error (MAE) between its pixels and the pixels in C is minimum. If the pixels in C and R are denoted as $C(x+k,\ y+l)$ and $R(x+i+k,\ y+j+l)$ for $0 \le k < M$ and $0 \le l < N$, then the MAE is defined as

$$MAE(i,j) = \frac{1}{MN}\sum_{k=0}^{M-1}\sum_{l=0}^{N-1}|C(x+k,y+l)-R(x+i+k,y+j+l)|,\quad -p \le i,j \le p \qquad (7.9)$$

The choice of N, M, and p is a compromise between accuracy and computation complexity. Small values (from 4 to 8) are preferable since the local smoothness constraint of the motion vector is easily met, while bigger values are more efficient for fast algorithms [74].

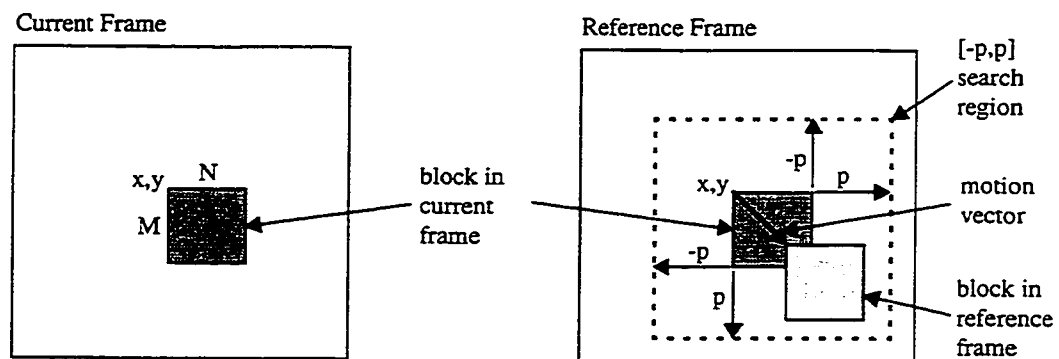

Figure 7.6 Motion Estimation

Algorithm 7.9 and Figure 7.7 illustrates the CRAM implementation of full-search motion estimation of 1024 x1024 8-bit picture frames using a block size of 4 pixels. The CRAM execution time of 19.39 ms represents a speedup of 1437 over the PC (27.87 s) and 1089 over the workstation (21.12 s). Note that computations in the 3 x 3 block neighborhood requires that data be shifted from the neighboring PEs. This uses a shift algorithm similar to that used for pixel neighborhoods (Algorithm 7.3).

```
/* on uniprocessor */
for each of the (image size/block size) blocks
  for each of the (2p+1)^2 search locations around
  the block
    reset MAE to zero;
    for each of the (p x p) pixels in the block
      accumulate MAE between block pixel and the
      corresponding pixel in the search area;
    end for
    if MAE is less than the minimum MAE so far
      store MAE as minimum MAE;
      save the search location as the motion vector
      of the block;
    end if
  end for
end for
```

```
/* on CRAM */
// reference and current frames are stored one block
// per PE
for each of the 9 search blocks around current block
  if not already in temp variables
    in parallel
      shift blocks required for pixel comparison into
      temp variables;
    end in parallel
  end if
  for each search location in the search block
    in parallel
      reset MAE to zero;
    end in parallel
    for each of the (p x p) pixels in the block
      in parallel
        accumulate MAE between block pixel and
        the corresponding pixel in the search area;
      end in parallel
    end for
    in parallel
      if MAE is less than the minimum MAE so far
        store MAE as minimum MAE;
        save the search location as the motion
        vector of the block;
      end if
    end in parallel
  end for
end for
```

Algorithm 7.9 Motion Estimation

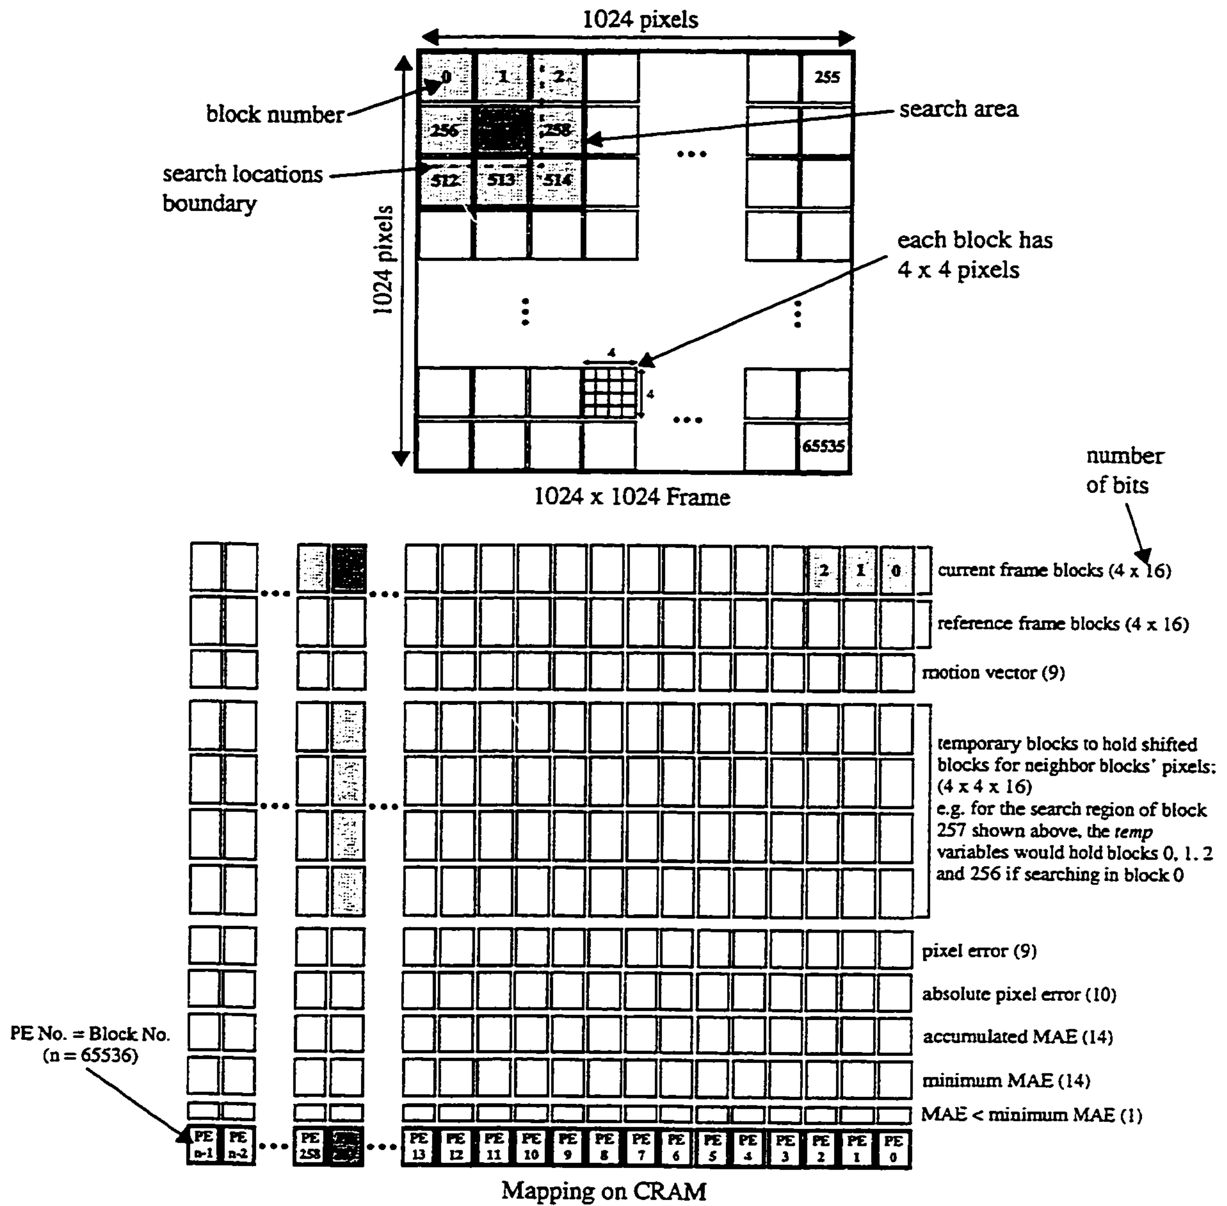

Figure 7.7 CRAM Implementation of Motion Estimation

7.6 Performance Analysis

7.6.1 Applications Performance Summary

Table 7.4 summarizes the performance of CRAM compared to the two uniprocessor systems for the applications describes in this chapter. Two cases of CRAM are considered. The first represents the case where CRAM is used as either part of the system main memory, or as the graphics or video RAM. Therefore the total execution time does not include the time to transfer application data between the host and CRAM. The second case is when CRAM is used just as an accelerator card for massively parallel applications. In this case, data is normally resident on the host and hence the total time to execute an application on CRAM has to account for the I/O overhead of transferring data between the host and CRAM. As noted before, even for the case that CRAM is not used as the system main memory or video RAM, it is not always the case that data has to be transferred between the host and CRAM for each operation. This is especially true for basic operations such as low-level image processing and basic searches because a number of such operations might need to be applied to the same data before the data is used elsewhere. For example, image enhancement operations such as brightness and contrast

Application	Work-station Execution time (ms)	PC Execution time (ms)	CRAM without I/O Overhead		CRAM with I/O Overhead	
			Execution time (ms)	Speedup over PC	Execution time (ms)	Speedup over PC
Brightness adjustment	2.72	7.44	0.0091	818	1.679	4.4
Spatial average filtering	31.78	53.13	0.3896	136	2.06	25.8
Edge enhancement	14.83	31.4	0.2474	127	1.92	16.4
Conversion to binary image	2.31	5.23	0.0093	562	1.679	3.1
Multiple-threshold segmentation	56.54	138.41	0.2686	515	1.94	71.4
Equal-to search	2.94	5.96	0.0138	431	6.694	0.89
Maximum search	3.84	6.85	0.015	457	6.695	1.02
Greater-than search	3.54	7.06	0.0128	551	6.693	1.06
Between-limits search	4.67	8.73	0.0196	445	6.7	1.3
Least means squared match	68.96	81.7	0.9686	84	14.33	5.7
Vector quantization	38222.4	74945.9	12.72	5892	19.4	3863
Motion estimation	21121.4	27869.4	19.39	1437	72.83	383

Table 7.4 Applications Performance Summary

enhancement are usually followed by filtering [11], [69]. Similarly, to find the spread of records from the maximum and minimum values in a database, both maximum and minimum searches, and some computations and record updating have to be performed on the same data. Therefore, even for the case where I/O overhead is to be included, operations might give higher speedups than the ones shown for this case in Table 7.4. Only the first operation on the data suffers the low speedup due to I/O overhead.

From Table 7.4, it is evident that CRAM yields substantial speedups over the uniprocessor for a variety of practical applications. There are higher speedups for applications that have a high degree of computations within each parallelized operation. These include Vector Quantization and Motion Estimation. Low speedups are evident in applications that either have multiplications and divisions, or use PE neighborhood computations, or have SIMD complexity of $O(n)$ or greater.

7.6.2 Controller and System Performance

Previous CRAM performance measurements [9], [4], [5] have been based on an ideal CRAM system with no control and host overhead. In this case, the reported CRAM execution times have merely been the time that the PEs are busy executing CRAM microinstructions. In a practical CRAM system, CRAM code is run from the host computer. Because of the smaller bandwidth at the host external bus, the challenge is to design a CRAM controller such that the practical CRAM system still results in significantly better performance when compared to a uniprocessor system. This section analyzes the performance of the controller and the whole system by comparing the execution times of a few practical applications on three CRAM systems: an ideal system (no control and host overhead), a practical system (described in this thesis), and an unoptimized system (without the performance-enhancement features of the controller, i.e. no instruction FIFO, no read/write buffers, and no buffer-based constant unit).

On a practical CRAM system, the total execution time for an application is made up of PE execution time, control overhead, and host overhead. PE execution time is the time that the PEs are busy executing CRAM microinstructions. This is also the total execution time on an ideal CRAM system. The other two components are defined only for the case when

no CRAM instructions are executing, i.e. when the functions described by these overheads are not executed in parallel with the execution of CRAM microinstructions. Host overheads include the time that the host executes the CRAM compiler code, sequential components of the application code, and other activities due to its interaction with CRAM. The later component depends heavily on the type of interface with the CRAM controller as will be shown later. Code execution delays on the host due to inherent activities on the system, are also bundled into host overheads. These delays are difficult to characterize. But the advantage of the simulation strategy adopted in this work (Section 6.7) is that you don't have to know or characterize these delays - they are automatically reflected into the execution time. Control overhead includes the time to load instructions into the CRAM FIFO, the overhead of filling the CRAM instruction queue, and the time to read/write data from/to the CRAM buffers and registers. Figure 7.8 shows these times for the applications described in this chapter.

