



Software Optimization Guide for AMD Family 15h Processors

Publication No.	Revision	Date
47414	3.02	April 2011

© 2010, 2011 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. (“AMD”) products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein may be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD’s Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD’s products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD’s product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD Arrow logo, and combinations thereof, AMD Athlon, AMD Opteron, 3DNow!, AMD Virtualization and AMD-V are trademarks of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

Linux is a registered trademark of Linus Torvald.

Microsoft and Windows are registered trademarks of Microsoft Corporation.

MMX is a trademark of Intel Corporation.

PCI-X and PCI Express are registered trademarks of the PCI-Special Interest Group (PCI-SIG).

Solaris is a registered trademark of Sun Microsystems, Inc.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Contents

Revision History	15
Chapter 1 Introduction	17
1.1 Intended Audience	17
1.2 Getting Started	17
1.3 Using This Guide	18
1.3.1 Special Information	19
1.3.2 Numbering Systems	19
1.3.3 Typographic Notation	20
1.4 Important New Terms	20
1.4.1 Multi-Core Processors	20
1.4.2 Internal Instruction Formats	20
1.4.3 Types of Instructions	21
1.5 Key Optimizations	22
1.5.1 Implementation Guideline	22
1.6 What's New on AMD Family 15h Processors	22
1.6.1 AMD Instruction Set Enhancements	23
1.6.2 Floating-Point Improvements	23
1.6.3 Load-Execute Instructions for Unaligned Data	25
1.6.4 Instruction Fetching Improvements	25
1.6.5 Instruction Decode and Floating-Point Pipe Improvements	26
1.6.6 Notable Performance Improvements	26
1.6.7 Large Page Support	27
1.6.8 AMD Virtualization™ Optimizations	27
Chapter 2 Microarchitecture of AMD Family 15h Processors	29
2.1 Key Microarchitecture Features	30
2.2 Microarchitecture of AMD Family 15h Processors	30
2.3 Superscalar Processor	31
2.4 Processor Block Diagram	31

2.5	AMD Family 15h Processor Cache Operations	32
2.5.1	L1 Instruction Cache	33
2.5.2	L1 Data Cache	33
2.5.3	L2 Cache	33
2.5.4	L3 Cache	33
2.6	Branch-Prediction	34
2.7	Instruction Fetch and Decode	34
2.8	Integer Execution	35
2.9	Translation-Lookaside Buffer	35
2.9.1	L1 Instruction TLB Specifications	35
2.9.2	L1 Data TLB Specifications	35
2.9.3	L2 Instruction TLB Specifications	35
2.9.4	L2 Data TLB Specifications	36
2.10	Integer Unit	36
2.10.1	Integer Scheduler	36
2.10.2	Integer Execution Unit	36
2.11	Floating-Point Unit	37
2.12	Load-Store Unit	38
2.13	Write Combining	39
2.14	Integrated Memory Controller	39
2.15	HyperTransport™ Technology Interface	40
2.15.1	HyperTransport Assist	41
Chapter 3	C and C++ Source-Level Optimizations	43
3.1	Declarations of Floating-Point Values	44
3.2	Using Arrays and Pointers	45
3.3	Use of Function Prototypes	47
3.4	Unrolling Small Loops	47
3.5	Expression Order in Compound Branch Conditions	48
3.6	Arrange Boolean Operands for Quick Expression Evaluation	49
3.7	Long Logical Expressions in If Statements	50

3.8	Pointer Alignment	51
3.9	Unnecessary Store-to-Load Dependencies	52
3.10	Matching Store and Load Size	53
3.11	Use of const Type Qualifier	56
3.12	Generic Loop Hoisting	56
3.13	Local Static Functions	59
3.14	Explicit Parallelism in Code	59
3.15	Extracting Common Subexpressions	62
3.16	Sorting and Padding C and C++ Structures	63
3.17	Replacing Integer Division with Multiplication	64
3.18	Frequently Dereferenced Pointer Arguments	65
3.19	32-Bit Integral Data Types	66
3.20	Sign of Integer Operands	67
3.21	Improving Performance in Linux [®] Libraries	68
3.22	Aligning Matrices	69
Chapter 4	General 64-Bit Optimizations	71
4.1	64-Bit Registers and Integer Arithmetic	71
4.2	Using 64-bit Arithmetic for Large-Integer Multiplication	73
4.3	128-Bit Media Instructions and Floating-Point Operations	77
4.4	32-Bit Legacy GPRs and Small Unsigned Integers	77
Chapter 5	Instruction-Decoding Optimizations	79
5.1	Load-Execute Instructions for Floating-Point or Integer Operands	79
5.1.1	Load-Execute Integer Instructions	80
5.1.2	Load-Execute SIMD Instructions with Floating-Point or Integer Operands	81
5.2	32/64-Bit vs. 16-Bit Forms of the LEA Instruction	82
5.3	Take Advantage of x86 and AMD64 Complex Addressing Modes	82
5.4	Short Instruction Encodings	84
5.5	Partial-Register Writes	84
5.6	Using LEAVE for Function Epilogues	89
5.7	Alternatives to SHLD Instruction	90

5.8	Code Padding with Operand-Size Override and Multibyte NOP	92
Chapter 6	Cache and Memory Optimizations	95
6.1	Memory-Size Mismatches	95
6.2	Natural Alignment of Data Objects	97
6.3	Store-to-Load Forwarding Restrictions	98
6.4	Good Practices for Avoiding False Store-to-Load Forwarding	104
6.5	Prefetch and Streaming Instructions	105
6.6	Write-Combining	114
6.7	L1 Data Cache Bank Conflicts	114
6.8	Placing Code and Data in the Same 64-Byte Cache Line	116
6.9	Memory and String Routines	117
6.10	Stack Considerations	119
6.11	Cache Issues When Writing Instruction Bytes to Memory	120
6.12	Interleave Loads and Stores	121
Chapter 7	Branch Optimizations	123
7.1	Instruction Fetch	123
7.1.1	Instruction Fetch	123
7.1.2	Reduce Instruction Size	124
7.2	Branch Fusion	124
7.3	Branches That Depend on Random Data	125
7.4	Pairing CALL and RETURN	126
7.5	Nonzero Code-Segment Base Values	128
7.6	Replacing Branches	128
7.7	Avoiding the LOOP Instruction	130
7.8	Far Control-Transfer Instructions	130
7.9	Branches Not-Taken Preferable to Branches Taken	131
Chapter 8	Scheduling Optimizations	133
8.1	Instruction Scheduling by Latency	133
8.2	Loop Unrolling	133
8.3	Inline Functions	138

8.4	MOVZX and MOVSX	139
8.5	Pointer Arithmetic in Loops	139
8.6	Pushing Memory Data Directly onto the Stack	141
Chapter 9	Integer Optimizations	143
9.1	Replacing Division with Multiplication	143
9.2	Alternative Code for Multiplying by a Constant	147
9.3	Repeated String Instructions	150
9.4	Using XOR to Clear Integer Registers	151
9.5	Efficient 64-Bit Integer Arithmetic in 32-Bit Mode	152
9.6	Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants	159
9.7	Optimizing Integer Division	165
9.8	Efficient Implementation of Population Count and Leading-Zero Count	166
Chapter 10	Optimizing with SIMD Instructions	167
10.1	Ensure All Packed Floating-Point Data are Aligned	168
10.2	Explicit Load Instructions	168
10.3	Unaligned and Aligned Data Access	169
10.4	Moving Data Between General-Purpose and XMM/YMM Registers	169
10.5	Use SIMD Instructions to Construct Fast Block-Copy Routines	170
10.6	Using SIMD Instructions for Fast Square Roots and Divisions	171
10.7	Use XOR Operations to Negate Operands of SIMD Instructions	174
10.8	Clearing SIMD Registers with XOR Instructions	175
10.9	Finding the Floating-Point Absolute Value of Operands of SIMD Instructions	176
10.10	Accumulating Single-Precision Floating-Point Numbers Using SIMD Instructions	176
10.11	Complex-Number Arithmetic Using AVX Instructions	178
10.12	Optimized 4 X 4 Matrix Multiplication on 4 X 1 Column Vector Routines	181
10.13	Floating-Point-to-Integer Conversion	183
10.14	Reuse of Dead Registers	184
10.15	Floating-Point Scalar Conversions	185

10.16	Move/Compute Optimization	186
10.17	Using SIMD Instructions for Rounding	188
10.18	Using SIMD Instructions for Floating-Point Comparisons	189
Chapter 11	Multiprocessor Considerations	191
11.1	ccNUMA Optimizations	191
11.2	Writing Instruction Bytes to Memory on Multiprocessor Systems	200
11.3	Multithreading	202
11.4	Memory Barrier Operations	207
11.5	Optimizing Inter-Core Data Transfer	210
Chapter 12	Optimizing Secure Virtual Machines	215
12.1	Use Nested Paging	216
12.2	VMCB.G_PAT Configuration	217
12.3	State Swapping	217
12.4	Economizing Interceptions	218
12.5	Nested Page Size	219
12.6	Shadow Page Size	220
12.7	Setting VMCB.TLB_Control	220
12.8	TLB Flushes in Shadow Paging	221
12.9	Use of Virtual Interrupt VMCB Field	222
12.10	Avoid Instruction Fetch for Intercepted Instructions	223
12.11	Share IOIO and MSR Protection Maps	224
12.12	Obey CPUID Results	224
12.13	Using Time Sources	225
12.14	Paravirtualized Resources	226
Appendix A	Implementation of Write-Combining	227
A.1	Write-Combining Definitions and Abbreviations	227
A.2	Programming Details	228
A.3	Write-Combining Operations	228
A.4	Sending Write-Buffer Data to the System	229
A.5	Write Combining to MMI/O Devices that Support Write Chaining	229

Appendix B	Instruction Latencies	233
B.1	Understanding Instruction Entries	233
B.2	General Purpose and Integer Instruction Latencies	237
B.3	System Instruction Latencies	250
B.4	FPU Instruction Latencies	253
B.5	Amended Latency for Selected FMA Instructions	309
Appendix C	Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems	313
C.1.1	Support Under Linux [®]	313
C.1.2	Support under Microsoft [®] Windows [®]	314
C.1.3	Hardware Support for System Topology Discovery	315
C.1.4	Support for Stability	315
C.2	Tools and APIs for Memory Node Interleaving	316
C.2.1	Support under Linux [®]	316
C.2.2	Support under Solaris [™]	316
C.2.3	Support under Microsoft [®] Windows [®]	316
C.2.4	Memory Node Interleaving Configuration in the BIOS	316
Appendix D	NUMA Optimizations for I/O Devices	319
D.1	AMD64 System Topologies	319
D.2	Optimization Strategy	319
D.3	Identifying Nodes that Have Noncoherent HyperTransport [™] I/O Links	321
D.4	Access of PCI Configuration Register	327
D.5	I/O Thread Scheduling	329
D.6	Using Write-Only Buffers for Device Consumption	330
D.7	Using Interrupt Affinity	330
Appendix E	Remarks on the RDTSC(P) Instruction	333
Appendix F	Guide to Instruction-Based Sampling on AMD Family 15h Processors	335
F.1	Background	335
F.2	Overview	336
F.3	IBS fetch sampling	337
F.3.1	Taking an IBS fetch sample	337

F.3.2	Interpreting IBS fetch data	338
F.4	IBS op sampling	340
F.4.1	Taking an IBS op sample	340
F.4.2	Interpreting IBS op data	341
F.4.3	Interpreting IBS branch/return/resync op data	342
F.4.4	Interpreting IBS Load/Store Data	344
F.4.5	Interpreting IBS load/store Northbridge data	345
F.5	Software-based analysis	347
F.5.1	Derived events and post-processing	348
F.5.2	Derived events for IBS fetch data	349
F.5.3	Derived Events for all Ops	351
F.5.4	Derived events for IBS branch/return/resync ops	351
F.5.5	Derived events for IBS load/store operations	351
F.6	Derived Events for Northbridge Activity	353