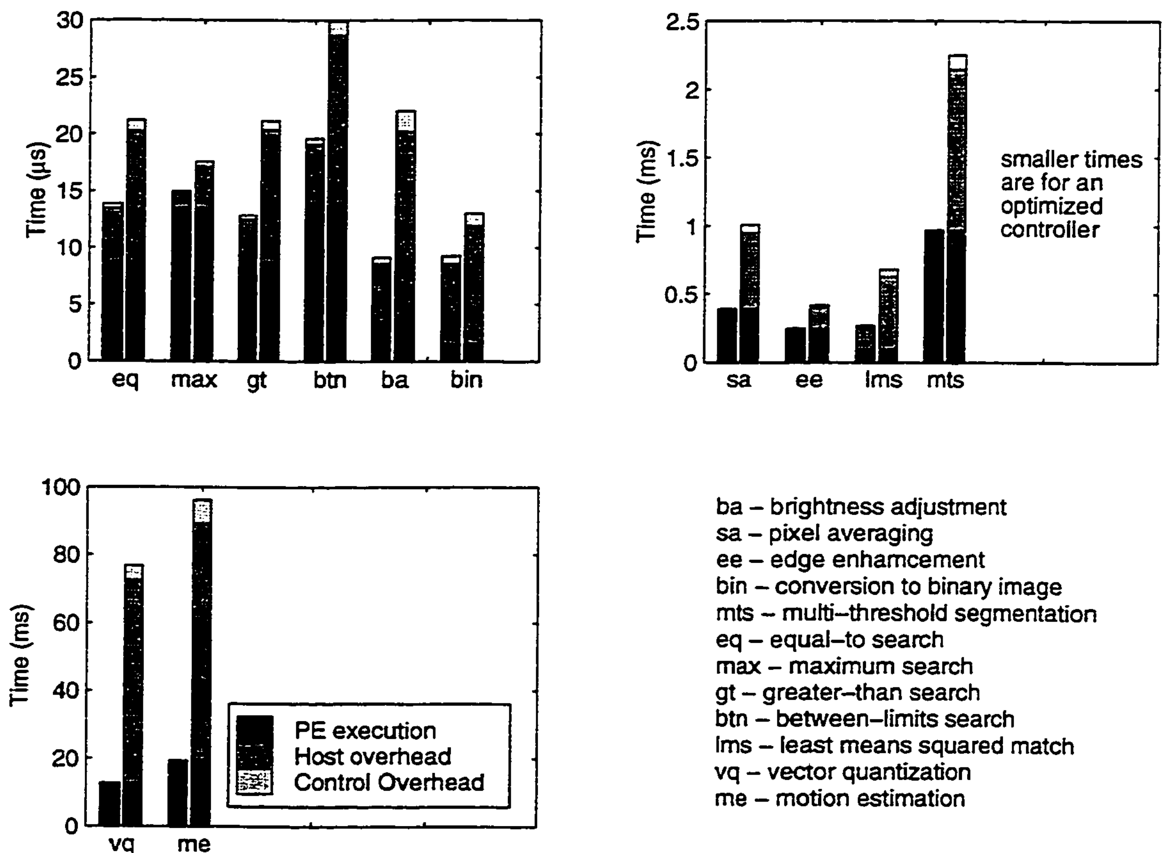

Figure 7.8 Effect of Controller on Applications Execution Times

From Figure 7.8, it can be seen that the performance-enhancement features of the controller help to reduce both the host and control overhead. Notice that the biggest component of the host overhead is the one due to the interaction of the host with CRAM. For an unoptimized controller, this component increases because the host does more data transfer retries, performs more conditional testing, executes more conditional code, sets up more data transfers, does more looping, etc. The control overhead also increases because the instruction queue is empty most of the time. This increases the overhead of filling the queue, and also means that most of the instruction-loading, transfer of data to/from CRAM registers and buffers, and other host activities are not done in parallel with the execution of CRAM instructions. Table 7.5 summarizes the impact of the controller performance-enhancement features on the performance of practical applications. For applications in which the PE execution time is very small, the difference between an ideal CRAM system and a practical CRAM system is very big because the total execution time of an application is dominated by host overheads. On the other hand, for applications with more parallel than sequential code, the CRAM system approaches an ideal system. This is good because naturally we would want to have the best CRAM performance possible when execution times are bigger. Otherwise the fact that the system efficiency is low for applications with very small PE execution time is not of much consequence because as

Application	**Ideal CRAM** Execution time (ms)	**CRAM with an Optimized Controller**		**CRAM with an Unoptimized Controller**	
		Execution time (ms)	PE Utilization (%)	Execution time (ms)	PE Utilization (%)
Brightness adjustment	0.0035	0.0091	38.46	0.0221	15.6
Spatial average filtering	0.3827	0.3896	98.22	2.1045	18.2
Edge enhancement	0.2408	0.2474	96.33	0.4152	58
Conversion to binary image	0.0015	0.0093	16.13	0.013	11.2
Multiple-threshold segmentation	0.0892	0.2686	33.21	0.6791	13.1
Equal-to search	0.0087	0.0138	63.04	0.0212	40.8
Maximum search	0.0134	0.015	89.3	0.0175	76.4
Greater-than search	0.0087	0.0128	68	0.0212	40.9
Between-limits search	0.0139	0.0196	70.92	0.0299	46.3
Least means squared match	0.9651	0.9686	99.64	2.2504	42.9
Vector quantization	12.4	12.72	97.5	76.69	16.2
Motion estimation	19.17	19.39	98.87	96.14	20

Table 7.5 Controller and System Performance

shown in Table 7.4, the actual CRAM speedup for such applications is even higher than for applications of equivalent complexity but higher system efficiency. The combination of high speedup and low system efficiency is just an indication that running the application on CRAM has reduced tremendously the execution time of the parallelized code such that the total execution time is now dominated by the sequential code and other host overheads.

7.6.3 Degree of Parallelism

As mentioned earlier, the performance of CRAM stems from the massive number of processors rather than the computational power of the individual processors. Therefore, it is important to assess how the number of PEs affect the performance of an application. Figure 7.9 shows how the CRAM speedup over the PC changes as the application is parallelized over a varying number of PEs. In this analysis, the vector size is matched to the number of PEs.

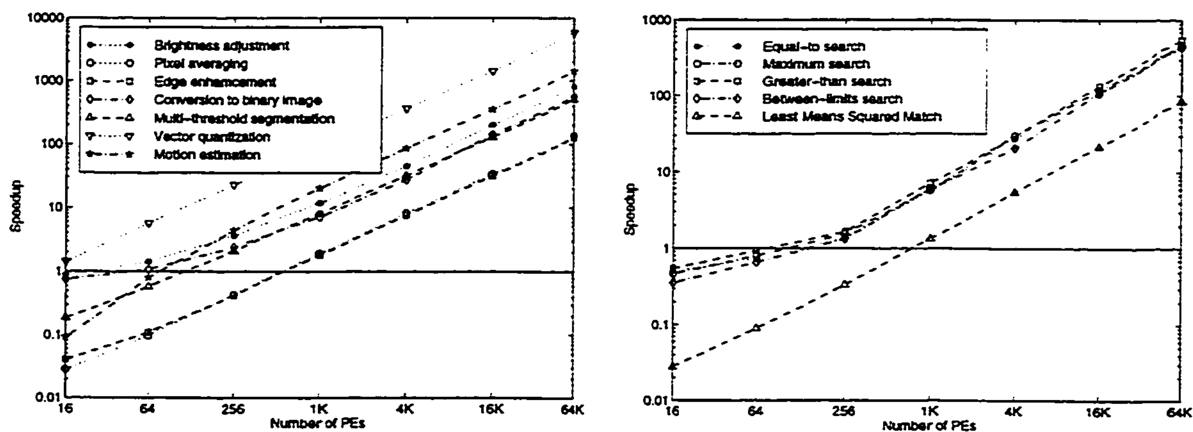

Figure 7.9 Effect of Degree of Parallelism on Applications Performance

There are three points to note from Figure 7.9. First, because of the high overhead of inter-processor communication, applications in which computations involve data from neighbor PEs exhibit lower speedups when compared to applications of similar complexity but with pointwise operations. Second, as expected, CRAM performs better than the uniprocessor PC if the number of PEs is large. However, what is important is to

determine the number of PEs at which this speedup is significant. It is evident from Figure 7.9 that this depends on the type and complexity of the application. Applications such as Vector Quantization and Motion Estimation, that have more operations within the parallelized code, generally attain reasonably significant speedup (greater than 10) at relatively smaller number of PEs (greater than 256). On the other hand, the more basic operations of low-level image processing and database searches attain such speedups at PE numbers greater than 4K. Lastly, while for large number of PEs the CRAM speedup of almost all applications increases linearly with the number of PEs, this is not the case when simple operations are parallelized over a small number (16 to 128) of PEs. This area of operation is magnified in Figure 7.10. In this case, increasing the number of PEs does not result into a proportional increase in speedup. The reason for this is that for such basic operations, the uniprocessor execution time for the main operations in the code (addition, assignment, etc.) is very small when compared to other host overheads. Therefore, the total execution time for the application becomes dominated by these overheads and hence increases slowly with the number of vectors. Also note that for this case, CRAM either performs worse than, or yields very small or no speedup, compared to the uniprocessor. In other words, the execution time of the parallelized code is so small such that the control and host overheads make the total execution time on CRAM almost equal to or bigger than that on the uniprocessor. Therefore, a CRAM system with less than 1K PEs would not be suitable for general-purpose parallel-processing, but could still be used for specific applications such as Vector Quantization and Motion Estimation.

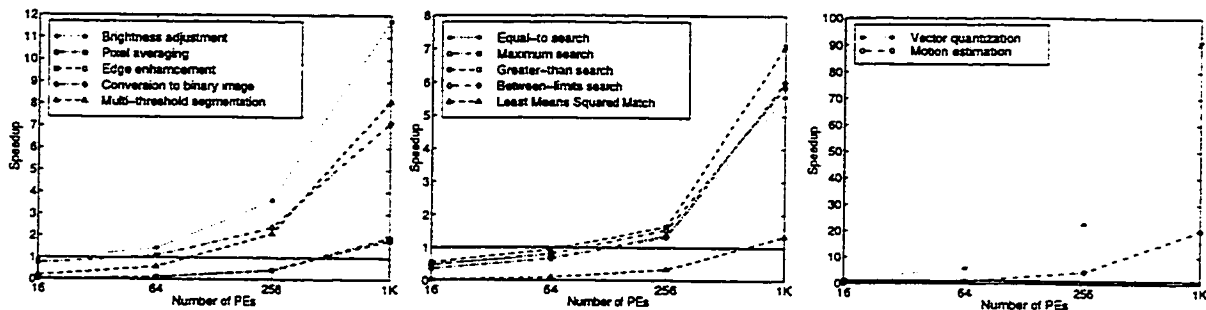

Figure 7.10 CRAM Speedup for Small Number of PEs

7.6.4 Comparison with other SIMD Systems

The main objective of CRAM is to allow conventional uniprocessor systems to attain high performance speedups when executing massively parallel applications, but with hardware far less than that of standard supercomputers. For this reason, the performance of CRAM has been compared consistently with uniprocessor systems, especially the PC. However, in order to justify the cost-effectiveness of CRAM, it is important to also know how CRAM compares with other SIMD and PIM systems. Table 7.6 shows the performance and other characteristics of two of the most popular SIMD supercomputers (the Connection Machine, CM-2, and the Massively Parallel Processor, MPP), four processor-in-memory (PIM) systems, and a 64K-PE CRAM. Performance measures used are those given in the literature for these systems. For many systems, these are in terms of Giga (10^9) Instructions per Second (GIPS) for 8-bit addition and multiplication. Where this is not available, the system application-specific performance measure is used.

From Table 7.6, we see that CRAM performance compares favorably with that of SIMD supercomputers, and CRAM has much less hardware than these systems. When compared to other PIM systems, CRAM has comparable or better performance, has the smallest hardware for all reported large-scale (greater than 4K PEs) PIM systems, and is the only general-purpose PIM system that has been implemented on the most widely used computer platform, namely the PC. It is also interesting to note that even when compared to the highly application-specific bit-parallel PIM systems such as the IMAP and FMPP-VQ64, CRAM performance is still comparable or even better, and it has less hardware when viewed on the basis of performance per PE area.

SYSTEM	No. of PEs	RAM bits/ PE	Inter-PE Communi-cation	PE Complexity	System Physical Attributes	PERFORMANCE				Types of Appli-cations
						8-bit add-ition (GIPS)	8-bit multip-lication (GIPS)	Others	Cycle Time (ns)	
MPP [6]	16K	1K	2-D grid	1000 transistors/ PE 1b PE, 6 1-bit regs, 1 full adder, shift reg & logic	large cabinet with 88 PCBs and 2100 PE chips plus memory chips	6.5	1.9	--	100	image processing
CM-2 [73]	64K	64K	2-D grid, hypercube network, router	1 PE chip has 16 PEs and is about 14,000 gates; 1b PE	large 1.5m x 1.5m cabinet, 152 PCBs, 4096 PE chips, 4096 floating-pt chips, 22528 RAM chips	4	3.3	--	--	general purpose
MIT-PP [14]	16K	128	2-D grid	1b PE with 8-to-1 mux, 5 latches	4 PCBs, 4 PE-RAM chips, 3 64Kbx32 RAM chips, 2 CPLDs, 10 reg chips, 16 64Kbx4 chips, 2 FPGAs	4.3	0.5	--	60	image processing
SRC-PIM [13]	32K	2K	parallel prefix & partitioned-	1b ALU, 3 regs, 3 muxes, & logic	external enclosure, 16 PE boards with 512 chips, 2 interface boards	--	--	3.2×10^{11} bit operations/s (BOPS)	100	general purpose
IMAP [11]	512	32K	1-D (shift left/right)	7000 transistors, 8b ALU & shifter, 14 8b regs,	1 board with 8 PE chips, 1 control chip, 1 video PCB, 1 host PCB, 2 mem chips	--	--	3x3 average filter on 512 x 512 in 0.42 ms	25	image processing
FMPP-VQ64 [17]	64	128	--	12b ALU with 3 12b registers	1 PC card with 1 PE chip and 1 control chip	--	--	53,000 nearest neighbor search (NNS)/s	40	VQ
CRAM	64K	4K	1-D (shift left/right)	1b PE with 75 transistors (8-to-1 mux, 3 1b regs)	1 PC card with 8 or 16 PE-RAM chips, 1 control chip or 1 (FPGA + CPLD + EPROM)	27	2	1.3×10^{12} BOPS 82×10^{6} NNS/s; 512x512 filter in 0.8 ms	50	general purpose

Table 7.6 System Comparisons

7.7 Summary

This chapter has looked at CRAM applications and analyzed the performance of the CRAM system. With regard to applications, the objective has not been to present an exhaustive review of all possible applications, but rather to use a few selected applications of varying execution models and complexity to study different CRAM performance issues. Using this approach, it has been shown that the type of applications most suited for CRAM are those that have fine grain parallelism and regular communication patterns. The implementation of different application structures on CRAM has also been demonstrated. For example, the use of the CRAM 1-D inter-PE communication network to implement 2-D communications during PE neighborhood computations has been described. This has also highlighted the inter-PE network as a major architectural bottleneck on CRAM. Even with this bottleneck, CRAM yields significant performance speedup over conventional uniprocessor systems when executing these massively parallel applications. The impact of the CRAM controller on the performance of a CRAM system has also been analyzed. In this regard, it has been shown that even with a bandwidth-limited host, a CRAM system still yields reasonably high performance because of the performance-enhancement features of the controller. Finally, it has been shown that even with less hardware and a simpler architecture, the performance of a CRAM system compares favorably with that of conventional supercomputers. Also, CRAM performs either better, slightly worse, or almost the same when compared to other (mostly application-specific) logic-in-memory systems. However, it has less hardware and is less complex than these systems.