Tables

Table 1.	Instructions, Macro-ops and Micro-ops	21
Table 2.	Optimizations by Rank.....	22
Table 3.	Prefetching Guidelines	107
Table 4.	DIV/IDIV Latencies.....	166
Table 5.	Single-Precision Floating-Point Scalar Conversion.....	185
Table 6.	Double-Precision Floating-Point Scalar Conversion	186
Table 7.	Write-Combining Completion Events.....	228
Table 8.	Mapping of Pipes to Floating-Point Units	234
Table 9.	Latency Formats.....	236
Table 10.	General Purpose and Integer Instruction Latencies	237
Table 11.	System Instruction Latencies	250
Table 12:	FPU Instruction Latencies.....	253
Table 13.	FPU Amended Latencies.....	309
Table 14.	Size of Base Address Register	323
Table 15.	IBS Hardware Event Flags.....	338
Table 16.	Event Flag Combinations.....	339
Table 17.	IbsOpData MSR Event Flags and Counts.....	343
Table 18.	Execution Status Indicated by IbsOpBrnMisp and IbsOpBrnTaken Flags.....	343
Table 19.	Execution Status Indicated by IbsOpReturn and IbsOpMispReturn Flags.....	343
Table 20.	IbsOpData3 Register Information.....	344
Table 21.	IbsOpData2 Register Fields	346
Table 22.	Northbridge Request Data Source Field	347
Table 23.	IBS Northbridge Event Data	347
Table 24.	An IBS Fetch Sample.....	348
Table 25.	2-D Table of IBS Fetch Samples	348
Table 26.	New Events Derived from Combined Event Flags.....	349
Table 27.	Derived Events for All Ops.....	351
Table 28.	Derived Events to Measure Branch, Return and Resync Ops.....	351

Table 29.	Derived Events for Ops That Perform Load and/or Store Operations	352
Table 30.	IBS Northbridge Derived Events	353

Figures

Figure 1.	AMD Family 15h Processors Block Diagram	32
Figure 2.	Integer Execution Unit Block Diagram	36
Figure 3.	Floating-Point Unit	38
Figure 4.	Load-Store Unit	39
Figure 5.	Memory-Limited Code	112
Figure 6.	Processor-Limited Code	112
Figure 7.	Simple SMP Block Diagram.....	192
Figure 8.	AMD 2P System	193
Figure 9.	Dual AMD Family 15h Processor Configuration.....	193
Figure 10.	Block Diagram of a ccNUMA AMD Family 15h Quad-Core Multiprocessor System.....	194
Figure 11.	Link Type Registers F0x[F8, D8, B8, 98]	322
Figure 12.	MMIO Base Low Address Registers F1x[B8h, B0h, A8h, A0h, 98h, 90h, 88h, 80h]	324
Figure 13.	MMIO Limit Low Address Registers F1x[1BCh, 1B4h, 1ACh, 1A4h, 9Ch, 94h, 8Ch, 84h]	324
Figure 14.	MMIO Base/Limit High Address Registers F1x[1CCh, 1C8h, 1C4h, 1C0h, 19Ch, 198h, 194h, 190h, 18Ch, 188h]	324
Figure 15.	Configuration Map Registers F1x[E0h, E4h, E8h, ECh]	326
Figure 16.	Configuration Address Register (0CF8h)	328
Figure 17.	Configuration Data Register (0CFCh)	328
Figure 18.	Histogram for the IBS Fetch Completed Derived Event	349

Revision History

Date	Rev.	Description
April 2011	3.02	Initial Public Release

Chapter 1 Introduction

This guide provides optimization information and recommendations for AMD Family 15h processors. These optimizations are designed to yield software code that is fast, compact, and efficient. Toward this end, the optimizations in each of the following chapters are listed in order of importance.

This chapter covers the following topics:

Topic	Page
Intended Audience	17
Getting Started	17
Using This Guide	18
Important New Terms	20
Key Optimizations	22
What's New on AMD Family 15h Processors	22

1.1 Intended Audience

This book is intended for compiler and assembler designers, as well as C, C++, and assembly-language programmers writing performance-sensitive code sequences. This guide assumes that you are familiar with the AMD64 instruction set and the AMD64 architecture (registers and programming modes). For complete information on the AMD64 architecture and instruction set, see the multivolume *AMD64 Architecture Programmer's Manual* available from AMD.com. Individual volumes and their order numbers are provided below.

Title	Order Number
Volume 1: <i>Application Programming</i>	24592
Volume 2: <i>System Programming</i>	24593
Volume 3: <i>General-Purpose and System Instructions</i>	24594
Volume 4: <i>128-Bit and 256-Bit Media Instructions</i>	26568
Volume 5: <i>64-Bit Media and x87 Floating-Point Instructions</i>	26569
Volume 6: <i>128-Bit and 256-Bit XOP, FMA4 and CVT16 Instructions</i>	43479

1.2 Getting Started

More experienced readers may skip to “Key Optimizations” on page 22, which identifies the most important optimizations, and to “What’s New on AMD Family 15h Processors” on page 22 for a quick review of key new performance enhancement features introduced with AMD Family 15h processors.

1.3 Using This Guide

Each of the remaining chapters in this document focuses on a particular general area of relevance to software optimization on AMD Family 15h processors. Each chapter is organized into a set of one or more recommended related optimizations pertaining to a particular issue. These sections are divided into three sections:

- **Optimization**—Specifies the recommended action required for achieving the optimization under consideration.
- **Application**—Specifies the type of software for which the particular optimization is relevant (*i.e.*, to 32-bit software or 64-bit software or to both).
- **Rationale**—Provides additional explanatory technical information regarding the particular optimization. This section usually provides illustrative C, C++, or assembly code examples as well.

The chapters that follow cover the following topics:

- Chapter 2, “Microarchitecture of AMD Family 15h Processors,” discusses the internal design, or microarchitecture, of the AMD Family 15h processor and provides information about translation-lookaside buffers and other functional units that, while not part of the main processor, are integrated on the chip.
- Chapter 3, “C and C++ Source-Level Optimizations,” describes techniques that you can use to optimize your C and C++ source code.
- Chapter 4, “General 64-Bit Optimizations,” presents general assembly-language optimizations that can improve the performance of software designed to run in 64-bit mode. The optimizations in this chapter apply *only* to 64-bit software.
- Chapter 5 “Instruction-Decoding Optimizations,” discusses optimizations designed to maximize the number of instructions that the processor can decode at one time.
- Chapter 6 “Cache and Memory Optimizations,” discusses how to take advantage of the large L1 caches and high-bandwidth buses.
- Chapter 7, “Branch Optimizations,” discusses improving branch prediction and minimizing branch penalties.
- Chapter 8, “Scheduling Optimizations,” discusses improving instruction scheduling in the processor.
- Chapter 9, “Integer Optimizations,” discusses integer performance.
- Chapter 10, “Optimizing with SIMD Instructions,” discusses the 64-bit and 128-bit SIMD instructions used to encode floating-point and integer operations.
- Chapter 11, “Multiprocessor Considerations,” discusses processor/core selection and related issues for applications running on multiprocessor/multicore cache coherent non-uniform memory access (ccNUMA) configurations.

- Chapter 12, “Optimizing Secure Systems,” discusses ways to minimize the performance overhead imposed by the virtualization of a guest.
- Appendix A, “Implementation of Write-Combining,” describes how AMD Family 15h processors perform memory write-combining.
- Appendix B, “Instruction Latencies,” provides a complete listing of all AMD64 instructions with each instruction’s decode type, execution latency, and—where applicable—the pipes and throughput used in the floating-point unit.
- Appendix C, “Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems” provides information on tools for programming in NUMA environments.
- Appendix D, “NUMA Optimizations for I/O Devices” provides information on the association of particular I/O devices with a specific nodes in a NUMA system.
- Appendix E, “Remarks on the RDTSC(P) Instruction” provides information on using the RDTSC and RDTSCP instructions to load the value of the time stamp counter (TSC).

1.3.1 Special Information

Special information in this guide is marked as follows:

❖ This symbol appears next to the most important, or key, optimizations.

1.3.2 Numbering Systems

The following suffixes identify different numbering systems:

This suffix	Identifies a
b	Binary number. For example, the binary equivalent of the number 5 is written 101b.
d	Decimal number. Decimal numbers are followed by this suffix only when the possibility of confusion exists. In general, decimal numbers are shown without a suffix.
h	Hexadecimal number. For example, the hexadecimal equivalent of the number 60 is written 3Ch.

1.3.3 Typographic Notation

This guide uses the following typographic notations for certain types of information:

This type of text	Identifies
<i>italic</i>	Placeholders that represent information you must provide. Italicized text is also used for the titles of publications and for emphasis.
monowidth	Program statements and function names.

1.4 Important New Terms

This section defines several important terms and concepts used in this guide.

1.4.1 Multi-Core Processors

AMD Family 15h processors have multiple compute units, each containing its own L2 cache and two cores. The cores share their compute unit's L2 cache. Each core incorporates the complete x86 instruction set logic and L1 data cache. Compute units share the processor's L3 cache and Northbridge (see Chapter 2, Microarchitecture of AMD Family 15h Processors).

1.4.2 Internal Instruction Formats

AMD Family 15h processors perform four types of *primitive operations*:

- Integer (arithmetic or logic)
- Floating-point (arithmetic)
- Load
- Store

The AMD64 instruction set is complex. Instructions have variable-length encoding and many perform multiple primitive operations. AMD Family 15h processors do not execute these complex instructions directly, but, instead, decode them internally into simpler fixed-length instructions called *macro-ops*. Processor schedulers subsequently break down macro-ops into sequences of even simpler instructions called *micro-ops*, each of which specifies a single primitive operation.

A *macro-op* is a fixed-length instruction that:

- Expresses, at most, one integer or floating-point operation and one load and/or store operation.
- Is the primary unit of work managed (that is, dispatched and retired) by the processor.

A *micro-op* is a fixed-length instruction that:

- Expresses one and only one of the primitive operations that the processor can perform (for example, a load).
- Is executed by the processor's execution units.

Table 1 on page 21 summarizes the differences between AMD64 instructions, macro-ops, and micro-ops.

Table 1. Instructions, Macro-ops and Micro-ops

Comparing	AMD64 instructions	Macro-ops	Micro-ops
Complexity	Complex A single instruction may specify one or more of each of the following operations: <ul style="list-style-type: none"> • Integer or floating-point • Load • Store 	Average A single macro-op may specify—at most—one integer or floating-point operation and one of the following operations: <ul style="list-style-type: none"> • Load • Store • Load and store to the same address 	Simple A single micro-op specifies only one of the following primitive operations: <ul style="list-style-type: none"> • Integer or floating-point • Load • Store
Encoded length	Variable (instructions are different lengths)	Fixed (all macro-ops are the same length)	Fixed (all micro-ops are the same length)
Regularized instruction fields	No (field locations and definitions vary among instructions)	Yes (field locations and definitions are the same for all macro-ops)	Yes (field locations and definitions are the same for all micro-ops)

1.4.3 Types of Instructions

Instructions are classified according to how they are decoded by the processor. There are three types of instructions:

Instruction Type	Description
FastPath Single	Decodes directly into one macro-op in microprocessor hardware.
FastPath Double	Decodes directly into two macro-ops in microprocessor hardware.
Microcode	Decodes into one or more (usually three or more) macro-ops using the on-chip microcode-engine ROM (MROM).

1.5 Key Optimizations

While all of the optimizations in this guide help improve software performance, some of them have more impact than others. Optimizations that offer the most improvement are called *key* optimizations.

❖ This symbol appears next to the most important (key) optimizations.

1.5.1 Implementation Guideline

Concentrate your efforts on implementing key optimizations before moving on to other optimizations.

Table 2 lists the key optimizations. These optimizations are discussed in detail in later sections of this book.

Table 2. Optimizations by Rank

Rank	Optimization
1	Load-Execute Instructions for Floating-Point or Integer Operands (See section 5.1 on page 79.)
2	Write-Combining (See section 6.6 on page 114.)
3	Branches That Depend on Random Data (See section 7.3 on page 125.)
4	Loop Unrolling (See section 8.2 on page 133.)
5	Pointer Arithmetic in Loops (See section 8.5 on page 139.)
6	Explicit Load Instructions (See section 10.2 on page 168.)
7	Reuse of Dead Registers (See section 10.14 on page 184.)
8	ccNUMA Optimizations (See section 11.1 on page 191.)
9	Multithreading (See section 11.3 on page 202.)
10	Prefetch and Streaming Instructions (See section 6.5 on page 105.)
11	Memory and String Routines (See section 6.9 on page 117.)
12	Floating-Point Scalar Conversions (See sections 10.15 on page 185.)

1.6 What's New on AMD Family 15h Processors

AMD Family 15h processors introduce several new features that can significantly enhance software performance when compared to the previous AMD64 microprocessors. The following section provides a summary of these performance improvements. Throughout this discussion, it is assumed that readers are familiar with the software optimization guide for the previous AMD64 processors and the terminology used there.

1.6.1 AMD Instruction Set Enhancements

The AMD Family 15h processor has been enhanced with the following new instructions:

- XOP and AVX support—Extended Advanced Vector Extensions provide enhanced instruction encodings and non-destructive operands with an extended set of 128-bit (XMM) and 256-bit (YMM) media registers
- FMA instructions—support for floating-point fused multiply accumulate instructions
- Fractional extract instructions—extract the fractional portion of vector and scalar single-precision and double-precision floating-point operands
- Support for new vector conditional move instructions.
- VPERMILx instructions—allow selective permutation of packed double- and single-precision floating point operands
- VPHADDx/VPSUBx—support for packed horizontal add and subtract instructions
- Support for packed multiply, add and accumulate instructions
- Support for new vector shift and rotate instructions

Support for these instructions is implementation dependent. See the *CPUID Specification*, order# 25481, and the *AMD64 Architecture Programmer's Manual Updates Application Note*, order# 33633, for additional information.

1.6.2 Floating-Point Improvements

AMD Family 15h processors add support for 128-bit floating-point execution units. As a result, the throughput of both single-precision and double-precision floating-point SIMD vector operations has improved by 2X over the previous generation of AMD processors.

Users may notice differences in the results of programs when using the fused multiply and add FMAC. These differences do not imply that the new results are less accurate than using the ADD and MUL instructions separately. These differences result from the combination of an ADD and a MUL into a single instruction. As separate instructions, ADD and MUL provide a result which is accurate to $\frac{1}{2}$ a bit in the least significant bit for the precision provided. However, the combined result of the ADD and the MUL is *not* accurate to $\frac{1}{2}$ a bit.

By fusing these two instructions into a “single” instruction, a fused multiply accumulate (FMAC), an accurate result is provided that *is* within $\frac{1}{2}$ a bit in the in least significant bit. Thus the difference between performing “separate” ADDs and MULs and doing a “single” FMAC is the cause of differences in the least significant bit of program results.

Performance Guidelines for Vectorized Floating-Point SIMD Code

While 128-bit floating-point execution units imply better performance for vectorized floating-point SIMD code, it is necessary to adhere to several performance guidelines to realize their full potential:

- Avoid writing less than 128 bits of an XMM register when using certain initializing and non-initializing operations.

A floating-point XMM register is viewed as one 128-bit register internally by the processor. Writing to a 64-bit half of a 128-bit XMM register results in a merge dependency on the other 64-bit half. Therefore the following replacements are advised on AMD Family 15h processors:

- Replace `MOVLPx/MOVHPx reg, mem` pairs with `MOVUPx reg, mem`, irrespective of the alignment of the data. On AMD Family 15h processors, the `MOVUPx` instruction is just as efficient as `MOVAPx`, which is designed for use with aligned data. Hence it is advised to use `MOVUPx` regardless of the alignment.
- Replace `MOVLPD reg, mem` with `MOVSD reg, mem`.
- Replace `MOVSD reg, reg` with `MOVAPD reg, reg`.

However, there are also several instructions that initialize the lower 64 or 32 bits of an XMM register and zero out the upper 64 or 96 bits and, thus, do not suffer from such merge dependencies. Consider, for example, the following instructions:

```
MOVSD xmm, [mem64]
MOVSS xmm, [mem32]
```

When writing to a register during the course of a non-initializing operation on the register, there is usually no additional performance loss due to partial register reads and writes. This is because in the typical case, the partial register that is being written is also a source to the operation. For example, `addsd xmm1, xmm2` does not suffer from merge dependencies.

There are often cases of non-initializing operations on a register, in which the partial register being written by the operation is not a source for the operation. In these cases also, it is preferable to avoid partial register writes. If it is not possible to avoid writing to a part of that register, then you should schedule any prior operation on any part of that register well ahead of the point where the partial write occurs.

Examples of non-initializing instructions that result in merge dependencies are `SQRTSD`, `CVTPI2PS`, `CVTSI2SD`, `CVTSS2SD`, `MOVLHPS`, `MOVHLPS`, `UNPCKLPD` and `PUNPCKLQDQ`.

For additional details on this optimization see “Partial-Register Writes” on page 84, “Explicit Load Instructions” on page 168, “Unaligned and Aligned Data Access” on page 169, and “Reuse of Dead Registers” on page 184.

- Legacy SIMD instructions themselves always merge the upper `YMM[255:128]` bits. AMD family 15h processors keep track of two Zero bits: one for double-precision floating-point values (`ZD = (dest[127:64]==0)`), and one for single-precision floating-point values (`ZS = (dest[127:32]==0)`). `ZS` implies a `ZD`. Most SIMD instructions are merging destination bits `[127:64]` or `[127:32]` for scalar double and single respectively. Some operations force the output to 0 for these bits—that is, when we set the `ZD/ZS` bits. We then propagate them through dependency chains, so that for a few key operations we can break the false dependency. (Most

merging operations have real dependencies on the lower bits, and any dependency on the upper bits are irrelevant).

In the past, the combination of MOVLPD/MOVHPD instructions was used instead of MOVAPD (or MOVUPD). Without optimization, the MOVLPD/MOVHPD instruction pair would have a false dependency on a previous loop iteration, while the MOVAPD instruction would not. By optimizing, we can convert those cases that we can detect and remove false dependencies resulting from the use of MOVLPD/MOVHPD. In the long run, it is still better to avoid the issue and use the MOVAPD instruction in the first place, instead of MOVLPD/MOVHPD.

- In the event of a load following a previous store to a given address for aligned floating-point vector data, use 128-bit stores and 128-bit loads instead of MOVLPX/MOVHPX pairs for storing and loading the data. This allows store-to-load forwarding to occur. Using MOVLPX/MOVHPX pairs is still recommended for storing unaligned floating-point vector data. Additional details on these restrictions can be obtained in “Store-to-Load Forwarding Restrictions” on page 98.
- To make use of the doubled throughput of both single-precision and double-precision floating-point SIMD vector operations, a compiler or an application developer can consider either increasing the unrolling factor of loops that include such vector operations and/or performing other code transformations to keep the floating-point pipeline fully utilized.

1.6.3 Load-Execute Instructions for Unaligned Data

Use load-execute instructions instead of discrete load and execute instructions when performing SIMD integer, SIMD floating-point and x87 computations on floating-point source operands. This is recommended regardless of the alignment of packed data on AMD Family 15h processors. (The use of load-execute instructions under these circumstances was only recommended for aligned packed data on the previous AMD64 processors.) This replacement is only possible if the misaligned exception mask (MM) is set. See the *AMD CPUID Specification*, order# 25481, and the *AMD64 Architecture Programmer's Manual Updates Application Note*, order# 33633, for additional information on SIMD misaligned access support. This optimization can be especially useful in vectorized SIMD loops and may eliminate the need for loop peeling due to nonalignment. (See “Load-Execute Instructions for Floating-Point or Integer Operands” on page 79.)

1.6.4 Instruction Fetching Improvements

While previous AMD64 processors had a single 32-byte fetch window, AMD Family 15h processors have two 32-byte fetch windows, from which four μ ops can be selected. These fetch windows, when combined with the 128-bit floating-point execution unit, allow the processor to sustain a fetch/dispatch/retire sequence of four instructions per cycle. Most instructions decode to a single μ op, but fastpath double instructions decode to two μ ops. ALU instructions can also issue four μ ops per cycle and microcoded instructions should be considered single issue. Thus, there is not necessarily a one-to-one correspondence between the decode size of assembler instructions and the capacity of the 32-byte fetch window and the production of optimal assembler code requires considerable attention to the details of the underlying programming constraints.

Assembly language programmers can now group more instructions together but must still concern themselves with the possibility that an instruction may span a 32-byte fetch window. In this regard, it is also advisable to align hot loops to 32 bytes instead of 16 bytes, especially in the case of loops for large SIMD instructions. See Chapter 7, “Branch Optimizations” on page 123 for details.

1.6.5 Instruction Decode and Floating-Point Pipe Improvements

Several integer and floating-point instructions have improved latencies and decode types on AMD Family 15h processors. Furthermore, the FPU pipes utilized by several floating-point instructions have changed. These changes can influence instruction choice and scheduling for compilers and hand-written assembly code. A comprehensive listing of all AMD64 instructions with their decode types, decode type changes from previous families of AMD processors, and execution latencies and FPU pipe utilization data are available in Appendix C.

1.6.6 Notable Performance Improvements

Several enhancements to the AMD64 architecture have resulted in significant performance improvements in AMD Family 15h processors, including:

- Improved performance of shuffle instructions
- Improved data transfer between floating-point registers and general purpose registers
- Improved floating-point register to floating-point register moves
- Optimization of repeated move instructions
- More efficient PUSH/POP stack operations
- 1-Gbyte paging

These are discussed in the following paragraphs and elsewhere in this document.

Improved Bandwidth Decode Type for Shuffle Instructions

The floating-point logic in AMD Family 15h processors uses three separate execution positions or pipes called FADD, FMUL and FSTORE. This is illustrated in Figure 1 on page 32 in Appendix A. Current AMD Family 15h processors support two SIMD logical/shuffle units, one in the FMUL pipe and another in the FADD pipe, while previous AMD64 processors have only one SIMD logical/shuffle unit in the FMUL pipe. As a result, the SIMD shuffle instructions can be processed at twice the previous bandwidth on AMD Family 15h processors. Furthermore, the PSHUFD and SHUFPx shuffle instructions are now DirectPath instructions instead of VectorPath instructions on AMD Family 15h processors and take advantage of the 128-bit floating point execution units. Hence, these instructions get a further 2X boost in bandwidth, resulting in an overall improvement of 4X in bandwidth compared to the previous generation of AMD processors.

It's more efficient to use SHUFPx and PSHUFD instructions over combinations of more than one MOVLHPS/MOVHLPS/UNPCKx/PUNPCKx instructions to do shuffle operations.

Data Transfer Between Floating-Point Registers and General Purpose Integer Registers

We recommend using the MOVD instruction when moving data from an MMX™ or XMM register to a GPR. However, when moving data from a GPR to an MMX or XMM floating-point register, it is advisable to use separate store and load instructions to move the data from the source register to a temporary location in memory and then from memory into the destination register, taking the memory latency into account when scheduling them.

The performance of the CVTSS2SI, CVTTSS2SI, CVTSD2SI, CVTTSD2SI instructions that are used to convert floating-point data to integer data has improved. For additional details see “Floating-Point-to-Integer Conversion” on page 183.

Floating-Point Register-to-Register Moves

On previous AMD processors, floating-point register-to-register moves could only go through the FADD and FMUL pipes. On AMD Family 15h processors, floating-point register-to-register moves can also go through the FSTORE pipe, thereby improving overall throughput.

Repeated String Instructions

REP instructions have been optimized on AMD Family 15h processors. See “Repeated String Instructions” on page 150 for details on how to take advantage of these optimizations.

1.6.7 Large Page Support

AMD Family 15h processors now have better large page support, having incorporated new 1GB paging and 2MB and 4KB paging improvements.

The L1 data TLB and L2 data TLB now support 1GB pages, a benefit to applications making large data-set random accesses.

The L1 instruction TLB, L1 data TLB and L2 data TLB have increased the number of entries for 2MB pages. This improves the performance of software that uses 2MB code or data or code mixed with data virtual pages.

The L1 data TLB has also increased the number of entries for 4KB pages.

For additional details on the actual number of TLB entries, see section A.10, “Translation-Lookaside Buffer” on page 35.

1.6.8 AMD Virtualization™ Optimizations

Chapter 12, “Optimizing Secure Virtual Machines” covers optimizations that minimize the performance overhead imposed by the virtualization of a guest in AMD Virtualization™ technology (AMD-V™). Topics include:

- The advantages of using nested paging instead of shadow paging

- Guest page attribute table (PAT) configuration
- State swapping
- Economizing Interceptions
- Nested page and shadow page size
- TLB control and flushing in shadow pages
- Instruction Fetch for Intercepted (REP) INS instructions
- Sharing IOIO and MSR protection masks
- CPUID
- Time resources
- Paravirtualized resources

Chapter 2 Microarchitecture of AMD Family 15h Processors

An understanding of the terms *architecture*, *microarchitecture*, and *design implementation* is important when discussing processor design.

The *architecture* consists of the instruction set and those features of a processor that are visible to software programs running on the processor. The architecture determines what software the processor can run. The AMD64 architecture of the AMD Family 15h processors is compatible with the industry-standard x86 instruction set.

The term *microarchitecture* refers to the design features used to reach the target cost, performance, and functionality goals of the processor. The AMD Family 15h processor employs a decoupled decode/execution design approach. In other words, decoders and execution units operate essentially independently; the execution core uses a small number of instructions and a simplified circuit design implementation to achieve fast single-cycle execution with fast operating frequencies.

The *design implementation* refers to a particular combination of physical logic and circuit elements that comprise a processor that meets the microarchitecture specifications.