Chapter 8

Conclusions and Future Work

This thesis describes the system design issues for a computational-RAM logic-in-memory parallel processing machine. This includes the architecture and implementation of the controller for the CRAM processing elements, the interface of CRAM to the host computer, and the development of CRAM system software tools. This chapter provides a summary of the thesis and presents some ideas for future work.

8.1 Summary

Integrating several 1-bit processing elements at the sense amplifiers of a standard RAM improves the performance of massively-parallel applications because of the inherent parallelism and high data bandwidth inside the memory chip. However, implementing such a system on a host computer such as the PC poses several challenges in controlling the processing elements, exchanging data between the host and CRAM, and writing CRAM application programs. This is so because of the small bandwidth at the host system buses, the differences in the bit-serial and bit-parallel data formats on CRAM and the host, respectively, and the constrained size of a PC expansion card. The CRAM controller and software tools developed in this thesis provides solutions to these system design challenges.

The area of the controller has been minimized by using a simplified microprogram

sequencer and grouping microinstructions such that the size of the control store is halved. This reduces system cost, allows easy implementation of a CRAM system on a single PC card, and facilitates future integration of the controller and the PE array on the same chip. Several architectural features are used to enhance the performance of PE control and interaction with the host computer. A FIFO-based instruction queue unit improves PE utilization, especially for instructions with a small number of microinstructions. Buffers reduce the time of transferring data between the host and CRAM, simplifies synchronization, and eases electrical and physical loading of the host bus. A new constant broadcast unit improves the performance of operations with constants, reduces the size of the control store, and simplifies the use of variable-size constants. A new parallel array-based data corner-turning scheme greatly reduces the time of software data transposition, thus removing the need for the high-cost hardware used to convert data between bit-serial and bit-parallel formats. These performance-enhancement features also minimize the effect of the host on the performance of CRAM. This allows the implementation of CRAM on a variety of platforms, including those with slow host buses and processors such as ISA computers and embedded systems with slow microcontrollers. In general, the high-performance, minimal-hardware and general-purpose architecture of the CRAM controller enhances the use of CRAM as a general-purpose parallel-processing system.

PCI and ISA CRAM controller prototypes have been implemented in a Xilinx XC4013EPQ240-2 FPGA. A 64-PE ISA CRAM system prototype has been built and tested as an expansion card in a 133 MHz Pentium PC under both Linux and MSDOS. These prototypes are used to demonstrate working models of the CRAM concept.

Four high-level software tools are used to hide low-level architectural details of the CRAM system, allowing software developers and system analysts to focus on implementation details. The CRAM C++ compiler, which is simply a library of CRAM parallel variables, is used to write CRAM programs using the standard C++ language and standard C++ compilers. The CRAM assembler provides optional low-level control of CRAM hardware. The CRAM microcode assembler is used for generating CRAM microinstructions. The CRAM C++ simulator simulates the behavior of a CRAM system and can be used by both hardware and software designers to analyze CRAM architectural features as well as the performance and implementation of applications. A CRAM system

VHDL simulation model allows logic designers to easily simulate the behavioral, gate-level, or post-place-and-route VHDL designs of the controller and CRAM chips by running the actual CRAM system assembly or machine code.

Comparing a 20 MHz 32 Mbytes 64K-PE PCI CRAM system with a 133 MHz 32 Mbytes Pentium PC and a 167 MHz 64 Mbytes SUN Sparc Ultra workstation show that CRAM yields significant performance speedup over conventional uniprocessor systems when executing massively-parallel applications. Comparisons with conventional supercomputers and other logic-in-memory systems show that CRAM has comparable speed but uses less hardware and has a less complex architecture.

8.2 Future Work

Two new ideas are suggested for the continuation of this work. These are an on-chip CRAM controller and a MIMD-SIMD CRAM system, as well as a pipelined constant-sensitive PE architecture. Other ideas for general future work are also suggested in this section.

8.2.1 An On-Chip Controller and a MIMD-SIMD CRAM System

With the general architectures of CRAM and its controller now established, another interesting step is to explore the possibility and analyze the pros and cons of putting the controller on the same chip as the PE/memory array. One obvious disadvantage is that CRAM now becomes less compatible with standard RAM, both in fabrication (more logic) and pinout (an extended pinout to support system bus signals). But there are many potential gains of this approach. The data access time between the controller and CRAM may be significantly reduced. Part of the CRAM memory, say the first 4 Kbytes, may be used as the control store. This would reduce the size of the controller logic, and may improve performance by dynamically changing the size of the microprogram memory (on a program-by-program basis) to allow more application-specific functions/algorithms to be implemented as microcode. An on-chip controller also reduces the number of chips on the CRAM system PCB. This may reduce system cost and also increase the CRAM size (number of CRAM chips) that can be implemented on a single card.

An on-chip controller also gives rise to a new possible architecture in which two or more CRAM chips are connected to form a MIMD-SIMD CRAM system. This is illustrated Figure 8.1. Each CRAM chip is connected to the host bus, possibly memory-mapped to a specific host address space, and therefore may operate on a different instruction/program (in normal SIMD stye) than the other chips. A 1-bit register on the controller may allow the chip(s) to be connected as one big SIMD array or as components of a MIMD system. Other issues to be studied include inter-controller communication and the number of PEs per controller (such as transistors per controller versus number of PEs x transistors per PE).

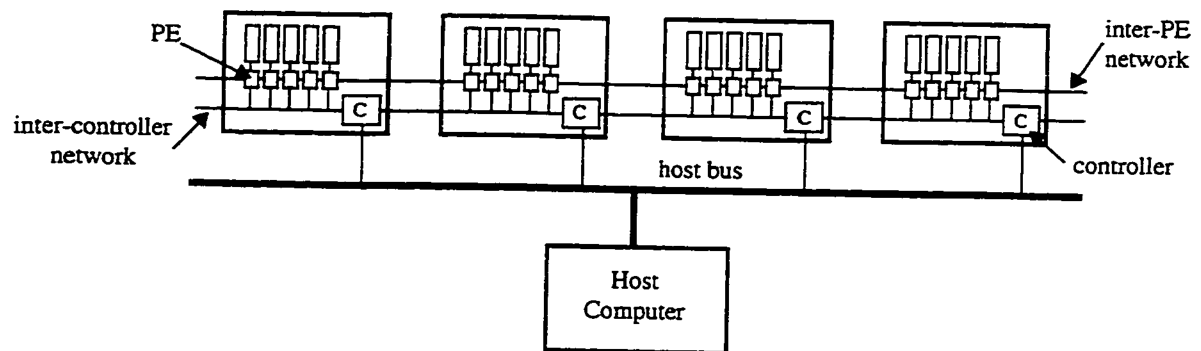

Figure 8.1 A MIMD-SIMD CRAM System

8.2.2 A Pipelined Constant-Sensitive PE

With an on-chip controller, the address and data buses of CRAM would no longer be multiplexed with PE control opcodes (TTOP/COP). This would allow an internal PE memory access to be executed in parallel with a PE operation. This can be accomplished by putting a latch at the output of the PE. This pipelining would reduce both the execution time and the size of the microprogram memory. The PE may also be enhanced to directly support operations with constants by multiplexing one or all three PE source registers with the bit coming from the constant broadcast unit. This would remove the need for an LDK instruction, thus reducing the execution of n-bit operations by n clock cycles. It may also slightly reduce the size of the control store.

Other improvements not related to an on-chip controller involves the inter-PE communication network. In the current architecture, the output of a PE shift operation is

hardwired to a specific register of the neighbor PE. This may be improved by making the output of a PE during shift operations equal to the output of its neighbor PE. This would allow the shifted value to be written to any of the three destination registers, thus reducing the number of microinstructions in such routines as data transposition. The inter-PE communication network may also be upgraded to a 2-D network. All these PE improvements will have to be carefully weighted against the increase in the PE area so that the total area of computational logic is still a small fraction (may be less than 15%) of the total area of the PE/RAM array.

8.2.3 Other Work

There are several other ideas for the continuation of this work. A bigger CRAM system prototype needs to be built so that more applications can be tested. This may require fabricating bigger CRAM chips, with the errors in the current chips corrected (A critical error is the bus-tie because it allows CRAM chips to be connected together to form bigger systems). An ASIC controller may also be fabricated for use in the prototype.

The software tools, especially the C++ library, also need to optimized so that the time that the host processor takes to execute this code is minimized. This might require changing the definition, functionality, hierarchy or code-ordering of some classes, or simply applying standard code optimizations, such as loop unrolling and code in-lining. The library may also be extended to support floating-point variables.

On the architecture and implementation of the controller, several things could be done. The control store could be enlarged so that more frequently used functions are implemented as microcode. The instruction queue unit could be extended to support looping and/or branching (with possible significant increase in the size of the controller). Variable-length instructions, or other instruction formats, could be explored. Finally, for embedded application-specific uses, an architecture using a hard-wired controller could be implemented with most of the general-purpose features stripped out. CRAM systems targeted for other buses, such as AGP and Sbus could also be designed.

References

[1] Harold S. Stone, "A Logic-in-Memory Computer", *IEEE Transactions on Computers*, pp 73-78, January, 1970.

[2] J. L. Hennessy and D. A. Patterson, *Computer Organization and Design*, Second Edition, Morgan Kaufmann Publishers, San Francisco, 1997.

[3] D. G. Elliott, M. Snelgrove, and M. Stumm, "Computational RAM: A Memory-SIMD Hybrid and Its Application to DSP", *Custom Integrated Circuits Conference*, May, 1992.

[4] D. G. Elliott, *Computational RAM: A Memory-SIMD Hybrid*, Ph.D. Thesis, University of Toronto, December, 1997.

[5] Thinh M. Le, *Visual Communications on a Memory-Embedded Array Processor: The Computational RAM*, Ph.D. Thesis Proposal, University of Ottawa, 1998.

[6] J. L. Potter, *The Massively Parallel Processor*, MIT Press, Cambridge, 1985.

[7] Harold S. Stone, *High-Performance Computer Architecture*, Addison-Wesley Publishing Co., 1990.

[8] D. G. Elliott, W. M. Snelgrove, C. Cojocaru, and M. Stumm, "A PetaOp/s is currently feasible by Computing in RAM", *PetaFLOPS Frontier Workshop*, February, 1995.

[9] R. McKenzie, *A DRAM-Compatible Parallel Processor for Real-Time Video*, Masters Thesis, Carleton University, December, 1997.

[10] Noboyuki Yamashita, et al, "A 3.84 GIPS Integrated Memory Array Processor with 64 Processing Elements and a 2-Mb SRAM", *IEEE Journal of Solid-State Circuits*, Vol. 29, No. 11, pp 1336-1343, November, 1994.

[11] Shin'ichiro Okazaki, et al, "A Compact Real-Time Vision System Using Integrated Memory Array Processor Architecture", *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 5, No. 5, pp 446-452, October 1995.

[12] Yoshiharu Aimoto, et al, "A 7.68 GIPS 3.84 GB/s 1W Parallel Image-Processing RAM Integrating a 16 Mb DRAM and 128 Processors", *IEEE International Solid-State Circuits Conference*, pp 372-373, February, 1996.

[13] M. Gokhale, B. Holmes, and K. Iobst, "Processing in Memory: The Terasys Massivel Parallel PIM Array", *IEEE Computer*, pp 22-31, April, 1995.

[14] J. C. Gealow, F. P. Herrmann, L. T. Hsu, and C. G. Sodini, "System Design for Pixel-Parallel Image Processing", *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, pp 32-41, March, 1996.

[15] David Patterson, et al, "A Case for Intelligent RAM", *IEEE Micro*, pp 34-44, April,1997.

[16] Kazutoshi Kobayashi, et al, "A Memory-Based Parallel Processor for Vector Quantization: FMPP-VQ", *IEICE Trans. Electron*, Vol. E80-C, No. 7, pp 970-975, July, 1997.

[17] Kazuhiko Terada, "An LSI for Low Bit-Rate Image Compression Using Vector Quantization", *IEICE Trans. Electron*, Vol. E81-C, No. 5, pp 1-7, May, 1998.

[18] R. Torrance, et al, "A 33 GB/s 13.4 Mb Integrated Graphics Accelerator and Frame Buffer", *IEEE International Solid-State Circuits Conference*, pp 340-341, February,

1998.

[19] D. G. Elliott and W. M. Snelgrove, "CRAM: Memory with a Fast SIMD Processor", *Proceedings of the Canadian Conference on VLSI*, pp 3.3.1-3.3.6, October, 1990.

[20] D. G. Elliott and W. M. Snelgrove, *Method and Apparatus for a Single Instruction Operating Multiple Processors on a Memory Chip*, U. S. Patent 5546343, Issued August 1996, Filed October, 1990.

[21] D. G. Elliott, M. Stumm, and W. M. Snelgrove, "Computational RAM: The Case of SIMD Computing in Memory", Workshop on Mixing Logic and DRAM, ISCA '97, pp 6.1-6.7, June, 1997.

[22] Christian Cojocaru, *Computational RAM: Implementation and Bit-Prallel Architecture*, Master's Thesis, Carleton University, January, 1995.

[23] H. L. Kalter, et al, "A 50-ns 16-Mb DRAM with a 10-ns Data Rate and On-Chip ECC", *IEEE Journal of Solid-State Circuits*, Vol. 25, No. 5, October, 1990.

[24] Betty Prince, *Semiconductor Memories*, Second Edition, Chichester, England, John Wiley and Sons, 1996.

[25] Peter M. Nyasulu and W. Martin Snelgrove, "Architecture and Implementation of a Computational RAM Controller", *International Conference on Massively Parallel Computing Systems*, Colorado Springs, April, 1998.

[26] W. M. Loucks, M. Snelgrove, and S. G. Zaky, "A Vector Processor Based on One-Bit Microprocessors", *IEEE Micro*, pp 53-62, February, 1982.

[27] C. M. Habiger and R. M. Lea, "Hybrid-WSI: A Massively Parallel Computing technology?", *Computer*, Vol. 26, No. 4, pp 50-61, April, 1993.

[28] Kenneth E. Batcher, "Design of a Massively Parallel Processor", *IEEE Transactions on Computers*, Vol. C-29, No. 9, pp 836-840, September, 1980.

[29] Tom Blank, "The MasPar MP-1 Architecture", *Proceedings of the IEEE COMPCON Spring*, pp 20-24, February, 1990.

[30] J. Mick and J. Brick, *Bit-Slice Microprocessor Design*, McGraw-Hill Book Company, 1980.

[31] C. Fernstrom, I. Kruzela, and B. Svensson, *LUCAS Associative Array Processor*, Springer-Verlag, 1985.

[32] John R. Nickolls, "The Design of the MasPar MP-1: A Cost-Effective Massively Parallel Computer", *Proceedings of the IEEE COMPCON Spring*, pp 25-28, February, 1990.

[33] W. D. Hillis, *The Connection Machine*, Cambridge Massachussetts, MIT Press, 1985.

[34] Tom Shanley and Don Anderson, *PCI System Architecture*, Third Edition, Addison-Wesly Co, 1995.

[35] Warren Andrews, "PCI: The Universal Socket", *RTC Magazine*, Vol. VI, No. 2, pp 19-23, February, 1998.

[36] *PCI Local Bus Specification*, Rev. 2.1, June, 1995.

[37] *Accelerated Graphics Port Interface Specification*, Rev. 1.0, Intel Corporation, July, 1996.

[38] Edward Solari, *ISA & EISA Theory and Operation*, Annabooks, San Diego, CA, 1994.

[39] Wade D. Peterson, *The VME Handbook*, Third Edition, VITA, Scottdale, AZ, 1993.

[40] *Using Select-RAM Memory in XC4000 Series FPGAs*, Xilinx Application Note,

July, 1996.
[41] Jeffrey C. Gealow, *An Integrated Computing Structure for Pixel-Parallel Image Processing*, PhD Thesis, Massachusetts Institute of Technology, June 1997.
[42] Peter M. Nyasulu, Ralph Mason, W. Martin Snelgrove, and Duncan Elliott, "Minimizing the Effect of the Host Bus on the Performance of a Computational RAM Logic-in-Memory Parallel-Processing System", to be presented at *IEEE Custom Integrated Circuits Conference*, San Diego, May, 1999.
[43] R. G. Lange, "High-Level Language for Associative and Parallel Computation with STARAN" *Proceedings of the International Conference on Parallel Processing*, pp 170-176, August, 1976.
[44] Peter Christy, "Software to Support Massively Parallel Computing on the MasPar MP-1", *Proceedings of the IEEE COMPCON Spring*, pp 29-33, February, 1990.
[45] Bjarne Stroustrup, *The C++ Programming Language*, Third Edition, Addison-Wesley, Reading, Massachussetts, 1997.
[46] Kenneth E. Batcher, "The Multidimensional Access Memory in STARAN", *IEEE Transactions on Computers*, February, 1977, pp174-177.
[47] Kenneth E. Batcher, "The Flip Network in STARAN", *IEEE Transactions on Computers*, February, 1977, pp 65-71.
[48] A. Abnous, et al, "Design and Implementation of the 'Tiny RISC' Microprocessor", *Microprocessors and Microsystems*, Vol. 16, 1992, pp 187-193.
[49] S. R. Wang, et al, "A 66 MHz PA-RISC Embedded Controller with 80486DX-Like Bus Interface", *IEEE COMPCON*, Spring, 1994, pp 53-57.
[50] J. Schack, et al, "KISS-16V1: A 16 Bit Signal Processor", *IEEE International Symposium on Circuits and Systems*, Vol. 4, 1991, pp1905-1908.
[51] Peter M. Nyasulu, *Microprocessor Design for Instrument Control*, Masters Thesis, Carleton University, August, 1995.
[52] D. Gajski, "Introduction to High-Level Synthesis", *IEEE design and Test of Computers*, Winter, 1994, pp 45-54.
[53] V. Nagasamy, et al, "Specification, Planning, and Synthesis in a VHDL Design Environment", *IEEE Design and Test of Computers*, 1992, pp 58-68.
[54] C. Jay, "VHDL and Synthesis Tools Provide a Generic Design Entry Platform into FPGAs, PLDs, and ASICs", *Microprocessors and Microsystems*, Vol. 17, Sept., 1993, pp 391-398.
[55] Bob Reese, "Use of VHDL Synthesis in an Advanced Digital Design Course", *IEEE SouthEast Con'92*, Vol. 2, 1992, pp 509-512.
[56] L. Clonts, et al, "Writing Area-Efficient Hardware Descriptions for Logic Synthesis", *IEEE SouthEast Con'92*, Vol. 2, 1992, pp 505-508.
[57] Darren Jones, et al, "Verification Techniques for a MIPS Compatible Embedded Control Processor", *IEEE International Conference on Computer Design*, Oct., 1991, pp 329-332.
[58] Tim Brodnax, "The PowerPC 601 Design Methodology", *IEEE International Conference on Computer Design*, October, 1993, pp 248-252.
[59] Kevin Skahill, *VHDL for Programmable Logic*, Addison-Wesley Publishing, CA, 1996.
[60] *Programmable Data Book*, Xilinx Inc., 1990.
[61] *PCI S5920 developer's Kit User Manual and Technical Reference Manual*, Applied

Micro Circuits Corporation, Revision 1.3, April, 1998.

[62] *S5920 PCI Target Interface Data Book*, Applied Micro Circuits Corporation, 1997.

[63] Tom Shanley and Don Anderson, *ISA System Architecture*, Third Edition, Addison-Wesly Co, 1995.

[64] Tom Shanley, *EISA System Architecture*, New Revised Edition, Mindshare Press, Richardson, TX, 1993.

[65] Thinh M. Le and S. Panchanathan, "Computational-RAM Implementation of an Adaptive Vector Quantization for Video Compression", *IEEE Transactions on Consumer Electronics*, Vol. 41, No. 3, pp 738-747, August, 1995.

[66] Thinh M. Le, S. Panchanathan, and W. M. Snelgrove, "Computational-RAM Implementation of Mean-Average Scalable Vector Quantization for Real-Time Progressive Image Transmission", *Proceedings of the 1996 Canadian Conference on Electrical and Computer Engineering*, Vol. 1, pp 442-445, May, 1996.

[67] T. M. Le, W. M. Snelgrove, and S. Panchanathan, "Computational RAM Implementation of MPEG-2 for Real-Time Encoding", *SPIE Proceedings of Multimedia Hardware Architectures '97*, pp 182-193, February, 1997.

[68] T. M. Le, W. M. Snelgrove, and S. Panchanathan, "Fast Motion Estimation Using Feature Extraction and XOR Operations", *SPIE Proceedings of Multimedia Hardware Architectures '98*, January, 1998.

[69] Zahid Hussain, *Digital Image Processing: Practical Applications of Parallel Processing Techniques*, Ellis Horwood Limited, Chichester, 1991.

[70] Loren Heiny, *Advanced Graphics Programming Using C/C++*, John Wiley & Sons Inc., 1993.

[71] Anil K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1989.

[72] Robert Heaton, Donald Blevins, and Edward Davis, "A Bit-Serial VLSI Array Processing Chip for Image Processing", *IEEE Journal of Solid-State Circuits*, Vol. 25, No. 2, April, 1990.

[73] Lewis W. Tucker and George G. Robertson, "Architecture and Applications of the Connection Machine", *Computer*, Vol. 21, No. 8, pp 26-38, August, 1988.

[74] V. Bhaskaran and K. Konstantinides, *Image and Video Compression Standards: Algorithms and Architectures*, Second Edition, Kluwer Academic Publishers, Boston, 1997.

[75] K. Hwang and F. Briggs, *Computer Architecture and Paralell Processing*, McGraw-Hill Inc., New York, 1984.

[76] A. Bjorck, R. J. Plemmons, and H. Schneider, *Large Scale Matrix Problems*, Elsevier North Holland Inc., New York, 1981.

Appendix A

CRAM Controller Architectural Details

A.1 Microinstruction Word

This appendix describes bit settings for the CRAM microinstruction word shown in Figure 4.8.

A.1.1 Next Address Instruction (NAI)

This 4-bit field (Table A.1) selects the next address of the control store (i.e. the next microinstruction).

Next address Instruction (NAI)	Control Word Bits [31:28]	Description
NEXTI	0000	Go to the first microinstruction of the instruction in IR. This also implements the jump to address zero (JZ) or WAIT/NOP.
NEXTu	0001	Go to the next sequential microinstruction (at address uPC+1).
SETLP	0010	Set looping to begin at next microinstruction (LPR ← uPC+1).
LOOPN	0011	Loop back to the address in LPR, else next microinstruction.
LOOPE	0100	Loop back to the address in LPR, else end.
JSR	0101	Conditional jump to subroutine; save return address (SBR ← uPC+1).
JMP	0110	Conditional jump to address in microinstruction bits [7:0].
	0111-1111	Reserved.

Table A.1 Next Address Instructions

A.1.2 Condition Select (COND)

This field (Table A.2) selects condition to be used in a conditional jump or loop instruction.

Condition	Bits [27:25]	Description	Use
NOCND	000	Always true	unconditional execution
CNZ	001	Loop Counter is not zero	n-looping
CN1	010	Loop counter is not one	(n-1)-looping
GOR	011	CRAM chip Global-OR (Bus-Tie) status	general
SLO	100	CRAM chip shift left output of (n-1)th PE	general
SRO	101	CRAM chip shift right output of PE 0	general
DZ	110	Byte read from CRAM contains all zeroes	max/min searches, etc.
	111	Reserved	Reserved

Table A.2 Microsequencer Conditions

A.1.3 CRAM Function and Read/Write Bit

Microinstruction bits [23:22] indicates the CRAM function being executed in the current microinstruction. the settings are shown in Table A.3. Even though a NOP means that there is no operation on the CRAM chip, a controller instruction (CCI and ACCI, described in the next section) may be embedded in the microinstruction. Bit [21] indicates a read (1) or a write (0).

Bits [23:22]	CRAM Function
00	No Operation (NOP)
01	PE Operation
10	Internal CRAM read/write (data is read from CRAM memory to PE register M, or written from the PE ALU output to CRAM memory)
11	External CRAM read/write (transfer of data between CRAM and controller)

Table A.3 CRAM Functions

A.1.4 External TTOP (EXTTOP)

If EXTTOP is set, TTOP is selected to be the value of operand 2 (OPR2) of the instruction word. Otherwise TTOP is selected as bits [7:0] of the current microinstruction.

A.1.5 Address Select (ASEL)

Bits [20:19] select the value to be put on the CRAM chip address bus. They can be set to select either AR0 (ASEL="00"), AR1 ("01"), AR2 ("10") or bits [15:8] ("11"). Bits [15:8] represent either COP or a system mask/temporary address (TMPA).

A.1.6 CRAM Controller Instructions (CCI and ACCI)

CRAM Controller instructions (CCI) are coded in microinstruction bits [18:16]. These are mainly used for loading parameter registers or incrementing address registers. Table A.4 describes these instructions. To minimize the size of the microinstruction word, all controller instructions that do not execute simultaneously with a CRAM instruction are coded on bits [3:0]. These auxiliary controller instructions (ACCI) are decoded only if CCI is set to ACCI (i.e. bits [18:16] = "111"). These are described in Table A.5.

Instructio (CCI)	Bits [18:16]	Description
NOCCI	000	No controller instruction
INCA	001	After using it, increment the currently-selected address register
DECA	010	After using it, decrement the currently-selected address register
INCB	011	After using it, increment the bank address register
LDK	100	Load a constant's bit into a PE register
-	101	Reserved
XCOP	110	COP is from external (value of instruction operand 1)
ACCI	111	Execute an auxiliary controller instruction

Table A.4 CRAM Controller Instructions

Auxiliary Instruction (ACCI)	Bits [3:0]	Description
	0000	Reserved
LDCBA	0001	Load CRAM bank address (CBA) register
LDIBA	0010	Load the read/write buffer internal address (WIBA & RIBA) registers
LDWLEN	0011	Load the word length (WLEN) register
LDAX0	0100	Load address extension register 0 (AX0)
LDAX1	0101	Load address extension register 1 (AX1)
LDAX2	0110	Load address extension register 2 (AX2)
SETINT	0111	Load (set) the interrupt number
LDSIC	1000	Load the shift-input selection code
LDBAI	1001	Load the buffer address increment registers
	1010-1111	Reserved

Table A.5 Auxiliary Controller Instructions

A.2 Command, Status, and Parameter Registers

A.2.1 Command Register (CCR)

The command register is memory-mapped and can be written and read by the host processor on a byte basis. Writing a 1 to a bit of the command register will toggle that bit. The implemented CCR bits are shown in Table A.6. The IRQn bits are provided in this prototype to allow flexibility in the choice of the ISA interrupt line to use (The PC interrupts can sometimes be very populated and unshareable).

Name	CCR Bit	Function	Value on Reset
Control Store Ready (CSRDY)	14	1 - The control store is ready 0 - Control store not ready	0
Extended PE (XPE)	13	1 - CRAM has extended PEs 0 - CRAM has baseline PEs	0
Bit-Parallel CRAM (BPL)	12	1 - CRAM operates in bit-parallel mode 0 - CRAM operates in bit-serial mode	0
External Timing (XCTRL)	11	1 - CRAM PEs use external timing 0 - CRAM PEs use internal timing	0
Interrupt Request Lines (IRQ15, IRQ11, IRQ10, IRQ9)	3-0	Indicates (if set) the interrupt line that the CRAM system will use. Only one may be set.	IRQ15 = 1 Others = 0

Table A.6 Command Register

A.2.2 Status Register (CSR)

The status register (CSR) is a 16-bit memory-mapped read-only register. Table A.7 shows the implemented bits of CSR. IRQP is automatically cleared when the host processor reads the byte that contains it (CSR[16:8]).

Name	CSR Bit	Function
Global-OR (GOR)	0	Value of the CRAM Bus-Tie (Global-OR)
Shift-Left Output (SLO)	1	Value shifted out of PE_{n-1} during CRAM shift-left operation.
Shift-Right Output (SRO)	2	Value shifted out of PE_0 during CRAM shift-right operation.
Instructions Executing (IXEC)	3	1 - Instructions are still executing or pending in the queue 0 - No instructions executing; the queue is empty
Interrupt Pending (IRQP)	4	1 - The CRAM system has issued an interrupt 0 - No interrupt has been issued by the CRAM system

Table A.7 Status Register

A.3 PCI Configuration Registers and Commands

This appendix describes how PCI configuration registers are implemented for CRAM. Since CRAM is not a PCI-to-PCI bridge, it implements a PCI device configuration header type zero shown in Figure A.1. The shaded area represents registers that are mandatory. **Header Type** identifies the format of the device configuration header (bits 6:0), and also indicates if the device is a single or multi-function (bit 7). The CRAM controller is a single-function PCI device (i.e. header type register bit 7 is hardwired to 0), and it implements a header type 0 (i.e. header type register bits 6:0 are hardwired to "0000000").