This appendix covers the following topics:

Topic	Page
Key Microarchitecture Features	30
Microarchitecture of AMD Family 15h Processors	30
Superscalar Processor	31
Processor Block Diagram	31
AMD Family 15h Processor Cache Operations	32
Branch-Prediction	34
Instruction Fetch and Decode	34
Integer Execution	35
Translation-Lookaside Buffer	35
Integer Unit	36
Floating-Point Unit	37
Load-Store Unit	38
Write Combining	39
Integrated Memory Controller	39
HyperTransport™ Technology Interface	40

2.1 Key Microarchitecture Features

AMD Family 15h processors include many features designed to improve software performance. The internal design, or *microarchitecture*, of these processors provides the following key features:

- Integrated DDR3 memory controller with memory prefetcher
- 64-Kbyte L1 instruction cache and 16-Kbyte L1 data cache
- Shared L2 cache between cores of compute unit
- Shared L3 cache compute units on chip (for supported platforms)
- 32-byte instruction fetch
- Instruction predecode and branch prediction during cache-line fills
- Decoupled prediction and instruction fetch pipelines
- Four-way AMD64 instruction decoding (This is a theoretical limit. See section 2.3 on page 31.)
- Dynamic scheduling and speculative execution
- Two-way integer execution
- Two-way address generation
- Two-way 128-bit wide floating-point execution
- Legacy single-instruction multiple-data (SIMD) instruction extensions, as well as support for XOP, FMA4, VPERMIL x , and Advanced Vector Extensions (AVX).
- Superforwarding
- Prefetch into L2 or L1 data cache
- Deep out-of-order integer and floating-point execution
- HyperTransport™ technology

2.2 Microarchitecture of AMD Family 15h Processors

AMD Family 15h processors implement the AMD64 instruction set by means of *macro-ops* (the primary units of work managed by the processor) and *micro-ops* (the primitive operations executed in the processor's execution units). These are simple fixed-length operations designed to include direct support for AMD64 instructions and adhere to the high-performance principles of fixed-length encoding, regularized instruction fields, and a large register set. This enhanced microarchitecture enables higher processor core performance and promotes straightforward extensibility for future designs.

2.3 Superscalar Processor

The AMD Family 15h processors are aggressive, out-of-order, four-way superscalar AMD64 processors. They can theoretically fetch, decode, and issue up to four AMD64 instructions per cycle using decoupled fetch and branch prediction units and three independent instruction schedulers, consisting of two integer schedulers and one floating-point scheduler.

These processors can fetch 32 bytes per cycle and can scan two 16-byte instruction windows for up to four micro-ops, which can be dispatched together in a single cycle. However, this is a theoretical limit. The actual number of micro-ops that are dispatched may be lower, depending on a number of factors, such as whether the processor is executing in fast or slow mode and whether instructions can be broken up into 16-byte windows. The processors move integer instructions through the replicated Integer clusters and floating point instructions through the shared floating point unit (FPU), as shown in Figure 1. on page 32.

2.4 Processor Block Diagram

A block diagram of the AMD Family 15h processors is shown in Figure 1 on page 32.




Figure 1. AMD Family 15h Processors Block Diagram

2.5 AMD Family 15h Processor Cache Operations

AMD Family 15h processors use four different caches to accelerate instruction execution and data processing:

- L1 instruction cache
- L1 data cache
- Share compute unit L2 cache
- Shared on chip L3 cache (on supported platforms)

2.5.1 L1 Instruction Cache

The out-of-order execution engine of AMD Family 15h processors contains a 64-Kbyte, 2-way set-associative L1 instruction cache. Each line in this cache is 64 bytes long. However, only 32 bytes are fetched in every cycle. Functions associated with the L1 instruction cache are instruction loads, instruction prefetching, instruction predecoding, and branch prediction. Requests that miss in the L1 instruction cache are fetched from the L2 cache or, subsequently, from the L3 cache or system memory.

On misses, the L1 instruction cache generates fill requests to a naturally aligned 64-byte line containing the instructions and the next sequential line of bytes (a prefetch). Because code typically exhibits spatial locality, prefetching is an effective technique for avoiding decode stalls. Cache-line replacement is based on a least-recently-used replacement algorithm.

Predecoding begins as the L1 instruction cache is filled. Predecode information is generated and stored alongside the instruction cache. This information is used to help efficiently identify the boundaries between variable length AMD64 instructions.

2.5.2 L1 Data Cache

The AMD Family 15h processor contains a 16-Kbyte, 4-way predicted L1 data cache with two 128-bit ports. This is a write-through cache that supports up to two 128 Byte loads per cycle. It is divided into 16 banks, each 16 bytes wide. In addition, the L1 cache is protected from single bit errors through the use of parity. There is a hardware prefetcher that brings data into the L1 data cache to avoid misses. The L1 data cache has a 4-cycle load-to-use latency. Only one load can be performed from a given bank of the L1 cache in a single cycle.

2.5.3 L2 Cache

The AMD Family 15h processor has one shared L2 cache per compute unit. This full-speed on-die L2 cache is mostly inclusive relative to the L1 cache. The L2 is a write-through cache. Every time a store is performed in a core, that address is written into both the L1 data cache of the core the store belongs to and the L2 cache (which is shared between the two cores). The L2 cache has an 18-20 cycle load to use latency.

Size and associativity of the AMD Family 15h processor L2 cache is implementation dependent. See the appropriate BIOS and Kernel Developer's Guide for details.

2.5.4 L3 Cache

The AMD Family 15h processor supports a maximum of 8MB of L3 cache per die, distributed among four L3 sub-caches which can each be up to 2MB in size. The L3 cache is considered a *non-inclusive victim cache* architecture optimized for multi-core AMD processors. Only L2 evictions cause allocations into the L3 cache. Requests that hit in the L3 cache can either leave the data in the L3 cache—if it is likely the data is being accessed by multiple cores—or remove the data from the L3 cache (and place it solely in the L1 cache, creating space for other L2 victim/copy-backs), if it is

likely the data is only being accessed by a single core. Furthermore, the L3 cache of the AMD Family 15h processor also features a number of micro-architectural improvements that enable higher bandwidth.

2.6 Branch-Prediction

To predict and accelerate branches, AMD Family 15h processors employ a combination of next-address logic, a 2-level branch target buffer (BTB) for branch identification and direct target prediction, a return address stack used for predicting return addresses, an indirect target predictor for predicting indirect jump and call addresses, a hybrid branch predictor for predicting conditional branch directions, and a fetch window tracking structure (BSR). Predicted-taken branches incur a 1-cycle bubble in the branch prediction pipeline when they are predicted by the L1 BTB, and a 4-cycle bubble in the case where they are predicted by the L2 BTB. The minimum branch misprediction penalty is 20 cycles in the case of conditional and indirect branches and 15 cycles for unconditional direct branches and returns.

The BTB is a tagged two-level set associative structure accessed using the fetch address of the current window. Each BTB entry includes information about a branch and its target. The L1 BTB contains 128 sets of 4 ways for a total of 512 entries, while the L2 BTB has 1024 sets of 5 ways for a total of 5120 entries.

The hybrid branch predictor is used for predicting conditional branches. It consists of a global predictor, a local predictor and a selector that tracks whether each branch is correlating better with the global or local predictor. The selector and local predictor are indexed with a linear address hash. The global predictor is accessed via a 2-bit address hash and a 12-bit global history.

AMD Family 15h processors implement a separate 512-entry indirect target array used to predict indirect branches with multiple dynamic targets.

In addition, the processors implement a 24-entry return address stack to predict return addresses from a near or far call. Most of the time, as calls are fetched, the next return address is pushed onto the return stack and subsequent returns pop a predicted return address off the top of the stack. However, mispredictions sometimes arise during speculative execution. Mechanisms exist to restore the stack to a consistent state after these mispredictions.

2.7 Instruction Fetch and Decode

AMD Family 15h processors can theoretically fetch 32B of instructions per cycle and send these instructions to the Decode Unit (DE) in 16B windows through the 16-entry (per-thread) Instruction Byte Buffer (IBB). The Decode Unit can only scan two of these 16B windows in a given cycle for up to four instructions. If four instructions partially or wholly exist in more than two of these windows, only those instructions within the first and second windows will be decoded. Aligning to 16B boundaries is important to achieve full decode performance.

2.8 Integer Execution

The integer execution unit for the AMD Family 15h processor consists of two components:

- the integer datapath
- the instruction scheduler and retirement control

These two components are responsible for all integer execution (including address generation) as well as coordination of all instruction retirement and exception handling. The instruction scheduler and retirement control tracks instruction progress from dispatch, issue, execution and eventual retirement.

The scheduling for integer operations is fully data-dependency driven; proceeding out-of-order based on the validity of source operands and the availability of execution resources.

Since the Bulldozer core implements a floating point co-processor model of operation, most scheduling and execution decisions of floating-point operations are handled by the floating point unit. However, the scheduler does track the completion status of all outstanding operations and is the final arbiter for exception processing and recovery.

2.9 Translation-Lookaside Buffer

A *translation-lookaside buffer* (TLB) holds the most-recently-used page mapping information. It assists and accelerates the translation of virtual addresses to physical addresses.

The AMD Family 15h processors utilize a two-level TLB structure.

2.9.1 L1 Instruction TLB Specifications

The AMD Family 15h processor contains a fully-associative L1 instruction TLB with 48 4-Kbyte page entries and 24 2-Mbyte or 1-Gbyte page entries. 4-Mbyte pages require two 2-Mbyte entries; thus, the number of entries available for 4-Mbyte pages is one half the number of 2-Mbyte page entries.

2.9.2 L1 Data TLB Specifications

The AMD Family 15h processor contains a fully-associative L1 data TLB with 32 entries for 4-Kbyte, 2-Mbyte, and 1-Gbyte pages. 4-Mbyte pages require two 2-Mbyte entries; thus, the number of entries available for 4-Mbyte pages is one half the number of 2-Mbyte page entries.

2.9.3 L2 Instruction TLB Specifications

The AMD Family 15 processor contains a 4-way set-associative L2 instruction TLB with 512 4-Kbyte page entries.

2.9.4 L2 Data TLB Specifications

The AMD Family 15h processor contains an L2 data TLB and page walk cache (PWC) with 1024 4-Kbyte, 2-Mbyte or 1-Gbyte page entries (8-way set-associative). 4-Mbyte pages require two 2-Mbyte entries; thus, the number of entries available for 4-Mbyte pages is one half the number of 2-Mbyte page entries.

2.10 Integer Unit

The *integer unit* consists of two components, the *integer scheduler*, which feeds the integer execution pipes, and the *integer execution* unit, which carries out several types of operations discussed below. The integer unit is duplicated for each thread pair.

2.10.1 Integer Scheduler

The scheduler can receive and schedule up to four micro-ops (μ ops) in a dispatch group per cycle. The scheduler tracks operand availability and dependency information as part of its task of issuing μ ops to be executed. It also assures that older μ ops which have been waiting for operands are executed in a timely manner. The scheduler also manages register mapping and renaming.

2.10.2 Integer Execution Unit

There are four integer execution units per core. Two units which handle all arithmetic, logical and shift operations (EX). And two which handle address generation and simple ALU operations (AGLU). Figure 2 shows a block diagram for one integer cluster. There are two such integer clusters per compute unit.




Figure 2. Integer Execution Unit Block Diagram

Macro-ops are broken down into micro-ops in the schedulers. Micro-ops are executed when their operands are available, either from the register file or result buses. Micro-ops from a single operation can execute out-of-order. In addition, a particular integer pipe can execute two micro-ops from different macro-ops (one in the ALU and one in the AGLU) at the same time. (See Figure 1 on page 32.) The scheduler can receive up to four macro-ops per cycle. This group of macro-ops is called a dispatch group.

EX0 contains a variable latency non-pipelined integer divider. EX1 contains a pipelined integer multiplier. The AGLUs contain a simple ALU to execute arithmetic and logical operations and generate effective addresses. A load and store unit (LSU) reads and writes data to and from the L1 data cache. The integer scheduler sends a completion status to the ICU when the outstanding micro-ops for a given macro-op are executed. (For more information on the LSU, see section 2.12 on page 38.)

L1 DTLB has been increased to 64M for AMD Family 15h Models 10h-1fh processors. For AMD Family 15h models 20h to 2fh processors, the L1 DTLB size has increased from 32 entries to 64 entries.

The LZCNT and POPCNT operations are handled in a pipelined unit attached to EX0, as shown in Figure 1 on page 32.

2.11 Floating-Point Unit

The AMD Family 15h processor floating point unit (FPU) was designed to provide four times the raw FADD and FMUL bandwidth as the original AMD Opteron and Athlon 64 processors. It achieves this by means of two 128-bit fused multiply-accumulate (FMAC) units which are supported by a 128-bit high-bandwidth load-store system. The FPU is a coprocessor model that is shared between the two cores of one AMD Family 15h compute unit. As such it contains its own scheduler, register files and renamers and does not share them with the integer units. This decoupling provides optimal performance of both the integer units and the FPU. In addition to the two FMACs, the FPU also contains two 128-bit integer units which perform arithmetic and logical operations on AVX, MMX and SSE packed integer data.