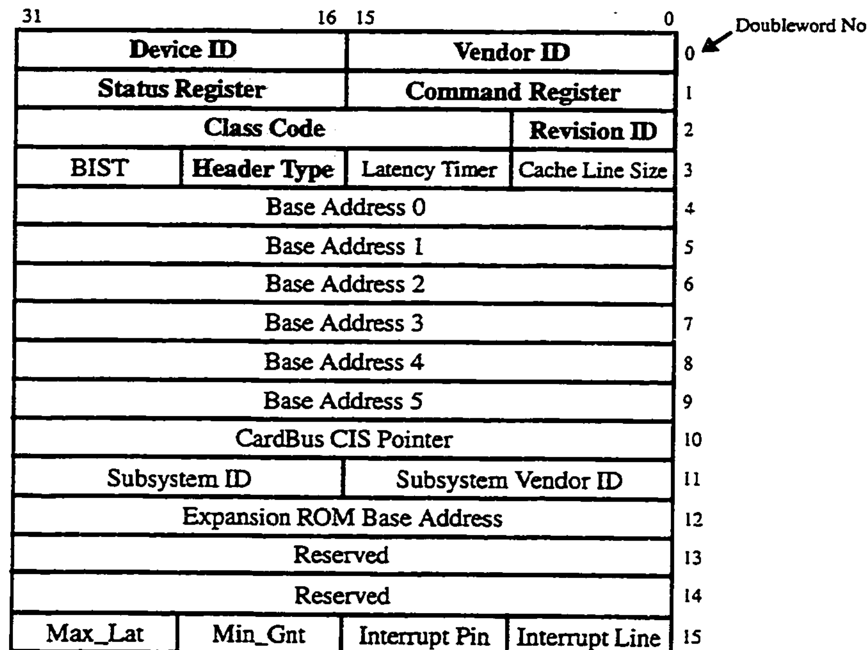

Figure A.1 PCI Device's Configuration Header

- **Vendor ID Register:** This16-bit register identifies the manufacturer of the device and is assigned by the PCI SIG (Special Interest Group). Since our design is just for prototyping purposes, an arbitrary value of 0x1008 is hardwired as the Vendor ID.
- **Device ID Register:** This is a 16-bit register. It is assigned by the device manufacturer to identify the type of the device. The CRAM device ID is hardwired to 0x0000.
- **Revision ID Register:** This 8-bit register is assigned by the device manufacturer to identify the revision number of the device. CRAM revision ID is hardwired to 0x00.
- **Command Register:** The 16-bit command register defines how a PCI device responds to PCI accesses. Table A.8 describes the nine defined bits of the command register and how they are implemented in the CRAM controller. Only bits 8, 6 and 1 are implemented as real registers. This means that for CRAM, these command bits can be set to either 1 or 0 using a PCI configuration cycle. The rest of the command bits have fixed values for CRAM implementation, and hence are just hardwired.

Bit	Function	CRAM Implementation
0	I/O Access Enable 1 - Device responds to PCI I/O accesses 0 - Disable I/O	Hardwired to 0. (All CRAM units/registers are memory-mapped)
1	Memory Access Enable 1 - Device responds to PCI memory accesses 0 - Disable	Register (RESET# to 0; set to 1 to allow memory accesses)
2	Master Enable 1 - Device acts as a bus master 0 - Device can't be master	Hardwired to 0 (CRAM only acts as a PCI target)
3	Special Cycle Recognition 1 - Device monitors special cycles 0 - Special cycles are ignored	Hardwired to 0
4	Memory Write and Invalidate Enable 1 - Master can generate memory write and invalidate 0 - Master only uses memory write commands	Hardwired to 0 (This is only for PCI bus masters)
5	VGA Palette Snoop Enable 1 - Device must perform palette snooping 0 - Disable	Hardwired to 0 (This is for VGA compatible devices only)
6	Parity Error Response 1 - Device can report parity errors (using PERR#) 0 - Device ignores parity errors	Register (RESET# to 0; set to 1 if CRAM is to report parity errors)
7	Wait Cycle Enable 1 - Device does address/data stepping 0 - No address/data stepping	Hardwired to 0
8	System Error Enable 1 - Device can report address parity errors (SERR#) 0 - Disables use of SERR#	Register (RESET# to 0; set to 1 if CRAM is to address report parity errors)
9	Fast Back-to-Back Enable 1 - Can perform fast back-to-back transactions 0 - Disable	Hardwired to 0 (This is for PCI bus masters only)
15:10	Reserved by PCI	Hardwired to "000000"

Table A.8 CRAM PCI Command Registers

- **Status Register:** This tracks the status of PCI transactions on a PCI device. Any bit of the status register can be cleared by writing a 1 to it using a PCI configuration write cycle. See Table A.9.

Bit	Function	CRAM Implementation
4:0	Reserved by PCI	Hardwired to "00000"
5	66 MHz Capable 1 - Device is capable of running at 66 MHz 0 - Device can only run at 33 MHz	Hardwired to 0
6	UDF Supported 1 - Device supports User Defined Features 0 - Doesn't support UDFs	Hardwired to 0
7	Fast Back-to-Back Capable 1 - Supports fast back-to-back transactions 0 - Isn't capable	Hardwired to 1 (CRAM can handle fast back-to-back transactions)
8	Data Parity Reported 1 - Master reported the parity error 0 - Didn't report any parity error	Hardwired to 0 (For bus masters only)
10:9	Device Select (DEVSEL#) Timing 00 - Fast DEVSEL# timing 01 - Medium DEVSEL# timing 10 - Slow DEVSEL# timing 11 - Reserved	Hardwired to "01" (CRAM can claim a PCI transaction 2 clock cycles after FRAME# was asserted)
11	Signaled Target Abort 1 - Target terminated cycle with Target Abort 0 - Didn't	Hardwired to 0 (CRAM never uses a Target Abort; only a Disconnect and a Retry are used to terminate transactions)
12	Received Target Abort 1 - Transaction was terminated by a Target Abort 0 - Didn't	Hardwired to 0 (Only indicated by a bus master whose transaction was terminated)
13	Received Master Abort 1 - Transaction was terminated by a Master Abort 0 - Didn't	Hardwired to 0 (Again, for bus masters only)
14	Signaled System Error (SERR#) 1 - Device signaled an address parity error 0 - Did not	Register (RESET# to 0; set to 1 if CRAM detects and issues a system error)
15	Detected Parity Error 1 - Device detected a data parity error 0 - Didn't	Register (RESET# to 0; set to 1 if CRAM detects a data parity error)

Table A.9 CRAM PCI Status Register

- **Class Code register:** The class code is a 24-bit read-only register that identifies the function of the PCI device. The CRAM controller is a Memory Controller (bits 23:16 hardwired to 0x05), and it is not a RAM or Flash Memory Controller (bits 15:8 hardwired to 0x80). Like all PCI Memory Controllers, the CRAM controller has no specific programming interface (bits 7:0 hardwired to 0x00).
- **Base Address Register 0:** Base address registers allow units on a PCI device to be automatically allocated a memory or I/O address space. Memory is prefetchable if reads have no side effects on the con-

tents of the locations being read, a read always returns all four bytes regardless of the byte enable settings, posted and merged writes do not cause errors, and the memory is not cached by the host processor. Table A.10 describes CRAM implementation of the base address register 0.

Bit	Function	CRAM Implementation
0	Memory or I/O Space 0 - Memory; 1- I/O	Hardwired to 0 (All CRAM units are memory-mapped)
2:1	Memory Location 00 - Anywhere in 32-bit address space 01 - Below 1 MB 10 - Anywhere in 64-bit address space 11 - Reserved	Hardwired to "00" (Current CRAM controller is on a 32-bit PCI bus)
3	Prefetchable Bit 0 - Non-prefetchable 1 - Prefetchable	Hardwired to 1 (CRAM memory is prefetchable)
31:4	Base Address	23:4 is hardwired to 0x00000 (CRAM address space is 16 MBytes); 31:24 are registers (RESET# to 0; initialized by the host processor to assign the base address of the CRAM card

Table A.10 CRAM PCI Base Address Register 0

A.3.1 PCI Commands

Table A.11 shows all PCI commands and the ones supported by the CRAM controller interface unit.

CBE#[3:0]	PCI Command	Supported in CRAM
0000	Interrupt Acknowledge	No (Only for Host/PCI bridges)
0001	Special Cycle	No (Broadcast messages are ignored by the CRAM controller)
0010 0011	I/O read I/O Write	No (All resources are memory-mapped) No
0100-0101	Reserved	No
0110 0111	Memory Read Memory Write	Yes Yes
1000-1001	Reserved	No
1010 1011	Configuration Read Configuration Write	Yes Yes
1100	Memory Read Multiple	Yes
1101	Dual-Access Cycle	No (No 64-bit addressing supported)
1110 1111	Memory Read Line Memory Write and Invalidate	Yes Yes

Table A.11 PCI Command Types

Appendix B

CRAM System PCBs and Pinouts

B.1 ISA CRAM System

B.1.1 PCB Layout

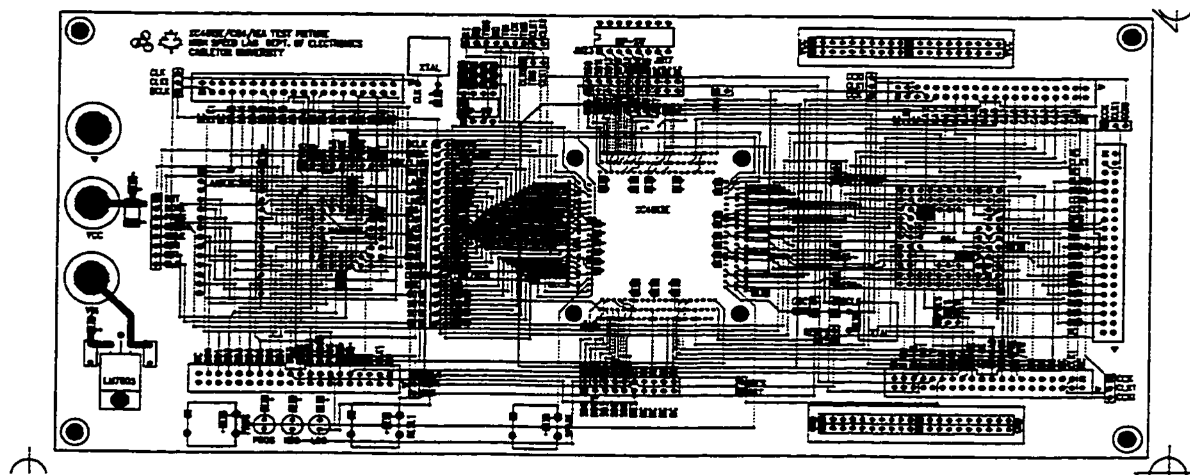

Figure B.1 ISA CRAM System PCB Layout

B.1.2 Pinout

CRAM Controller FPGA (Xilinx XC4013E-PQ240)

PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL
P1	GND	P27	BALE	P59	GND	P103	MCK	P151	GND	P205	IOB1
P2	BCLK	P28	SA<4>	P60	M0	P105	SLORI	P152	D<4>	P207	XCTRL
P4	MWTC	P29	GND	P61	VCC	P106	GND	P154	D<5>	P209	RSTB
P5	MRDC	P30	VCC	P62	M2	P108	SROLI	P156	D<6>	P211	GND
P6	LA<17>	P31	SA<5>	P63	SD<0>	P110	CCK	P159	D<7>	P212	VCC
P7	LA<18>	P32	SA<6>	P64	SD<1>	P117	A<10>	P161	VCC	P213	BTOUT
P8	LA<19>	P33	SA<7>	P65	NOWS	P119	GND	P162	A<0>	P222	VCC
P9	IRQ15	P34	SA<8>	P66	SD<2>	P120	DONE	P164	A<1>	P223	JMPR<17>
P10	LA<20>	P35	SA<9>	P67	SD<3>	P121	VCC	P166	GND	P224	JMPR<18>
P11	LA<21>	P36	SA<10>	P68	SD<4>	P122	PROGRAM	P167	A<2>	P225	JMPR<19>
P12	IRQ11	P38	SA<11>	P69	SD<5>	P123	CCKI	P169	A<3>	P226	JMPR<20>
P13	LA<22>	P39	SA<12>	P70	IRQ9	P126	A<11>	P171	A<4>	P227	GND
P14	GND	P40	VCC	P71	SD<6>	P128	A<12>	P173	A<5>	P228	JMPR<21>
P15	IRQ10	P41	SA<13>	P72	SD<7>	P130	A<13>	P175	A<6>	P229	JMPR<22>
P16	LA<23>	P42	SA<14>	P73	RESET	P132	A<14>	P176	A<7>	P230	JMPR<23>
P18	SBHE	P43	SA<15>	P75	GND	P134	A<15>	P177	DIN	P232	SD<15>
P19	VCC	P44	SA<16>	P80	VCC	P135	GND	P179	CCLK	P233	SD<14>
P20	M16	P45	GND	P89	INIT	P140	VCC	P180	VCC	P234	SD<13>
P21	OSC	P46	SA<17>	P90	VCC	P141	D<0>	P182	GND	P235	SD<12>
P23	SA<0>	P47	SA<18>	P91	GND	P144	D<1>	P192	A<8>	P236	SD<11>
P24	SA<1>	P48	SA<19>	P97	RAM	P146	D<2>	P194	A<9>	P237	SD<10>
P25	SA<2>	P51	CHRDY	P100	OPS	P148	D<3>	P196	GND	P238	SD<9>
P26	SA<3>	P58	M1	P101	VCC	P150	VCC	P201	VCC	P239	SD<8>
								P202	BTEN	P240	VCC

Download CPLD (MACH 210A-10JC)

PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL
P1	GND	P12	GND	P23	GND	P34	GND
P2	INIT	P13	CLK	P24	A<0>	P35	N/C
P3	DONE	P14	D<0>	P25	A<1>	P36	A<8>
P4	PROGRAM	P15	D<1>	P26	A<2>	P37	A<9>
P5	DIN	P16	D<2>	P27	A<3>	P38	A<10>
P6	CCLK	P17	D<3>	P28	A<4>	P39	A<11>
P7	FSM<0>	P18	D<4>	P29	A<5>	P40	A<12>
P8	FSM<1>	P19	D<5>	P30	A<6>	P41	A<13>
P9	N/C	P20	D<6>	P31	A<7>	P42	A<14>
P10	N/C	P21	D<7>	P32	N/C	P43	A<15>
P11	N/C	P22	VCC	P33	N/C	P44	VCC

Download EPROM (AM27C512)

PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL
P1	A<15>	P8	A<2>	P15	D<3>	P22	OE (GND)
P2	A<12>	P9	A<1>	P16	D<4>	P23	A<11>
P3	A<7>	P10	A<0>	P17	D<5>	P24	A<9>
P4	A<6>	P11	D<0>	P18	D<6>	P25	A<8>
P5	A<5>	P12	D<1>	P19	D<7>	P26	A<13>
P6	A<4>	P13	D<2>	P20	CE (GND)	P27	A<14>
P7	A<3>	P14	GND	P21	A<10>	P28	VCC

Appendix C

CRAM Software Details

C.1 CRAM Assembly Language

C.1.1 CRAM Memory Variable Instructions

These are instructions that operate on CRAM memory variables (cvar). They are divided into the following groups:

- **3-Address Arithmetic and Logical Instructions:** In these instructions, AR0 is address of the destination operand, and AR1 and AR2 are the addresses of the two source operands. These instructions include ADD, SUB, XOR, OR, and AND.

  Syntax: *mnemonic* #AR0, #AR1, AR2;
  Example: ADD #24, #16, #8; // mem[24] ⇐ mem[16] + mem[8]

- **2-Address, 1 Immediate Value, Arithmetic and Logical Instructions:** AR0 is the address of destination operand, AR1 is the address of source operand, and *Imm* is an immediate integer value. Example are ADDI, SUBI, ISUB, XORI, ORI, and ANDI. ISUB is subtract from immediate value.

  Syntax: *mnemonic* #AR0, #AR1, #*Imm*;
  mnemonic #AR0, #*Imm*, #AR1; // syntax for ISUB
  Example: ADDI #24, #16, #10; // mem[24] ⇐ mem[16] + 10

- **2-Address Arithmetic and Logical Instructions:** AR0 and AR1 are addresses of destination and source operands, respectively. Examples are INC, DEC, NEG, and NOT.

  Syntax: *mnemonic* #AR0, #AR1;
  Example: INC #24, #16; // mem[24] ⇐ mem[16] + 1

- **Compare Instructions:** These instructions compare two operands at addresses AR1 and AR2. The result of the comparison is put in a cbool or cboolref variable at address AR0. Examples are LT, GT, EQ, LTE, GTE, and NEQ.

  Syntax: *mnemonic* #AR0, #AR1, #AR2;
  Example: LT #24, #16, #8; // mem[24] ⇐ (mem[16] < mem[8])

- **Compare Immediate Instructions:** These instructions compare an operand at address AR1 with an immediate integer, *Imm*. The result is put in a cbool or cboolref variable at address AR0. Examples are LTI, GTI, EQI, LTEI, GTEI, and NEQI.

  Syntax: *mnemonic* #AR0, #AR1, #*Imm*;
  Examples: EQI #8, #12, #16; // mem[8] ⇐ (mem[12] == 16)

- **Minimum/Maximum Search Instructions:** These instructions search for the minimum or maximum value of a cvar variable. The result is a cbool or cboolref (at mem[AR0]) that is set for a PE that contains a minimum or maximum element. The search instructions are MIN and MAX. In both cases, AR1 is the address of the $(n-1)^{th}$ bit of the search source operand (n is the number of bits of the operand).

  Syntax: *mnemonic* #AR0, #AR1;

Example: MAX #17, #16; // mem[17] ⇐ (is maximum element of mem[16])

- **Shift Instructions:** These instructions shift or rotate the positions of the elements of a cvar variable at address AR1, and puts the results at address AR0. If specified, the elements are shifted n positions, otherwise they are shifted one position. Shift instructions are SL, SR, RTL, and RTR. These instructions should not be confused with C++ shift operators (<< and >>) which shift the bits of the individual elements themselves.

  Syntax: *mnemonic* #AR0, #AR1 [, #*n*];
  Example: SL #16, #8, #5; // mem[16] ⇐ (mem[8] shifted left 5 columns)

- **Load Immediate Instruction:** The load immediate instruction, MVI, loads a constant integer, *Imm*, into all the elements of a cvar variable at address AR0.

  Syntax: MVI #AR0, #*Imm*;
  Example: MVI #16, #25; // (All elements of mem[16]) ⇐ 25

- **Copy Instruction:** The cvar copy instruction, MOV, copies the elements of a cvar at address AR1 into a cvar at address AR0. MOVR is a variation of MOV that does the copying in reverse starting with the last bit of the operands. Therefore for MOVR, AR0 and AR1 should point to the last bit of the operands.

  Syntax: MOV #AR0, #AR1; // mem[AR0] ⇐ mem[AR1]

- **Set/Clear Instructions:** MSET and MCLR sets (sets value to 0xFF...) and clears (sets value to 0x000...), respectively, the elements of a cvar at address AR0.

  Syntax: *mnemonic* #AR0;
  Example: MCLR #18; // mem[18] ⇐ 0

- **Copy PE Register Instructions:** The MOVPE instruction copies the contents of a PE register into a cbool or cboolref variable at address AR0. A variant of MOVPE instruction, the NMOVPE instruction, loads all the bits of a cvar variable with the contents of a PE register. Optionally, the source operand can be the inverse of a PE register by preceding the register name with an exclamation mark (!). Valid PE source registers are X, Y, and M.

  Syntax: *mnemonic* #AR0, [!]*PE_reg*;
  Example: MOVPE #16, !X; // mem[16] ⇐ $\overline{\text{PE.X}}$

- **Boolean Instructions:** All instructions involving cvar CRAM variables require that the correct value of the bit length of the operands be present in the WLEN register. This might require loading this value using the LDWLEN instruction before issuing a cvar instruction. But since instructions involving CRAM boolean variables (cbool and cboolref) are very common but always require a bit-length of one, a few of the cvar instructions are provided with their boolean equivalents. For these instructions, the bit-length of 1 has already been incorporated in their microinstructions. Therefore there is no need of setting up WLEN before issuing boolean instructions. The names of these instructions are made from their parent instructions, with a B added at the end. Boolean instructions include ANDB, ORB, NOTB, and MOVB. Their syntax is the same as their parent instructions.

C.1.2 CRAM PE Register Instructions

These are instructions that operate on PE registers. They include instructions to set or clear registers, and instructions to load PE registers from either other PE registers, or from cbool or cboolref variables. The following sections describe the PE register instructions.

- **Set/Clear Register:** CLR and ST instructions set the PE register to 0 and 1 respectively. Valid destination registers (*PE_reg*) are X, Y, and W.

Syntax: *mnemonic PE_reg* [*PE_reg*] [*PE_reg*];
Example: CLR X Y; // X[PE] ⇐ 0, Y[PE] ⇐ 0

- **Load from cvar:** The LDPE instruction loads PE registers with the contents of a cbool or cboolref variable at address AR0. Again, the destination register can be X, Y, or W. Optionally, you can load the registers with the inverse of the source operand by preceding the address register with an exclamation mark (!).

  Syntax: LDPE *PE_reg* [*PE_reg*] [*PE_reg*], [!]#AR0;
  Example: LDPE X, !#16; // X[PE] ⇐ $\overline{\text{mem[16]}}$

- **Register-to-Register Copy:** The TXPE instruction copies the contents of a PE register to another. The destination registers (*PE_dst*) can be X, Y, or W. The source register (*PE_src*) can be either X, Y, M, or their inverse (register name preceded by !).

  Syntax: TXPE *PE_dst* [*PE_dst*] [*PE_dst*], [!]*PE_src*;
  Example: TXPE X, !Y; // X[PE] ⇐ $\overline{\text{Y[PE]}}$

C.1.3 CRAM Controller Instructions

CRAM controller instructions include instructions to load, increment and decrement controller registers, as well as instructions for reading and writing data between CRAM and the controller. The following sections explain.

- **Load Register:** These are instructions to load specific registers of the controller (Table A.4 and Table A.5). *Imm* is the integer value to be loaded into the register. Load register instructions include LDCBA, LDIBA, LDWLEN, LDAX0, LDAX1, LDAX2, LDSIC, and SETINT.

  Syntax: *mnemonic #Imm*;
  Example: LDAX0 #5; // AX0 ⇐ 5

- **Increment/Decrement Register:** Registers AR0, AR1, and AR2 can be incremented and decremented by INCA and DECA instructions, respectively. Register CBA can also be incremented by INCA, but it can not be decremented.

  Syntax: *mnemonic ctrl_reg*;
  Example: INCA AR1; // AR1 ⇐ AR1 + 1

- **Memory Read/Write:** CRAM memory bytes starting at address [AR0] can be read and put into the controller read buffer at address [RIBA] using the READ instruction. Similarly, bytes can be transferred from the controller write buffer to the CRAM memory using the WRITE instruction. In both cases, like in any other instructicn, the number of bytes to be transferred is defined by the contents of the WLEN register, which normally represents the number of bits for the cvar variable being accessed.

  Syntax: READ AR0; // read_buffer[RIBA] ⇐ mem[AR0]
  WRITE AR0; // mem[AR0] ⇐ write_buffer[WIBA]

- **Conditional Read-Bank:** The conditional read-bank instruction (RDBNK) is used to scan a cbool or cboolref variable to check if there is a 1 at any PE position. This is done by first reading the byte of the variable corresponding to the first eight PEs. If this byte contains a 1 on any of its eight bits, the byte is saved in the DTR register and the RDBNK instruction terminates. Otherwise the bank address register (CBA) is incremented to point to the byte corresponding to the next eight PEs, and the cycle is repeated. This is repeated until either a 1 is found or all the bits of the variable have been scanned. The results of the scan (CBA and DTR) can be used by the host processor to compute the index of a PE that yielded a true value to a search or comparison.

  Syntax: RDBNK AR0; // DTR ⇐ mem[CBA:AR0], if mem[CBA:AR0] != 0

C.2 CRAM Microcode Assembly Language

C.2.1 Microroutine Format

Each microroutine begins with its name followed by a colon. This is the name that is used in the CRAM assembly language, e.g. ADD, LDPE. The name declaration is followed by one or more microinstructions. The format of the microroutine is illustrated below. The different types of microinstructions are described in subsequent sections.

Example:
```
MOVPE: // instruction name; moves PE register into cbool
RST, EXT_TTOP;
CWRITE AR0, END;
```

C.2.2 PE Microinstructions

These instructions act on PE registers. They can be optionally mixed with some controller instructions.

Syntax: *mnemonic* [*PE_dst* [*PE_dst* ...]] [, [!]*PE_src* [[!]*PE_src* ...]]
[, EXT_TTOP] [, EXT_COP] [, *nai*] [, IF *cond*];

where,

PE_dst - is either Y, X, W, LX, RY, or B. LX is the X register of the left-neighbor PE, and hence may not be used simultaneously with X as the destination register. Similarly, RY is the Y register of the right-neighbor PE. RY and Y can not be used simultaneously. B denotes a bus-tie (Global-OR) operation on the results of a PE operation.

PE_src - is either X, Y, M. The inverse of the registers can be used by preceding the register name with the optional exclamation mark (!).

EXT_TTOP - shows that the PE truth table opcode (TTOP) for this microinstruction comes as OPR2 of the (macro)instruction. This TTOP is specified by the source operands of the instruction USE clause (Section C.2.5).

EXT_COP - shows that the PE control opcode (COP) for this microinstruction comes as OPR1 of the (macro)instruction. The actual value of COP is derived from the destination operands of the instruction USE clause (Section C.2.5).

nai - is any of the next address instructions (NAI) shown in Table A.1. Usually, the NEXTu instruction is not explicitly specified, but is assumed in all cases where no *nai* is specified. The keyword END can be (and is usually) used as an alias for NEXTI.

cond - is the condition if *nai* is a conditional instruction such as LOOPE, LOOPN, JMP, or JSR. Valid conditions are listed in Table A.2.

Example: ACARRY X, X !Y M, LOOPE, IF CNZ; // X $\Leftarrow$ carry(X, !Y, M); loop if CNZ $\neq$ 0

C.2.3 Memory Access Microinstructions

The five CRAM memory access microinstructions are CREAD, CWRITE, READ, WRITE, and RDBNK. CREAD reads data from a CRAM memory row to the PEs M registers. CWRITE writes data from the PEs ALU outputs to a CRAM memory row. READ and WRITE transfers data between CRAM (pointed to by CBA:AX0:AR0) and the controller read/write buffers. RDBNK is used for conditionally reading a byte from CRAM as described in Appendix C.1.3.

Syntax: *mnemonic addr_reg* [, *cci*]] [, *nai*] [, IF *cond*];

where,

addr_reg - is either AR0, AR1, or AR2, and contains address of the CRAM memory row.

cci - is the controller instruction to increment or decrement the *addr_reg* or CBA after the microinstruction has finished executing. Valid instructions are INCA, DECA, and INCB.

nai and *cond* - are the next address instruction and its condition (see Section C.2.2).

Example: CREAD AR0, SETLP; // PE.M $\Leftarrow$ mem[AR0]; LPR $\Leftarrow$ μPC + 1

C.2.4 Controller Microinstructions

These are microinstructions for loading controller registers (Appendix C.1.3).

Syntax: *mnemonic* [, EXT_COP]] [, *nai*] [, IF *cond*];

where,

EXT_COP - denotes that the actual register to be loaded will be identified in the USE clause. Note that COP is used here since the actual controller load-register instruction is identified by microinstruction bits [15:8], which is also the position of PE COP.

nai and *cond* - are the next address instruction and its condition (see Section C.2.2).

Example: LDIBA, EXT_COP; // IBA ⇐ OPR0;

C.2.5 USE Clause

When an instruction microroutine has EXT_TTOP and/or EXT_COP in one of its microinstructions, another instruction can use this microroutine by simply specifying its specific PE source registers (TTOP) and/or PE destination registers or auxiliary controller instruction (COP). The USE clause is used to specify the full microinstruction that should substitute the microinstruction with EXT_TTOP/EXT_COP

Syntax: *instr_name*:

USE *prev_instr*: *new_μinstr*

where,

instr_name - name of instruction whose microroutine is being defined.

prev_instr - name of previously defined instruction whose microroutine should be used.

new_μinstr - the microinstruction that should be used to substitute the microinstruction in the *prev_instr* microroutine. Note that even though the full microinstruction is used, only its components that are relevant in forming TTOP and/or COP are used. The rest (NAI, COND, etc.) use the parameters specified in the original microinstruction.

Example: MSET:// instruction name; used to set the memory to FF..