A 128-bit integer multiply accumulate (IMAC) unit is incorporated into FPU pipe 0. The IMAC performs integer fused multiply and accumulate, and similar arithmetic operations on AVX, MMX and SSE data. A crossbar (XBAR) unit is integrated into FPU pipe 1 to execute the permute instruction along with shifts, packs/unpacks and shuffles. There is an FPU load-store unit which supports up to two 128-bit loads and one 128-bit store per cycle.

FPU Features Summary and Specifications:

- The FPU can receive up to four ops per cycle. These ops can only be from one thread, but the thread may change every cycle. Likewise the FPU is four wide, capable of issue, execution and completion of four ops each cycle. Once received by the FPU, ops from multiple threads can be executed.
- Within the FPU, up to two loads per cycle can be accepted, possibly from different threads.
- There are four logical pipes: two FMAC and two packed integer. For example, two 128-bit FMAC and two 128-bit integer ALU ops can be issued and executed per cycle.
- Two 128-bit FMAC units. Each FMAC supports four single precision or two double-precision ops.
- FADDs and FMULs are implemented within the FMAC's.

- x87 FADDs and FMULs are also handled by the FMAC.
- Each FMAC contains a variable latency divide/square root machine.
- Only 1 256-bit operation can issue per cycle, however an extra cycle can be incurred as in the case of a FastPath Double if both micro ops cannot issue together.

Figure 3 shows a block diagram of the dataflow through the FPU.

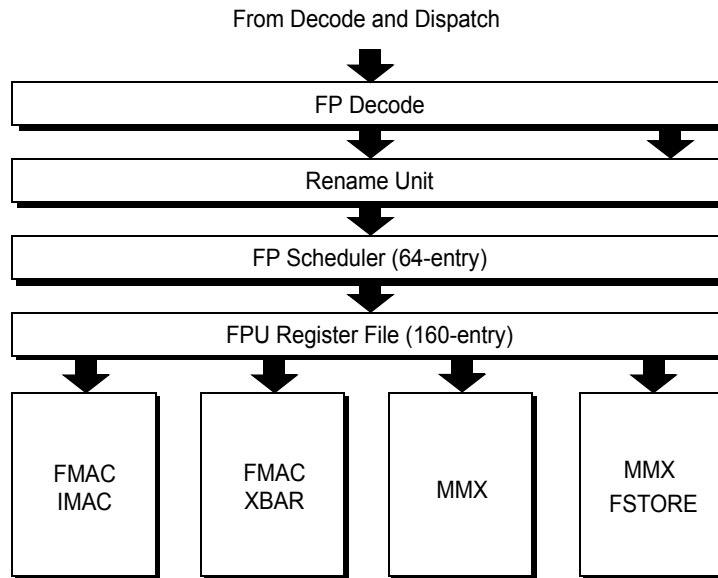


Figure 3. Floating-Point Unit

2.12 Load-Store Unit

The AMD family 15h processor load-store (LS) unit handles data accesses. There are two LS units per compute unit, or one per core. The LS unit supports two 128-bit loads/cycles and one 128-bit store/cycle. There is a 24 entry store queue. This queue buffers stored data until it can be written to the data cache. The load queue has 40 entries and holds load operations until after the load has been completed and delivered to the integer unit or the FPU. The LS unit is composed of two largely independent pipelines enabling the execution of two memory operations per cycle.

Finally, the LS unit helps ensure that the architectural load and store ordering rules are preserved (a requirement for AMD64 architecture compatibility).




Figure 4. Load-Store Unit

2.13 Write Combining

AMD Family 15h processors provide four write-combining data buffers that allow four simultaneous streams. For details, see Appendix A, “Implementation of Write-Combining,” on page 227.

A Write Coalescing Cache (WCC) has been incorporated into the AMD family 15h microarchitecture. The WCC is 4 KB in size and is 4-way set associative. Stores to cacheable memory and, thus, to the L2 cache are coalesced in this cache.

2.14 Integrated Memory Controller

AMD Family 15h processors provide integrated low-latency, high-bandwidth DDR3 memory controllers.

The memory controller supports:

- DRAM chips that are 4, 8, and 16 bits wide within a DIMM.
- Interleaving memory within DIMMs.

- ECC checking with single symbol correcting and double symbol detecting. (See the *BIOS and Kernel Developer's Guide for AMD Family 15h Processors*, order# 42301.)
- Dual-independent 64-bit channel operation.
- Optimized scheduling algorithms and access pattern predictors to improve latency and achieved bandwidth, particularly for interleaved streams of read and write DRAM accesses.
- A data prefetcher.

Prefetched data is held in the memory controller itself and is not speculatively filled into the L1, L2, or L3 caches. This prefetcher is able to capture both positive and negative stride values (both unit and non-unit) of cache-line size, as well as some more complicated access patterns.

For specifications on a certain processor's memory controller, see the data sheet for that processor. For information on how to program the memory controller, see the *BIOS and Kernel Developer's Guide for AMD Family 15h Processors*, order# 42301.

2.15 HyperTransport™ Technology Interface

HyperTransport technology is a scalable, high-speed, low-latency, point-to-point, packetized link that:

- Enables high data transfer rates.
- Simplifies connectivity by replacing legacy buses and bridges.
- Reduces latencies and bottlenecks within systems.

When compared with traditional technologies, HyperTransport technology allows much faster data-transfer rates. For more information on HyperTransport technology, see the *HyperTransport I/O Link Specification*, available at www.hypertransport.org.

On AMD Family 15h processors, HyperTransport technology provides the link to I/O devices. Some processor models—for example, those designed for use in multiprocessing systems—also utilize HyperTransport technology to connect to other processors. See the *BIOS and Kernel Developer's Guide* for your particular processor for details concerning HyperTransport technology implementation details.

In addition to supporting previous HyperTransport interfaces, AMD Family 15h processors support a newer version of the HyperTransport standard: HyperTransport3. HyperTransport3 increases the aggregate link bandwidth to a maximum of 25.6 Gbyte/s (16-bit link). HyperTransport3 also adds HyperTransport Retry which improves RAS by allowing detection and retransmission of packets corrupted in transit.

Additional features in the AMD Family 15h HyperTransport implementation may include:

- HyperTransport link bandwidth balancing, allowing multiple HyperTransport links to be teamed to carry coherent traffic.

- HyperTransport Link Splitting, which allowing a single 16-bit link to be split into two 8-bit links.

These features allow for further optimized platform designs that are capable of increasing system bandwidth and reducing latency.

2.15.1 HyperTransport Assist

Multisocket-capable AMD family 15h processors incorporate HyperTransport assist technology (also referred to in some documents as probe filtering). HyperTransport assist functionality may or may not be enabled by default on a specific platform implementation. However, the BIOS can enable HyperTransport assist on these platforms, if it is not enabled by default.

HyperTransport assist reduces the effective latency of memory access in multi-node systems by changing the coherence protocol from a broadcast style to a directory style. In cases requiring many memory accesses, in particular to local memory—as is common in NUMA-optimized applications—the probe and response latency required for maintaining coherence takes longer than the DRAM access. In these cases, HyperTransport assist can remove many of the probes and responses associated with DRAM access, thus decreasing effective latency to access system memory.

HyperTransport assist also increases the total coherent fabric bandwidth capability within the system by removing much probe and response traffic from the coherent HyperTransport links. It also streamlines probe and response handling throughout the L1/L2/L3 caches and elsewhere in the microarchitecture, which can lead to additional bandwidth improvements in systems with multiple processing nodes.

HyperTransport assist is enabled by partitioning the L3 cache physical storage into a section used as traditional (CPU-side) L3 cache, and a separate physical section for directory storage which is inaccessible to the CPUs. In effect, from the perspective of CPUs, systems with HyperTransport assist enabled have a smaller L3 cache. Typically, 1–2MB of L3 cache is reserved for use by HyperTransport assist technology. Thus, some amount of L3 capacity is traded for reduced latency on cache refills. While the benefit of this tradeoff can be workload-dependant, it is almost universally a win on larger (4+ node) systems. If a platform runs a specific workload, it may be worth evaluating performance with and without HyperTransport assist.

Enabling HyperTransport assist allows for performance benefits beyond the NUMA optimizations suggested in Appendix C, “Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems,” on page 313.

For information on HyperTransport assist implementation details for a specific processor, see the BIOS and Kernel Developer’s Guide for AMD Family 15h Processors, order #42301 (where it goes by the alternate name of probe filter).

Chapter 3 C and C++ Source-Level Optimizations

Although C and C++ compilers can often produce very efficient object code from naive source code, careful attention to coding details can lead to even better object code and therefore to improved performance. Many optimizations take advantage of the underlying mechanisms used by C and C++ compilers to translate source code into sequences of AMD64 instructions. This chapter includes guidelines for writing C and C++ source code that yields an approximation to the most highly efficient optimization.

This chapter covers the following topics:

Topic	Page
Declarations of Floating-Point Values	44
Using Arrays and Pointers	45
Use of Function Prototypes	47
Unrolling Small Loops	47
Expression Order in Compound Branch Conditions	48
Arrange Boolean Operands for Quick Expression Evaluation	49
Long Logical Expressions in If Statements	50
Pointer Alignment	51
Unnecessary Store-to-Load Dependencies	52
Matching Store and Load Size	53
Use of const Type Qualifier	56
Generic Loop Hoisting	56
Local Static Functions	59
Explicit Parallelism in Code	59
Extracting Common Subexpressions	62
Sorting and Padding C and C++ Structures	63
Replacing Integer Division with Multiplication	64
Frequently Dereferenced Pointer Arguments	65
32-Bit Integral Data Types	66
Sign of Integer Operands	67
Improving Performance in Linux® Libraries	68
Improving Performance in Linux® Libraries	68
Aligning Matrices	69

Additional Considerations

Source-code transformations interact with a compiler's code generator, making it difficult to control the generated machine code from the source level. It is even possible that source-code transformations aimed at improving performance may conflict with compiler optimizations. Depending on the compiler and the specific source code, it is possible for pointer-style code to compile into machine code that is faster than that generated from equivalent array-style code. Compare the performance of your code after implementing a source-code transformation with the performance of the original code to be sure that there is an improvement.

Some compilers provide proprietary declaration keywords that further allow the compiler to reduce possible aliasing. See *Compiler Usage Guidelines for 64-Bit Operating Systems on AMD64 Platforms Application Note*, order# 32035, for details.

3.1 Declarations of Floating-Point Values

Optimization

When working with single precision (`float`) values:

- Use the `f` or `F` suffix (for example, `3.14f`) to specify a constant value of type `float`.
- Use function prototypes for all functions that accept arguments of type `float`.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

C and C++ compilers treat floating-point constants and arguments as double precision (`double`) unless you specify otherwise. However, single precision floating-point values occupy half the memory space as double precision values and can often provide the precision necessary for a given computational problem.

This optimization also results in more efficient use of the XMM Streaming SIMD registers: four single precision values can be packed into a single XMM register, compared to two double precision values.

3.2 Using Arrays and Pointers

Optimization

Use array notation instead of pointer notation when working with arrays.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

C allows the use of either the array operator (`[]`) or pointers to access the elements of an array. However, the use of pointers in C makes work difficult for optimizers in C compilers. Without detailed and aggressive pointer analysis, the compiler has to assume that writes through a pointer can write to any location in memory, including storage allocated to other variables. (For example, `*p` and `*q` can refer to the same memory location, while `x[0]` and `x[2]` cannot.) Pointers make it difficult for compilers to detect the presence or absence of aliasing—with possible ambiguous access to a block of memory. The compiler sometimes must assume aliasing in the presence of pointers, which limits the opportunities for optimization. Array notation makes the task of the optimizer easier by reducing possible aliasing.