USE MOVEPE: SET;// by using the MOVPE routine with TTOP of SET

Appendix D

Applications Source Code

D.1 Low-Level Image Processing

```
/*********************************************************/
/* Brightness Adjustment (Uniprocessor)                 */
/*  Brightens the image by adding a value to each pixel value   */
/*********************************************************/

void brighten (IMAGE imageIn[IMAGE_SIZE][IMAGE_SIZE],
          IMAGE imageOut[IMAGE_SIZE][IMAGE_SIZE],
          unsigned char adjustment)
{
 int i, j;
 unsigned int new_val;

 for (i=0; i<IMAGE_SIZE; i++)
 {
  for (j=0; j<IMAGE_SIZE; j++)
  {
   new_val = imageIn[i][j] + adjustment;
   imageOut[i][j] = (new_val > 255) ? 255 : new_val;
  }
 }
}
```

```
/*********************************************************/
/* Brightness Adjustment (CRAM)                          */
/*  Brightens the image by adding a value to each pixel value   */
/*********************************************************/

void brighten (CRAM_IMAGE& imageIn, CRAM_IMAGE& imageOut,
          unsigned char adjustment)
{
 imageOut = imageIn + adjustment;
 cif (imageOut.bit(8))
  imageOut = 255;
 cend;
}
```

```
/*********************************************************/
/* Pixel Averaging (Uniprocessor)                        */
/*  Noise reduction by averaging pixel values in 3x3 neighborhood   */
/*********************************************************/

void average (IMAGE imageIn[IMAGE_SIZE][IMAGE_SIZE],
          IMAGE imageOut[IMAGE_SIZE][IMAGE_SIZE])
{
 int i, j;
 unsigned int ave;

 // leave the edges unchanged
 for (i=0; i<IMAGE_SIZE; i++)
 {
  imageOut[i][0] = imageIn[i][0];
  imageOut[i][IMAGE_SIZE-1] = imageIn[i][IMAGE_SIZE-1];
 }
 for (j=0; j<IMAGE_SIZE; j++)
 {
  imageOut[0][j] = imageIn[0][j];
  imageOut[IMAGE_SIZE-1][j] = imageIn[IMAGE_SIZE-1][j];
 }

 // average the other pixels (avoid edges!)
 for (i=1; i<(IMAGE_SIZE-1); i++)
 {
  for (j=1; j<(IMAGE_SIZE-1); j++)
  {
   ave = (imageIn[i-1][j-1] + imageIn[i-1][j] + imageIn[i-1][j+1] +
        imageIn[i][j-1] + imageIn[i][j] + imageIn[i][j+1] +
        imageIn[i+1][j-1] + imageIn[i+1][j] + imageIn[i+1][j+1])/9;
   imageOut[i][j] = ave;
  }
 }
}
```

```
/*********************************************************/
/* Pixel Averaging (CRAM)                                */
/*  Noise reduction by averaging pixel values in 3x3 neighborhood   */
/*********************************************************/

void average (CRAM_IMAGE& imageIn, CRAM_IMAGE& imageOut,
          cbool& not_edge_pes)
{
  // image shifted 1, IMAGE_SIZE-1, IMAGE_SIZE, IMAGE_SIZE+1
times
  cuchar shifted_image[4];

  cuint sum_left(10), sum_right(10), sum(11);

  // bring pixels from left and top, and add them
  PE.shiftr (shifted_image[0], imageIn, 1, 0);
  PE.shiftr (shifted_image[1], shifted_image[0], IMAGE_SIZE-2, 0);
  PE.shiftr (shifted_image[2], shifted_image[1], 1, 0);
  PE.shiftr (shifted_image[3], shifted_image[2], 1, 0);
  sum_left = (shifted_image[0] + shifted_image[1]) +
        (shifted_image[2] + shifted_image[3]);

  // bring pixels from right and bottom, and add them
  PE.shiftl (shifted_image[0], imageIn, 1, 0);
  PE.shiftl (shifted_image[1], shifted_image[0], IMAGE_SIZE-2, 0);
  PE.shiftl (shifted_image[2], shifted_image[1], 1, 0);
  PE.shiftl (shifted_image[3], shifted_image[2], 1, 0);
  sum_right = (shifted_image[0] + shifted_image[1]) +
        (shifted_image[2] + shifted_image[3]);

  // average
  sum = sum_left + sum_right;
  sum += imageIn;
  imageOut = sum/9;

  cif (!not_edge_pes)
   imageOut = imageIn;
  cend
}
```

```
/************************************************************/
/* Edge Enhancement  (Uniprocessor)                         */
/*  Enhances edges by subtracting the Laplacian of a pixel from the pixel */
/************************************************************/

void enhance_edges (IMAGE imageIn[IMAGE_SIZE][IMAGE_SIZE],
        IMAGE imageOut[IMAGE_SIZE][IMAGE_SIZE])
{
 int i, j;
 int diff;

 // leave the edges unchanged
 for (i=0; i<IMAGE_SIZE; i++)
 {
  imageOut[i][0] = imageIn[i][0];
  imageOut[i][IMAGE_SIZE-1] = imageIn[i][IMAGE_SIZE-1];
 }
 for (j=0; j<IMAGE_SIZE; j++)
 {
  imageOut[0][j] = imageIn[0][j];
  imageOut[IMAGE_SIZE-1][j] = imageIn[IMAGE_SIZE-1][j];
 }

 // subtract Laplacian from pixels (avoid edges!)
 for (i=1; i<(IMAGE_SIZE-1); i++)
 {
  for (j=1; j<(IMAGE_SIZE-1); j++)
  {
   diff = (imageIn[i][j] + 4*imageIn[i][j] - imageIn[i-1][j] -
          imageIn[i][j-1] - imageIn[i][j+1] - imageIn[i+1][j]);
   diff = abs(diff);
   imageOut[i][j] = (diff > 255) ? 255: diff;
  }
 }
}
```

```
/************************************************************/
/* Edge Enhancement  (CRAM)                                 */
/*  Enhances edges by subtracting Laplacian of a pixel from the pixel */
/************************************************************/

void enhance_edges (CRAM_IMAGE& imageIn,
      CRAM_IMAGE& imageOut, cbool& not_edge_pes)
{
 // image shifted 1, IMAGE_SIZE times
 cuchar shifted_image[2];
 cuint sum_left(9), sum_right(9), temp(12);
 cint laplacian (11);

 // bring pixels from left and top, and add them
 PE.shiftr (shifted_image[0], imageIn, 1, 0);
 PE.shiftr (shifted_image[1], shifted_image[0], IMAGE_SIZE-1, 0);
 sum_left = shifted_image[0] + shifted_image[1];

 // bring pixels from right and bottom, and add them
 PE.shiftl (shifted_image[0], imageIn, 1, 0);
 PE.shiftl (shifted_image[1], shifted_image[0], IMAGE_SIZE-1, 0);
 sum_right = shifted_image[0] + shifted_image[1];

 // find laplacian
 laplacian = 4*imageIn - (sum_right + sum_left);
 laplacian += imageIn;
 temp = abs(laplacian);

 cif(temp > 255)
  imageOut = 255;
 celse
  imageOut = temp;
 cend

  // leave edges unchanged
 cif (!not_edge_pes)
  imageOut = imageIn;
 cend
}
```

```
/************************************************************/
/* Segmentation (Thresholding)  (Uniprocessor)              */
/*  Segments the image by converting it to binary image     */
/************************************************************/

void tobinary (IMAGE imageIn[IMAGE_SIZE][IMAGE_SIZE],
        IMAGE imageOut[IMAGE_SIZE][IMAGE_SIZE],
        unsigned char threshold)
{
 int i, j;

 for (i=0; i<IMAGE_SIZE; i++)
 {
  for (j=0; j<IMAGE_SIZE; j++)
   imageOut[i][j] = (imageIn[i][j] >= threshold) ? 1 : 0;
 }
}
```

```
/************************************************************/
/* Segmentation (Thresholding)  (CRAM)                      */
/*  Segments the image by converting it to binary image     */
/************************************************************/

void tobinary (CRAM_IMAGE& imageIn, BINARY_IMAGE& imageOut,
        unsigned char threshold)
{
 if (threshold == 0)
  imageOut = (imageIn >= 0);
 else
  imageOut = (imageIn > (threshold-1));
}
```

```
/************************************************************/
/* Segmentation (Multiple-Thresholding)  (Uniprocessor)     */
/*  Segments the image by using multiple thresholding       */
/************************************************************/

void multi_threshold (IMAGE imageIn[IMAGE_SIZE][IMAGE_SIZE],
        IMAGE imageOut[IMAGE_SIZE][IMAGE_SIZE],
        unsigned char threshold[OBJECTS])
{
 int i, j, objm1;
 int obj; // objm1+1 - to reduce no of additions inside loop

 // for background
 for (i=0; i<IMAGE_SIZE; i++)
 {
  for (j=0; j<IMAGE_SIZE; j++)
   if (imageIn[i][j] < threshold[0]) imageOut[i][j] = 0;
 }

 // for objects 1 to N-1
```

```
/************************************************************/
/* Segmentation (Multiple-Thresholding)  (CRAM)             */
/*  Segments the image by using multiple thresholding       */
/************************************************************/

void multi_threshold (CRAM_IMAGE& imageIn, cuint& imageOut,
        unsigned char threshold[OBJECTS])
{
 int objm1, obj; // to reduce no of additions inside loop

 // for background
 cif (imageIn < threshold[0])
  imageOut = 0;
 cend

 // for objects 1 to N-1
 for (objm1=0, obj=1; objm1<(OBJECTS-1); objm1++,obj++)
 {
  cif ((imageIn >= threshold[objm1]) && (imageIn < threshold[obj]))
   imageOut = obj;
```

```
for (objm1=0, obj=1; objm1<(OBJECTS-1); objm1++,obj++)
{
 for (i=0; i<IMAGE_SIZE; i++)
 {
  for (j=0; j<IMAGE_SIZE; j++)
  {
    if ((imageIn[i][j] >= threshold[objm1]) &&
      (imageIn[i][j] < threshold[obj]))
     imageOut[i][j] = obj;
  }
 }
}

// for object N
for (i=0; i<IMAGE_SIZE; i++)
{
 for (j=0; j<IMAGE_SIZE; j++)
       if (imageIn[i][j] >= threshold[OBJECTS-1]) imageOut[i][j] =
OBJECTS;
 }
}
```

```
  cend
 }

 // for object N
 cif (imageIn >= threshold[OBJECTS-1])
  imageOut = OBJECTS;
 cend
}
```