Example

Avoid code, such as the following, which uses pointer notation:

```
typedef struct {
    float x, y, z, w;
} VERTEX;

typedef struct {
    float m[4][4];
} MATRIX;

void XForm(float *res, const float *v, const float *m, int numverts) {

    float dp;
    int i;
    const VERTEX* vv = (VERTEX *)v;

    for (i = 0; i < numverts; i++) {
        dp = vv->x * *m++;
        dp += vv->y * *m++;
        dp += vv->z * *m++;
        dp += vv->w * *m++;
    }
}
```

```

    *res++ = dp; // Write transformed x.

    dp = vv->x * *m++;
    dp += vv->y * *m++;
    dp += vv->z * *m++;
    dp += vv->w * *m++;

    *res++ = dp; // Write transformed y.

    dp = vv->x * *m++;
    dp += vv->y * *m++;
    dp += vv->z * *m++;
    dp += vv->w * *m++;

    *res++ = dp; // Write transformed z.

    dp = vv->x * *m++;
    dp += vv->y * *m++;
    dp += vv->z * *m++;
    dp += vv->w * *m++;

    *res++ = dp; // Write transformed w.

    ++vv; // Next input vertex
    m -= 16; // Reset to start of transform matrix.
}
}

```

Instead, use the equivalent array notation:

```

typedef struct {
    float x, y, z, w;
} VERTEX;

typedef struct {
    float m[4][4];
} MATRIX;

void XForm(float *res, const float *v, const float *m, int numverts) {
    int i;
    const VERTEX* vv = (VERTEX *)v;
    const MATRIX* mm = (MATRIX *)m;
    VERTEX* rr = (VERTEX *)res;
    for (i = 0; i < numverts; i++) {
        rr[i].x = vv[i].x * mm->m[0][0] + vv[i].y * mm->m[0][1] +
            vv[i].z * mm->m[0][2] + vv[i].w * mm->m[0][3];
        rr[i].y = vv[i].x * mm->m[1][0] + vv[i].y * mm->m[1][1] +
            vv[i].z * mm->m[1][2] + vv[i].w * mm->m[1][3];
        rr[i].z = vv[i].x * mm->m[2][0] + vv[i].y * mm->m[2][1] +
            vv[i].z * mm->m[2][2] + vv[i].w * mm->m[2][3];
        rr[i].w = vv[i].x * mm->m[3][0] + vv[i].y * mm->m[3][1] +
            vv[i].z * mm->m[3][2] + vv[i].w * mm->m[3][3];
    }
}

```

3.3 Use of Function Prototypes

Optimization

In general, use prototypes for all functions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Function prototypes can convey additional information to the compiler that might enable more aggressive optimizations and enable the compiler to catch potential runtime errors.

3.4 Unrolling Small Loops

Optimization

Completely unroll loops that have a small fixed loop count and a small loop body.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Many compilers do not aggressively unroll loops. Manually unrolling loops can benefit performance, especially if the loop body is small, making the loop overhead significant.

Unrolling loops increases the code size, which may decrease performance in rare cases.

Example

Avoid a small loop like this:

```
// 3D-transform: Multiply vector V by 4x4 transform matrix M.
for (i = 0; i < 4; i++) {
    r[i] = 0;
    for (j = 0; j < 4; j++) {
```

```
    r[i] += m[j][i] * v[j];  
  }  
}
```

Instead, replace it with its completely unrolled equivalent, as shown here:

```
r[0] = m[0][0] * v[0] + m[1][0] * v[1] + m[2][0] * v[2] + m[3][0] * v[3];  
r[1] = m[0][1] * v[0] + m[1][1] * v[1] + m[2][1] * v[2] + m[3][1] * v[3];  
r[2] = m[0][2] * v[0] + m[1][2] * v[1] + m[2][2] * v[2] + m[3][2] * v[3];  
r[3] = m[0][3] * v[0] + m[1][3] * v[1] + m[2][3] * v[2] + m[3][3] * v[3];
```

Related Information

For information on loop unrolling at the assembly-language level, see “Loop Unrolling” on page 133.

3.5 Expression Order in Compound Branch Conditions

Optimization

In the most active areas of a program, order the expressions in compound branch conditions to take advantage of short circuiting of compound conditional expressions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Branch conditions in C programs often consist of compound conditions consisting of multiple boolean expressions joined by the logical AND (&&) or logical OR (||) operators. C compilers guarantee short-circuit evaluation of these operators. In a compound logical OR expression, the first operand to evaluate to true terminates the evaluation, and subsequent operands are not evaluated at all. Similarly, in a logical AND expression, the first operand to evaluate to false terminates the evaluation. Hence, it is not always possible to swap the operands of logical OR and logical AND. This is especially true when the evaluation of one of the operands causes a side effect. In most cases the order of operands in such expressions is irrelevant.

When used to control conditional branches, expressions involving logical OR or logical AND are translated into a series of conditional branches. The ordering of the conditional branches is a function of the ordering of the expressions in the compound condition and can have a significant impact on performance. It is impossible to give an easy, closed-form formula on how to order the conditions. Overall performance is a function of the following factors:

- Probability of a branch misprediction for each of the branches generated
- Additional latency incurred due to a branch misprediction
- Cost of evaluating the conditions controlling each of the branches generated
- Amount of parallelism that can be extracted in evaluating the branch conditions
- Data stream consumed by an application (mostly due to the dependence of misprediction probabilities on the nature of the incoming data in data-dependent branches)

It is recommended to experiment with the ordering of expressions in compound branch conditions in the most active areas of a program (“hot spots,” consuming a great amount of execution time). Such hot spots can be found through the use of profiling by feeding a typical data stream to the program while doing the experiments.

3.6 Arrange Boolean Operands for Quick Expression Evaluation

Optimization

In expressions that use the logical AND (&&) or logical OR (||) operator, arrange the operands for quick evaluation of the expression:

If the expression uses this operator	Then arrange the operands from left to right in decreasing probability of being
&& (logical AND)	False
(logical OR)	True

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

C and C++ compilers guarantee short-circuit evaluation of the boolean operators && and ||. In an expression that uses &&, the first operand to evaluate to false terminates the evaluation; subsequent operands are not evaluated. In an expression that uses ||, the first operand to evaluate to true terminates the evaluation.

When used to control program flow, expressions involving && and || are translated into a series of conditional branches. This optimization minimizes the total number of conditions evaluated and branches executed.

Example 1

In the following code, the operands of `&&` are not arranged for quick expression evaluation because the first operand is not the condition case most likely to be false (it is far less likely for an animal name to begin with a ‘y’ than for it to have fewer than four characters):

```
char animalname[30];
char *p;

p = animalname;

if ((strlen(p) > 4) && (*p == 'y')) { ... }
```

Because the odds that the animal name begins with a ‘y’ are comparatively low, it is better to put that operand first:

```
if ((*p == 'y') && (strlen(p) > 4)) { ... }
```

Example 2

In the following code (assuming a uniform random distribution of `i`), the operands of `||` are not arranged for quick expression evaluation because the first operand is not the condition most likely to be true:

```
unsigned int i;

if ((i < 4) || (i & 1)) { ... }
```

Because it is more likely for the least-significant bit of `i` to be 1, it is better to put that operand first:

```
if ((i & 1) || (i < 4)) { ... }
```

3.7 Long Logical Expressions in If Statements

Optimization

In `if` statements, avoid long logical expressions that can generate dense conditional branches that violate the guideline described in “Instruction Fetch” on page 123. When long logical expressions are unavoidable, try to arrange them so that most of the implicit branches are not be taken.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

High branch density can lead to some branches not being identified by the branch predictor (as described in section “Instruction Fetch” on page 123). If these unpredicted branches are taken, they will incur a misprediction penalty.

Preferred for Data that Falls Mostly Within the Range

```
if (a <= max && a >= min && b <= max && b >= min)
```

If most of the data falls within the range, the branches will not be taken, so the above code is preferred. Otherwise, the following code is preferred.

Preferred for Data that Does Not Fall Mostly Within the Range

```
if (a > max || a < min || b > max || b < min)
```

3.8 Pointer Alignment

Optimization

Use the following technique to align a pointer to a 16-byte boundary in situations where alignment cannot be assured.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Examples

Dynamic memory allocation—accomplished through the use of the `malloc` library function in C—should always return a pointer that is suitably aligned for the largest base type (16-byte alignment). However, this may not always be the case. In this example, after memory allocation, use `np` instead of `p` to access the data. The pointer `p` is still needed in order to deallocate the storage later.

```
double *p;  
double *np;  
  
p = (double *)malloc(sizeof(double) * number_of_doubles + 15);  
np = (double *)(((ptrdiff_t)(p)) + 15L) & (-16L);
```

3.9 Unnecessary Store-to-Load Dependencies

Optimization

Avoid store-to-load dependencies.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

A store-to-load dependency exists when data is stored to memory, only to be read back shortly thereafter. For details, see “Store-to-Load Forwarding Restrictions” on page 98. The AMD Family 15h processor contains hardware to accelerate such store-to-load dependencies, allowing the load to obtain the store data before it has been written to memory. However, avoiding such dependencies and keeping the data in an internal register results in faster code.

It is especially important to avoid store-to-load dependencies if they are part of a long dependency chain, as may occur in a recurrence computation. If the dependency occurs while operating on arrays, many compilers are unable to optimize the code in a way that avoids the store-to-load dependency. In some instances the language definition may prohibit the compiler from using code transformations that would remove the store-to-load dependency. Therefore, it is recommended that the programmer remove the dependency manually, for example, by introducing a temporary variable that can be kept in a register. This can result in a significant performance increase.

Examples

Avoid

```
// In the following loops, each iteration writes to the memory location
// referenced by x[k], but then reads from this same location in the
// subsequent iteration, where it becomes x[k-1]. This creates a store-to-load
// dependency.
```

```
double x[VECLEN], y[VECLEN], z[VECLEN];
unsigned int k;

for (k = 1; k < VECLLEN; k++) {

x[k] = x[k-1] + y[k]; // x[k] is written to in iteration k and then
                    // read in iteration k+1

}

for (k = 1; k < VECLLEN; k++) {
```

```
x[k] = z[k] * (y[k] - x[k-1]); //x[k] is written to in iteration k and then read
in iteration k+1
}
```

Preferred

```
double x[VECLEN], y[VECLEN], z[VECLEN];
unsigned int k;
double t;
t = x[0];

// By using the temporary variable t to store the partial accumulated
// computations
// from each iteration, the restructured loops below are able to avoid
// having to read
// the memory location referenced by x[k] immediately after writing
// to this same location in the previous iteration.

for (k = 1; k < VECLLEN; k++) {

t = t + y[k];

x[k] = t; // x[k] is written to in iteration k, but not read in iteration k+1

}

t = x[0];
for (k = 1; k < VECLLEN; k++) {
t = z[k] * (y[k] - t);
x[k] = t; //x[k] is written to in iteration k, but not read in iteration k+1
}
```

3.10 Matching Store and Load Size

Optimization

Align memory accesses and match addresses and sizes of stores and dependent loads.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The AMD Family 15h processor contains a load-store buffer to speed up the forwarding of store data to dependent loads. However, this store-to-load forwarding (STLF) inside the load-store buffer

occurs, in general, only when the addresses and sizes of the store and the dependent load match, and when both memory accesses are aligned.

For details, see “Store-to-Load Forwarding Restrictions” on page 98.

It is impossible to control load and store activity at the source level in such a way as to avoid all cases that violate restrictions placed on store-to-load-forwarding. In some instances it is possible to spot such cases in the source code. Size mismatches can easily occur when different-size data items are joined in a union. Address mismatches could be the result of pointer manipulation.

The following examples show a situation involving a union of different-size data items. The examples show a user-defined unsigned 16.16 fixed-point type and two operations defined on this type. Function `fixed_add` adds two fixed-point numbers, and function `fixed_int` extracts the integer portion of a fixed-point number. Listing shows an inappropriate implementation of `fixed_int`, which, when used on the result of `fixed_add`, causes misalignment, address mismatch, or size mismatch between memory operands, such that no store-to-load forwarding in the load-store buffer takes place. The following examples shows how to properly implement `fixed_int` in order to allow store-to-load forwarding in the load-store buffer.

Examples

Avoid

```
typedef union {
    unsigned int whole;
    struct {
        unsigned short frac; /* Lower 16 bits are fraction. */
        unsigned short intg; /* Upper 16 bits are integer. */
    } parts;
} FIXED_U_16_16;

__inline FIXED_U_16_16 fixed_add(FIXED_U_16_16 x, FIXED_U_16_16 y) {
    FIXED_U_16_16 z;
    z.whole = x.whole + y.whole;
    return (z);
}

__inline unsigned int fixed_int(FIXED_U_16_16 x) {
    return((unsigned int)(x.parts.intg));
}
...
FIXED_U_16_16 y, z;
unsigned int q;
...
label1:
y = fixed_add (y, z);
q = fixed_int (y);

label2:
...
```

The object code generated for the source code between `label1` and `label2` typically follows one of these two variants:

```
; Variant 1
mov edx, DWORD PTR [z]
mov eax, DWORD PTR [y]      ;
add eax, edx                ;
mov DWORD PTR [y], eax      ; -+
mov EAX, DWORD PTR [y+2]    ; <+ Address mismatch--no forwarding in LSU
and EAX, 0FFFFh
mov DWORD PTR [q], eax

; Variant 2
mov  edx, DWORD PTR [z]
mov  eax, DWORD PTR [y]      ;
add  eax, edx                ;
mov  DWORD PTR [y], eax      ; -+
movzx eax, WORD PTR [y+2]    ; <+ Size and address mismatch--no forwarding in LSU
mov  DWORD PTR [q], eax
```

Some more sophisticated compilers may generate optimal machine code even for the previous example. These compilers provide various optional levels and types of optimizations that are controlled by compiler program flags. When compiled at a moderate level of optimization, such compilers may generate perfectly acceptable code from C++ code such as that listed above. For more information, see *Compiler Usage Guidelines for 64-Bit Operating Systems on AMD64 Platforms Application Note*, order# 32035.