D.2 Database Applications

```
/************************************************************/
/* EQUAL-TO SEARCH  (Uniprocessor)                          */
/*  An example of database equivalence searches: records with values  */
/*  equal to the search key are searched & replaced with a new value   */
/************************************************************/

void equal_to_search (RECORD_TYPE records[RECORDS],
       RECORD_TYPE key, RECORD_TYPE new_value)
{
 for (int rec=0; rec<RECORDS; rec++)
  if (records[rec] == key) records[rec] = new_value;
}
```

```
/************************************************************/
/* EQUAL-TO SEARCH   (CRAM)                                 */
/*  An example of database equivalence searches: records with values  */
/*  equal to the search key are searched & replaced with a new value   */
/************************************************************/

void equal_to_search (RECORD_TYPE& records, KEY_TYPE key,
       KEY_TYPE new_value)
{
 cif (records == key)
  records = new_value;
 cend
}
```

```
/************************************************************/
/* GREATER-THAN SEARCH  (Uniprocessor)                      */
/*  An example of database threshold searches: records with values   */
/*  greater than the threshold value are searched and replaced       */
/************************************************************/

void greater_than_search (RECORD_TYPE records[RECORDS],
       RECORD_TYPE threshold_value, RECORD_TYPE new_value)
{
 for (int rec=0; rec<RECORDS; rec++)
  if (records[rec] > threshold_value) records[rec] = new_value;
}
```

```
/************************************************************/
/* GREATER-THAN SEARCH   (CRAM)                             */
/*  An example of database threshold searches: records with values   */
/*  greater than the threshold value are searched and replaced       */
/************************************************************/

void greater_than_search (RECORD_TYPE& records,
       KEY_TYPE threshold_value,
       KEY_TYPE new_value)
{
 cif (records > threshold_value)
  records = new_value;
 cend
}
```

```
/************************************************************/
/* BETWEEN-LIMITS SEARCH  (Uniprocessor)                    */
/*  records with values between the two limits are searched & replaced  */
/************************************************************/

void between_limits_search (RECORD_TYPE records[RECORDS],
       RECORD_TYPE low_limit, RECORD_TYPE high_limit,
       RECORD_TYPE new_value)
{
 for (int rec=0; rec<RECORDS; rec++)
  if ((records[rec] > low_limit) && (records[rec] < high_limit))
   records[rec] = new_value;
}
```

```
/************************************************************/
/* BETWEEN-LIMITS SEARCH  (CRAM)                            */
/*  records with values between the two limits are searched & replaced  */
/************************************************************/

void between_limits_search (RECORD_TYPE& records,
       KEY_TYPE low_limit,
       KEY_TYPE high_limit, KEY_TYPE new_value)
{
 cif ((records > low_limit) && (records < high_limit))
  records = new_value;
 cend
}
```

```
/**********************************************************/
/* MAXIMUM SEARCH  (Uniprocessor)                          */
/*  An example of database extreme searches; records with the maximum*/
/*  values are searched and replaced with a new value;     */
/**********************************************************/

void maximum_search (RECORD_TYPE records[RECORDS],
        RECORD_TYPE new_value)
{
 int rec;
 int prev_match[RECORDS];
 int tail = 0;
 RECORD_TYPE max_record = records[0];
 prev_match[0] = -1;

 for (rec=1; rec<RECORDS; rec++)
 {
  if (records[rec] > max_record)
  {
   max_record = records[rec];
   prev_match[rec] = -1;
   tail = rec;
  }
  else if (records[rec] == max_record)
  {
   prev_match[rec] = tail;
   tail = rec;
  }
 }

 // replace records with new value
 rec = tail;
 do
 {
  records[rec] = new_value;
  rec = prev_match[rec];
 } while (rec != -1);
}
```

```
/**********************************************************/
/* MAXIMUM SEARCH (CRAM)                                   */
/*  An example of database extreme searches; records with the maximum*/
/*  values are searched and replaced with a new value;     */
/**********************************************************/

void maximum_search (RECORD_TYPE& records,
       KEY_TYPE new_value)
{
 cif (ismax(records))
  records = new_value;
 cend
}
```

```
/**********************************************************/
/* LMS (Uniprocessor)                                      */
/*  multi-criteria records that best match the search keye are searched */
/*  using the least mean squared match and matching records are  */
/*  replaced with new values                               */
/**********************************************************/

void lms (RECORD_TYPE records[RECORDS][VEC_SIZE],
        RECORD_TYPE key[VEC_SIZE],
        RECORD_TYPE new_data[VEC_SIZE])
{
 int rec, vec;
 int field_error; // error between records and search key fields
 int rec_error; // accumulated squared errors for the vector

 int prev_match[RECORDS];
 int tail = 0;
 int min_error;
 prev_match[0] = -1;

 // case record 0
 rec_error = 0;
 for (vec=0; vec<VEC_SIZE; vec++)
 {
  field_error = records[0][vec] - key[vec];
  rec_error += (field_error*field_error);
 }
 min_error = rec_error;

 for (rec=1; rec<RECORDS; rec++)
 {
  rec_error = 0;
  for (vec=0; vec<VEC_SIZE; vec++)
  {
   field_error = records[rec][vec] - key[vec];
   rec_error += (field_error*field_error);
  }
```

```
/**********************************************************/
/* LMS  (CRAM)                                             */
/*  multi-criteria records that best match the search keye are searched */
/*  using the least mean squared match and matching records are  */
/*  replaced with new values                               */
/**********************************************************/

void lms (RECORD_TYPE records[VEC_SIZE],
       KEY_TYPE key[VEC_SIZE],
       KEY_TYPE new_data[VEC_SIZE])
{
 // temporary CRAM variables
 cint field_error(17); // error between records and search key fields
 cuint rec_error(32);  // accumulated squared errors for the vector

 // find squared error between record and search key fields
 int vec;
 rec_error = 0;
 for (vec=0; vec<VEC_SIZE; vec++)
 {
  field_error = records[vec] - key[vec];
  rec_error += (field_error*field_error);
 }

 // update records with best LMS match
 cif (ismin(rec_error))
  for (vec=0; vec<VEC_SIZE; vec++) records[vec] = new_data[vec];
 cend
}
```

```
/*************************************************************/
/* Uniprocessor Motion Estimation                           */
/*  finds motion vectors of blocks between reference and current frames; */
/*  edged blocks are not coded (i.e. use Intra-frame information)      */
/*************************************************************/

void motion_estimation(
        unsigned char ref_frame[IMAGE_SIZE][IMAGE_SIZE],
        unsigned char cur_frame[IMAGE_SIZE][IMAGE_SIZE],
        unsigned int motion_vec[BLOCKS][2])
{
 int i, j, refi, refj;
 int blki, blkj, srchi, srchj;
 int vecx, vecy;
 int block_no = 0;
 int mae, min_mae;
 int MAX_MAE = 255*BLK_SIZE*BLK_SIZE;

 // do for each block at these positions
 for (blki=BLK_SIZE; blki<(IMAGE_SIZE-BLK_SIZE);
 blki+=BLK_SIZE)
 {
  for (blkj=BLK_SIZE; blkj<(IMAGE_SIZE-BLK_SIZE);
  blkj+=BLK_SIZE)
  {
   min_mae = MAX_MAE;

   // for each block there are (2*BLK_SIZE+1)^2 search posions
   for (srchi=(blki-BLK_SIZE); srchi<=(blki+BLK_SIZE); srchi++)
   {
     for (srchj=(blkj-BLK_SIZE); srchj<=(blkj+BLK_SIZE); srchj++)
     {
      mae = 0;

      // at each seach position, compare BLK_SIZExBLK_SIZE pixels
      for (i=blki, refi=srchi; i<(blki+BLK_SIZE); i++, refi++)
      {
       for (j=blkj, refj=srchj; j<(blkj+BLK_SIZE); j++, refj++)
       {
        mae += abs(cur_frame[i][j] - ref_frame[refi][refj]);
       }
      }

      // save if this has minimum distortion
      if (mae < min_mae)
      {
       min_mae = mae;
       vecx = srchj;
       vecy = srchi;
      }
     }
   }

   // save the motion vector
   motion_vec[block_no][0] = vecx;
   motion_vec[block_no][1] = vecy;
   block_no++;
  }
 }

}
```

```
/*************************************************************/
/* CRAM Motion Estimation                                   */
/*  finds motion vectors of blocks between reference and current frames; */
/*  edge blocks are not coded (i.e. use Intra-frame information):      */
/*  frames are divided into blocks. 1 block/PE               */
/*************************************************************/

void check_best_match (cuint& mae, cuint& min_mae,
            cuint& motion_vec, unsigned short* mvec);
void check_best_or_equal_match (cuint& mae, cuint& min_mae,
        cuint& motion_vec, unsigned short* mvec);

void motion_estimation(cuchar ref_blk[BLK_SIZE][BLK_SIZE],
            cuchar cur_blk[BLK_SIZE][BLK_SIZE],
            cuint& motion_vec)
{
 int i, j, srchi, srchj, refi, refj;
 unsigned short mvec = 0;

 cuchar temp0[BLK_SIZE][BLK_SIZE];
 cuchar temp1[BLK_SIZE][BLK_SIZE];
 cuchar temp2[BLK_SIZE][BLK_SIZE];
 cuchar temp3[BLK_SIZE][BLK_SIZE];

 cint diff(9);
 cuint mae(14), min_mae(14);

 min_mae = -1; // set to FFFF, max possible

 // shift blocks to get the 3 above and 1 to the left
 // temp3=left, temp2=top-right, temp1=top-centre, temp0=top-left
 for (i=0; i<BLK_SIZE; i++)
 {
  for (j=0; j<BLK_SIZE; j++)
  {
   PE.shiftl (temp3[i][j], ref_blk[i][j], 1, 0);
   PE.shiftl (temp2[i][j], temp3[i][j], BLOCK_COLS-2, 0);
   PE.shiftl (temp1[i][j], temp2[i][j], 1, 0);
   PE.shiftl (temp0[i][j], temp1[i][j], 1, 0);
  }
 }
/*************************************************************/
// use search positions in the top-left block
for (srchi=0; srchi<BLK_SIZE; srchi++)
{
 for (srchj=0; srchj<BLK_SIZE; srchj++)
 {
  // for each search position, compare all pixels
  mae = 0;
  for (i=0, refi=srchi; i<BLK_SIZE; i++, refi++)
  {
    for (j=0, refj=srchj; j<BLK_SIZE; j++, refj++)
    {
     if (refi < BLK_SIZE)
     {
      if (refj < BLK_SIZE) diff = cur_blk[i][j]-temp0[refi][refj];
      else diff = cur_blk[i][j]-temp1[refi][refj-BLK_SIZE];
     } else {
      if (refj < BLK_SIZE)
       diff = cur_blk[i][j]-temp3[refi-BLK_SIZE][refj];
      else diff = cur_blk[i][j]-
       ref_blk[refi-BLK_SIZE][refj-BLK_SIZE];
     }
     mae += abs(diff);
    }
  }
  check_best_match (mae, min_mae, motion_vec, &mvec);
 }
}

/*************************************************************/
// use search positions in the top-centre block
// top-left block no longer needed, transfer right block into temp0
for (i=0; i<BLK_SIZE; i++)
{
 for (j=0; j<BLK_SIZE; j++)
  PE.shiftr (temp0[i][j], ref_blk[i][j], 1, 0);
}
```

```
  for (srchi=0; srchi<BLK_SIZE; srchi++)
  {
   for (srchj=0; srchj<BLK_SIZE; srchj++)
   {
    // for each search position, compare all pixels
    mae = 0;
    for (i=0, refi=srchi; i<BLK_SIZE; i++, refi++)
    {
      for (j=0, refj=srchj; j<BLK_SIZE; j++, refj++)
      {
       if (refi < BLK_SIZE)
       {
        if (refj < BLK_SIZE) diff = cur_blk[i][j]-temp1[refi][refj];
        else diff = cur_blk[i][j]-temp2[refi][refj-BLK_SIZE];
       } else {
        if (refj < BLK_SIZE)
          diff = cur_blk[i][j]-ref_blk[refi-BLK_SIZE][refj];
        else diff = cur_blk[i][j]-temp0[refi-BLK_SIZE][refj-BLK_SIZE];
       }
       mae += abs(diff);
      }
    }
    check_best_match (mae, min_mae, motion_vec, &mvec);
   }
  }
/************************************************************/
// use search positions in the top-right block
 for (srchi=0; srchi<BLK_SIZE; srchi++)
 {
  // for each search position, compare all pixels
  mae = 0;
  for (i=0, refi=srchi; i<BLK_SIZE; i++, refi++)
  {
   for (j=0; j<BLK_SIZE; j++)
   {
     if (refi < BLK_SIZE) diff = cur_blk[i][j]-temp2[refi][j];
     else diff = cur_blk[i][j]-temp0[refi-BLK_SIZE][j];
     mae += abs(diff);
   }
  }
  check_best_match (mae, min_mae, motion_vec, &mvec);
 }
/************************************************************/
// use search positions in the left block
// top blocks no longer needed, transfer two botton blocks into temp1/2
for (i=0; i<BLK_SIZE; i++)
{
 for (j=0; j<BLK_SIZE; j++)
 {
  PE_shiftr (temp1[i][j], temp0[i][j], BLOCK_COLS-2, 0);
  PE_shiftr (temp2[i][j], temp1[i][j], 1, 0);
 }
}

for (srchi=0; srchi<BLK_SIZE; srchi++)
{
 for (srchj=0; srchj<BLK_SIZE; srchj++)
 {
  // for each search position, compare all pixels
  mae = 0;
  for (i=0, refi=srchi; i<BLK_SIZE; i++, refi++)
  {
    for (j=0, refj=srchj; j<BLK_SIZE; j++, refj++)
    {
     if (refi < BLK_SIZE)
     {
      if (refj < BLK_SIZE) diff = cur_blk[i][j]-temp3[refi][refj];
      else diff = cur_blk[i][j]-ref_blk[refi][refj-BLK_SIZE];
     } else {
      if (refj < BLK_SIZE)
        diff = cur_blk[i][j]-temp1[refi-BLK_SIZE][refj];
      else diff = cur_blk[i][j]-temp2[refi-BLK_SIZE][refj-BLK_SIZE];
     }
     mae += abs(diff);
    }
  }
  check_best_match (mae, min_mae, motion_vec, &mvec);
 }
}
/************************************************************/
// use search positions in the centre block
// left block no longer needed, transfer botton-right block into temp3
for (i=0; i<BLK_SIZE; i++)
{
 for (j=0; j<BLK_SIZE; j++)
  PE_shiftr (temp3[i][j], temp2[i][j], 1, 0);
}

for (srchi=0; srchi<BLK_SIZE; srchi++)
{
 for (srchj=0; srchj<BLK_SIZE; srchj++)
 {
  // for each search position, compare all pixels
  mae = 0;
  for (i=0, refi=srchi; i<BLK_SIZE; i++, refi++)
  {
    for (j=0, refj=srchj; j<BLK_SIZE; j++, refj++)
    {
     if (refi < BLK_SIZE)
     {
      if (refj < BLK_SIZE) diff = cur_blk[i][j]-ref_blk[refi][refj];
      else diff = cur_blk[i][j]-temp0[refi][refj-BLK_SIZE];
     } else {
      if (refj < BLK_SIZE)
        diff = cur_blk[i][j]-temp2[refi-BLK_SIZE][refj];
      else diff = cur_blk[i][j]-temp3[refi-BLK_SIZE][refj-BLK_SIZE];
     }
     mae += abs(diff);
    }
  }
  if ((srchi == 0) && (srchj == 0))
    check_best_or_equal_match (mae, min_mae, motion_vec, &mvec);
  else check_best_match (mae, min_mae, motion_vec, &mvec);
 }
}
/************************************************************/
// use search positions in the right block
for (srchi=0; srchi<BLK_SIZE; srchi++)
{
 // for each search position, compare all pixels
 mae = 0;
 for (i=0, refi=srchi; i<BLK_SIZE; i++, refi++)
 {
  for (j=0; j<BLK_SIZE; j++)
  {
    if (refi < BLK_SIZE) diff = cur_blk[i][j]-temp0[refi][j];
    else diff = cur_blk[i][j]-temp3[refi-BLK_SIZE][j];
    mae += abs(diff);
  }
 }
 check_best_match (mae, min_mae, motion_vec, &mvec);
}
/************************************************************/
// use search positions in the bottom-left block
for (srchj=0; srchj<BLK_SIZE; srchj++)
{
 // for each search position, compare all pixels
 mae = 0;
 for (i=0; i<BLK_SIZE; i++)
 {
  for (j=0, refj=srchj; j<BLK_SIZE; j++, refj++)
  {
    if (refj < BLK_SIZE) diff = cur_blk[i][j]-temp1[i][refj];
    else diff = cur_blk[i][j]-temp2[i][refj-BLK_SIZE];
    mae += abs(diff);
  }
 }
 check_best_match (mae, min_mae, motion_vec, &mvec);
}
```

```
/*************************************************************/
// use search positions in the bottom-centre block
for (srchj=0; srchj<BLK_SIZE; srchj++)
{
  // for each search position, compare all pixels
  mae = 0;
  for (i=0; i<BLK_SIZE; i++)
  {
    for (j=0, refj=srchj; j<BLK_SIZE; j++, refj++)
    {
      if (refj < BLK_SIZE) diff = cur_blk[i][j]-temp2[i][refj];
      else diff = cur_blk[i][j]-temp3[i][refj-BLK_SIZE];
      mae += abs(diff);
    }
  }
  check_best_match (mae, min_mae, motion_vec, &mvec);
}
/*************************************************************/
// use search positions in the bottom-right block
mae = 0;
for (i=0; i<BLK_SIZE; i++)
{
  for (j=0; j<BLK_SIZE; j++)
  {
    diff = cur_blk[i][j]-temp3[i][j];
    mae += abs(diff);
  }
}
check_best_match (mae, min_mae, motion_vec, &mvec);
}
/*************************************************************/
void check_best_match (cuint& mae, cuint& min_mae,
          cuint& motion_vec, unsigned short* mvec)
{
 // check if this search position is the best match so far
 cif (mae < min_mae)
 {
  min_mae = mae;
  motion_vec = *mvec;
 }
 cend
 (*mvec)++;
}

void check_best_or_equal_match (cuint& mae, cuint& min_mae,
      cuint& motion_vec, unsigned short* mvec)
{
 // special case for no movement in case similar object elsewhere
 cif ((mae < min_mae) || (mae == min_mae))
 {
  min_mae = mae;
  motion_vec = *mvec;
 }
 cend
 (*mvec)++;
}
```