Preferred

```
typedef union {
    unsigned int whole;
    struct {
        unsigned short frac; /* Lower 16 bits are fraction. */
        unsigned short intg; /* Upper 16 bits are integer. */
    } parts;
} FIXED_U_16_16;

__inline FIXED_U_16_16 fixed_add(FIXED_U_16_16 x, FIXED_U_16_16 y) {
    FIXED_U_16_16 z;
    z.whole = x.whole + y.whole;
    return(z);
}

__inline unsigned int fixed_int(FIXED_U_16_16 x) {
    return (x.whole >> 16);
}
...
FIXED_U_16_16 y, z;
unsigned int q;
...
label1:
y = fixed_add (y, z);
q = fixed_int (y);
```

```
label2:  
...
```

The object code generated for the source code between `label1` and `label2` typically looks like this:

```
mov edx, DWORD PTR [z]  
mov eax, DWORD PTR [y]  
add eax, edx  
mov DWORD PTR [y], eax ; -+  
mov eax, DWORD PTR [y] ; <+ Aligned (size/address match)--forwarding in LSU  
shr eax, 16  
mov DWORD PTR [q], eax
```

3.11 Use of const Type Qualifier

Optimization

For objects whose values will not be changed, use the `const` type qualifier.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Using the `const` type qualifier makes code more robust and may enable the compiler to generate higher-performance code. For example, under the C standard, a compiler is not required to allocate storage for an object that is declared `const`, if its address is never used.

3.12 Generic Loop Hoisting

Optimization

To improve the performance of inner loops, reduce redundant constant calculations (that is, loop-invariant calculations). This idea can also be extended to invariant control structures.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale and Examples

The following example demonstrates the use of an invariant condition in an `if` statement in a `for` loop. The second listing shows the preferred optimization.

Avoid

```
for (i...) {
    if (CONSTANT0) {
        DoWork0(i);    // Does not affect CONSTANT0.
    }
    else {
        DoWork1(i);    // Does not affect CONSTANT0.
    }
}
```

Preferred

```
if (CONSTANT0) {
    for (i...) {
        DoWork0(i);
    }
}
else {
    for (i...) {
        DoWork1(i);
    }
}
```

The preferred optimization in the preceding example tightens the inner loops by avoiding repetitious evaluation of a known `if` control structure. Although the branch would be easily predicted, the extra instructions and decode limitations imposed by branching are eliminated.

To generalize the preceding example further for multiple-constant control code, more work may be needed to create the proper outer loop. Enumeration of the constant cases reduces this to a simple `switch` statement.

Avoid

```
for (i...) {
    if (CONSTANT0) {
        DoWork0(i);    // Does not affect CONSTANT0 or CONSTANT1.
    }
    else {
        DoWork1(i);    // Does not affect CONSTANT0 or CONSTANT1.
    }

    if (CONSTANT1) {
        DoWork2(i);    // Does not affect CONSTANT0 or CONSTANT1.
    }
    else {
        DoWork3(i);    // Does not affect CONSTANT0 or CONSTANT1.
    }
}
```

```
}
```

Transform the loop in the preceding example (by using the `switch` statement) into:

Preferred

```
#define combine(c1, c2) (((c1) << 1) + (c2))
switch (combine(CONSTANT0 != 0, CONSTANT1 != 0)) {
    case combine(0, 0):
        for(i...) {
            DoWork0(i);
            DoWork2(i);
        }
        break;
    case combine(1, 0):
        for(i...) {
            DoWork1(i);
            DoWork2(i);
        }
        break;
    case combine(0, 1):
        for(i...) {
            DoWork0(i);
            DoWork3(i);
        }
        break;
    case combine( 1, 1 ):
        for(i...) {
            DoWork1(i);
            DoWork3(i);
        }
        break;
    default:
        break;
}
```

Some introductory code is necessary to generate all the combinations for the `switch` constant and the total amount of code has doubled. However, the inner loops are now free of `if` statements. In ideal cases where the `DoWorkn` functions are inlined, the successive functions have greater overlap, leading to greater parallelism than possible in the presence of intervening `if` statements.

The same idea can be applied to constant `switch` statements or to combinations of `switch` statements and `if` statements inside of `for` loops. The method used to combine the input constants becomes more complicated but benefits performance.

However, the number of inner loops can also substantially increase. If the number of inner loops is prohibitively high, then only the most common cases must be dealt with directly, and the remaining cases can fall back to the old code in the default clause of the `switch` statement. This situation is typical of run-time generated code. While the performance of run-time generated code can be improved by means similar to those presented here, it is much harder to maintain and developers must do their own code-generation optimizations without the help of an available compiler.

3.13 Local Static Functions

Optimization

Declare as `static` functions that are not used outside the file where they are defined.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Declaring a function as `static` forces internal linkage. Functions that are not declared as `static` default to external linkage, which may inhibit certain optimizations—for example, aggressive inlining—with some compilers. In C++, programmers can declare functions inside an anonymous namespace to achieve the same local scoping effect.

3.14 Explicit Parallelism in Code

Optimization

Where possible, break long dependency chains into several independent chains that can be executed in parallel to take advantage of the execution units in each pipeline.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Breaking long dependency chains into smaller parallelizable chains can improve performance. This is especially important for floating point code, due to the longer latency of floating point operations.

It should be remarked, however, that the reordered code may not produce identical computational results, since floating point operations are not associative. For this reason, most languages (including ANSI C) are bound by the requirement that floating-point expressions cannot be reordered, and compilers in those languages cannot reorder these expressions unless they support a flag for non-compliant reordering. In some cases, reordered floating-point code may lead to unexpected results,

but in most cases, the final result differs only in the least-significant bits. There are well-known numerical considerations in applying this optimization. For details, consult a reference on numerical analysis.

Example

When deciding how aggressively to unroll a loop, a key consideration is how many instructions need to be in-flight to keep all the pipelines and execution units busy in every iteration. For handling floating point code, there are two identical FP pipelines, each having a 5 cycle latency for the execution phase. This means that a total of ten floating point instructions need to be in flight in every loop iteration to keep both FP pipelines and the 5-stage FP adder busy on every cycle. Below is an example that shows how a loop can be unrolled to fully utilize the FP resources on every cycle. Notice that the assembly code shown for the unrolled loop body involves ten floating point additions in each iteration. This will keep both FP pipelines fully occupied at every cycle.

Original Loop (Avoid)

```
double a[100], sum;
int i;

sum = 0.0f;
for (i = 0; i < 100; i++) {
    sum += a[i];
}
```

Unrolled Loop Body Source (Preferred)

```

for (i = 0; i < 100; i += 10) {
sum1 += a[i];
sum2 += a[i+1];
sum3 += a[i+2];
sum4 += a[i+3];
sum5 += a[i+4];
sum6 += a[i+5];
sum7 += a[i+6];
sum8 += a[i+7];
sum9 += a[i+8];
sum10 += a[i+9];
}

```

Unrolled Loop Body (AVX Code)

```

.code
public loopunrollavx

loopunrollavx proc
    xor r11,r11
    xorpd xmm0,xmm0 ;xmm0 = 0
    vmovsd xmm0,QWORD PTR [rdx] ;xmm0 = sum
label1:

    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8] ;xmm0+=a[i]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+8] ;xmm0+=a[i+1]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+16] ;xmm0+=a[i+2]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+24] ;xmm0+=a[i+3]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+32] ;xmm0+=a[i+4]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+40] ;xmm0+=a[i+5]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+48] ;xmm0+=a[i+6]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+56] ;xmm0+=a[i+7]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+64] ;xmm0+=a[i+8]
    vaddsd xmm0,xmm0,QWORD PTR [r8+r11*8+72] ;xmm0+=a[i+9]
    add r11,10
    cmp r11,rcx ;i < counter?
    jl label1
    vmovntpd XMMWORD PTR [rdx],xmm0 ;*sum = XMM0
    ret
loopunrollavx endp

end

```

3.15 Extracting Common Subexpressions

Optimization

Manually extract common subexpressions from floating-point expressions, where C compilers may be unable to extract them due to the rules against reordering of floating-point expressions in the ANSI standard.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Specifically, the compiler cannot rearrange the computation according to algebraic equivalencies before extracting common subexpressions. Rearranging the expression may give different computational results due to the lack of associativity of floating-point operations, but the results usually differ in only the least-significant bits. However, since errors in the least significant bits can be magnified by later operations to the extent that they completely invalidate the calculation, the programmer should proceed with caution when implementing this sort of computation.

Examples

Avoid

```
double a, b, c, d, e, f;  
  
e = b * c / d;  
f = b / d * a;
```

Preferred

```
double a, b, c, d, e, f, t;  
  
t = b / d;  
e = c * t;  
f = a * t;
```

Avoid

```
double a, b, c, e, f;  
  
e = a / c;  
f = b / c;
```

Preferred

```
double a, b, c, e, f, t;
```

```
t = 1 / c;
```

```
e = a * t
```

```
f = b * t;
```

3.16 Sorting and Padding C and C++ Structures

Optimization

Sort and pad C and C++ structures to achieve natural alignment.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

In order to achieve better alignment for structures, many compilers have options that allow padding of structures to make their sizes multiples of words, doublewords, or quadwords. In addition, to improve the alignment of structure members, some compilers may allocate structure elements in an order that differs from the order in which they are declared. Unfortunately, some compilers may not offer any of these features, or their implementations might not work properly in all situations.

By sorting and padding structures at the source-code level, if the first member of a structure is naturally aligned, then all other members are naturally aligned as well. This allows, for example, arrays of structures to be perfectly aligned.

Sorting and Padding C and C++ Structures

To sort and pad a C or C++ structure, follow these steps:

1. Sort the structure members according to their type sizes, declaring members with larger type sizes ahead of members with smaller type sizes.
2. Pad the structure so the size of the structure is a multiple of the largest member's type size.

Examples

Avoid structure declarations in which the members are not declared in order of their type sizes and the size of the structure is not a multiple of the size of the largest member's type:

```
struct {
    char a[5];    \\ Smallest type size (1 byte * 5)
    long k;      \\ 4 bytes in this example
    double x;    \\ Largest type size (8 bytes)
} baz;
```

Instead, declare the members according to their type sizes (largest to smallest) and add padding to ensure that the size of the structure is a multiple of the largest member's type size:

```
struct {
    double x;    \\ Largest type size (8 bytes)
    long k;      \\ 4 bytes in this example
    char a[5];   \\ Smallest type size (1 byte * 5)
    char pad[7]; \\ Make structure size a multiple of 8.
} baz;
```

3.17 Replacing Integer Division with Multiplication

Optimization

Replace integer division with multiplication when there are multiple divisions in an expression. (This is possible only if no overflow will occur during the computation of the product. The possibility of an overflow can be determined by considering the possible ranges of the divisors.)

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Integer division is the slowest of all integer arithmetic operations.

Examples

Avoid code that uses two integer divisions:

```
int i, j, k, m;  
m = i / j / k;
```

Instead, replace one of the integer divisions with the appropriate multiplication:

```
m = i / (j * k);
```

3.18 Frequently Dereferenced Pointer Arguments

Optimization

Avoid dereferenced pointer arguments inside a function.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Because the compiler has no knowledge of aliasing between pointers, such dereferencing cannot be “optimized away.” Since data may not be maintained in registers, memory traffic can significantly increase.

Many compilers have an “assume no aliasing” optimization switch. This allows the compiler to assume that two different pointers always have disjoint contents and does not require copying of pointer arguments to local variables. If your compiler does not have this type of optimization, then copy the data referenced by the pointer arguments to local variables at the start of the function and if necessary copy them back at the end of the function. (Some compilers also provide keywords to provide the same aliasing information to the compiler. For details, see *Compiler Usage Guidelines for 64-Bit Operating Systems on AMD64 Platforms Application Note*, order# 32035.)

Examples

Avoid

```
// Assumes pointers are different and q != r.  
void isqrt(unsigned long a, unsigned long *q, unsigned long *r) {  
  
    *q = a;
```

```
if (a > 0) {
    while (*q > (*r = a / *q)) {
        *q = (*q + *r) >> 1;
    }
}
*r = a - *q * *q;
}
```

Preferred

```
// Assumes pointers are different and q != r.
void isqrt(unsigned long a, unsigned long *q, unsigned long *r) {

    unsigned long qq, rr;
    qq = a;
    if (a > 0) {
        while (qq > (rr = a / qq)) {
            qq = (qq + rr) >> 1;
        }
    }
    rr = a - qq * qq;
    *q = qq;
    *r = rr;
}
```

3.19 32-Bit Integral Data Types

Optimization

Use 32-bit integers instead of smaller sized integers (16-bit or 8-bit).

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

When choosing between 32-bit, 16-bit and 8-bit data types in cases where memory footprint is not a concern, using 32-bit integer types in 32-bit software (32-bit or 64-bit integer types in 64-bit software) avoids possible register-merging false dependencies due to partial register writes. See section 5.5, "Partial-Register Writes" on page 84 for details.

3.20 Sign of Integer Operands

Optimization

Where there is a choice of using either a signed or an unsigned type, take into consideration that some operations are faster with unsigned types while others are faster for signed types.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

In many cases, the type of data to be stored in an integer variable determines whether a signed or an unsigned integer type is appropriate. For example, to record the weight of a person in pounds, no negative numbers are required, so an unsigned type is appropriate. However, recording temperatures in degrees Celsius may require both positive and negative numbers, so a signed type is needed.

Integer-to-floating-point conversion using integers larger than 16 bits is faster with signed types, as the AMD64 architecture provides instructions for converting signed integers to floating-point but has no instructions for converting unsigned integers. In a typical case, a 32-bit integer is converted by a compiler to assembly as follows:

Example

Computing quotients and remainders in integer division by constants is faster when performed on unsigned types. The following typical case is the compiler output for a 32-bit integer divided by 4:

Avoid

```
int i;          =====>  mov eax, i
                  cdq
i = i / 4;      and edx, 3
                  add eax, edx
                  sar eax, 2
                  mov i,  eax
```

Preferred

```
unsigned int i; =====>  shr i, 2

i = i / 4;
```

In summary, use unsigned types for:

- Division and remainders
- Loop counters
- Array indexing

Use signed types for:

- Integer-to-floating-point conversion

3.21 Improving Performance in Linux[®] Libraries

Optimization

If symbol interposition is not important to a particular application, then you should control the visibility of the symbols in a shared object in such a way as to optimize internal references to other symbols in the library and minimize the symbol export table size.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Dynamically loadable libraries are a versatile feature of the Linux[®] operating system. These allow one or more symbols in one library to override an identical symbol in another library. Known as interposition, this ability makes customizations and probing seamless. Interposition is implemented by means of a procedure linkage table (PLT). The PLT is so flexible that even references to an overridden symbol inside its own library end up referencing the overriding symbol. However, the PLT imposes a performance penalty by requiring all function calls to public global routines to go through an extra step that increases the chances of cache misses and branch mispredictions. This is particularly severe for C++ classes whose methods refer to other methods in the same class.

When using `ld` to link a shared object, include the command line option `-Bsymbolic`.

If using a version of `gcc` prior to 4.0 to link a shared object, add the option `-Wl,-Bsymbolic` to the command-line. If using `gcc` 4.0 or later, add the option `-fvisibility=protected` to the command-line.

If finer control is desired, then it is possible to specify `-fvisibility=hidden` to `gcc` 4.0 or later and then add `__attribute__((visibility("default")))` to each symbol that should be exported. When building C++ shared objects, also consider using the `-fvisibility-inlines-hidden` option.

3.22 Aligning Matrices

Optimization

When using multi-dimensional arrays or matrices, make sure that each row or 2nd-order dimension starts at a 16-byte boundary.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Instead of creating matrices with arbitrary dimensions, make sure that the size in bytes of the low-order dimension is a multiple of 16 and that it starts at a 16-byte boundary. By doing so, when iterating over the elements of the matrix the compiler is presented with data properly aligned for low-cost vectorization.

For example, in:

```
double a [10][11],
        b [10][11];
int i, j;

for (j = 0; j < 10; j++)
    for (i = 0; i < 11; i++)
        b [j][i] = a [j][i] * M_1_PI;
```

Declare the matrices in this way:

```
__declspec (align (16))
    double a [10][ ((11 * sizeof (double) + 15) / 16) * 16 / sizeof (double)],
           b [10][ ((11 * sizeof (double) + 15) / 16) * 16 / sizeof (double)];
int i, j;

for (j = 0; j < 10; j++)
    for (i = 0; i < 11; i++)
        b [j][i] = a [j][i] * M_1_PI;
```

However, be aware of cache-bank conflicts for best performance. For more information, see section 6.7, "L1 Data Cache Bank Conflicts" on page 114.

Chapter 4 General 64-Bit Optimizations

The AMD x86-64 architecture provides a compatibility mode, which allows a 64-bit operating system to run existing 16-bit and 32-bit applications, and a 64-bit mode, which provides 64-bit addressing and expanded register resources to improve performance for recompiled 64-bit programs. This chapter presents general optimizations to improve the performance of software designed to run in 64-bit mode.

This chapter covers the following topics:

Topic	Page
64-Bit Registers and Integer Arithmetic	71
Using 64-bit Arithmetic for Large-Integer Multiplication	73
128-Bit Media Instructions and Floating-Point Operations	77
32-Bit Legacy GPRs and Small Unsigned Integers	77

4.1 64-Bit Registers and Integer Arithmetic

Optimization

Use 64-bit registers for 64-bit integer arithmetic.

Rationale

Using 64-bit registers instead of their 32-bit equivalents can dramatically reduce the amount of code necessary to perform 64-bit integer arithmetic.

Example 1

This code performs 64-bit addition using 32-bit registers:

```
; Add ECX:EBX to EDX:EAX, and place sum in EDX:EAX.
00000000 03 C3 add eax, ebx
00000002 13 D1 adc edx, ecx
```

Using 64-bit registers, the previous code can be replaced by one simple instruction (assuming that RAX and RBX contain the 64-bit integer values to add):

```
00000000 48 03 C3 add rax, rbx
```

Although the preceding instruction requires one additional byte for the REX prefix, it is still one byte shorter than the original code. More importantly, this instruction still has a latency of only one cycle, uses two fewer registers, and occupies only one decode slot.

Example 2

To perform the low-order half of the product of two 64-bit integers using 32-bit registers, a procedure such as the following is necessary:

```
; In:          [ESP+8]:[ESP+4] = multiplicand
;             [ESP+16]:[ESP+12] = multiplier
; Out:         EDX:EAX = (multiplicand * multiplier) % 2^64
; Modifies:   EAX, ECX, EDX, EFlags

llmul PROC
    mov edx, [esp+8]    ; multiplicand_hi
    mov ecx, [esp+16]   ; multiplier_hi
    or  edx, ecx       ; One operand >= 2^32?
    mov edx, [esp+12]   ; multiplier_lo
    mov eax, [esp+4]    ; multiplicand_lo
    jnz twomul         ; Yes, need two multiplies.
    mul edx             ; multiplicand_lo * multiplier_lo
    ret                ; Done, return to caller.

twomul:
    imul edx, [esp+8]   ; p3_lo = multiplicand_hi * multiplier_lo
    imul ecx, eax       ; p2_lo = multiplier_hi * multiplicand_lo
    add  ecx, edx       ; p2_lo + p3_lo
    mul  dword ptr [esp+12] ; p1 = multiplicand_lo * multiplier_lo
    add  edx, ecx       ; p1 + p2_lo + p3_lo = result in EDX:EAX
    ret                ; Done, return to caller.

llmul ENDP
```

Using 64-bit registers, the entire product can be produced with only one instruction:

```
; Multiply RAX by RBX. The 128-bit product is stored in RDX:RAX.
00000000 48 F7 EB imul rbx
```

Related Information

For more examples of 64-bit arithmetic using only 32-bit registers, see the example on page 75 and “Efficient 64-Bit Integer Arithmetic in 32-Bit Mode” on page 152.

4.2 Using 64-bit Arithmetic for Large-Integer Multiplication

Optimization

Use 64-bit arithmetic for integer multiplication that produces 128-bit or larger products.

Background

Large integer multiplications (those involving 128-bit or larger products) are utilized in a variety of applications, such as cryptography software, which figure prominently in e-commerce applications and secure transactions on the Internet. Processors cannot perform large-number multiplication natively; they must break the operation into chunks that are permitted by their architecture (32-bit or 64-bit additions and multiplications).

Rationale

Using 64-bit rather than 32-bit integer operations dramatically reduces the number of additions and multiplications required to compute large products. For example, computing a 1024-bit product using 64-bit arithmetic requires fewer than one quarter the number of instructions required when using 32-bit operations:

Comparing...	32-bit arithmetic	64-bit arithmetic
Number of multiplications	256	64
Number of additions with carry	509	125
Number of additions	255	63

In addition, the processor performs 64-bit additions just as fast as it performs 32-bit additions, and the latency of 64-bit multiplications is only slightly higher than for 32-bit multiplications. (The processor is capable of performing three independent 64-bit additions each clock cycle and a 64-bit multiplication every other clock cycle.)

Example

Consider the multiplication of two unsigned 64-bit numbers a and b , represented in terms of 32-bit components $a1:a0$ and $b1:b0$.

$$a = a1 * 2^{32} + a0$$

$$b = b1 * 2^{32} + b0$$

The product of a and b , calculated using the FOIL method of the polynomials above, can be expressed in terms of products of the 32-bit components, as follows:

Formula 3.1

$$c = (a1 * b1) * 2^{64} + (a1 * b0 + a0 * b1) * 2^{32} + (a0 * b0)$$

Each of the products of the components of a and b (for example, $a1 * b1$) is composed of 64 bits—an upper 32 bits and a lower 32 bits. It is convenient to represent these individual products as d , e , f , and g , as follows:

$$a0 * b0 = d1:d0 = d1 * 2^{32} + d0$$

$$a1 * b0 = e1:e0 = e1 * 2^{32} + e0$$

$$a0 * b1 = f1:f0 = f1 * 2^{32} + f0$$

$$a1 * b1 = g1:g0 = g1 * 2^{32} + g0$$

Substitution into Formula 3.1 above yields the following equation:

Formula 3.2

$$c = (g1 * 2^{32} + g0) * 2^{64} + (e1 * 2^{32} + e0 + f1 * 2^{32} + f0) * 2^{32} + (d1 * 2^{32} + d0)$$

Simplifying yields this equation:

Formula 3.3

$$c = g1 * 2^{96} + (e1 + f1 + g0) * 2^{64} + (d1 + e0 + f0) * 2^{32} + d0$$

It is convenient to represent the terms that are multiplied by each power of 2 as $c3$, $c2$, $c1$, and $c0$, as follows:

$$g1 = c3$$

$$e1 + f1 + g0 = c2$$

$$d1 + e0 + f0 = c1$$

$$d0 = c0$$

Substituting again yields:

Formula 3.4

$$c = c3 * 2^{96} + c2 * 2^{64} + c1 * 2^{32} + c0$$

The following procedure performs 64-bit unsigned integer multiplication, as previously illustrated using 32-bit integer operations:

```

; 32bitalu_64x64(int *a, int *b, int *c);
;
; TO ASSEMBLE INTO *.obj DO THE FOLLOWING:
;     ml.exe -coff -c 32bitalu_64x64.asm
;
.586
.K3D
.XMM
_DATA SEGMENT
tempESP dd 0
_DATA ENDS
_TEXT SEGMENT
ASSUME DS:_DATA
PUBLIC _32bitalu_64x64
_32bitalu_64x64 PROC NEAR
;=====
; Save the register state. Registers EAX, ECX, and EDX are considered volatile
; and assumed to be changed, while the registers below must be preserved.
push ebp
mov  ebp, esp
;=====
; Parameters passed into routine:
; [ebp+8] = ->a
; [ebp+12] = ->b
; [ebp+16] = ->c
;=====
push ebx
push esi
push edi
;=====
mov  esi,[ebp+8]    ; ESI = ->a
mov  edi,[ebp+12]  ; EDI = ->b
mov  ecx,[ebp+16]  ; ECX = ->c
push ebp
mov  [tempESP], esp
;=====
; Multiply 64-bit numbers a and b, each of which is composed of two 32-bit
; components:
; a = a1 * 2^32 + a0
; b = b1 * 2^32 + b0
mov  eax,[esi]     ; EAX = a0
mov  edx,[edi]     ; EDX = b0
mul  edx           ; EDX:EAX = a0*b0 = d1:d0
mov  ebx,edx       ; EDX = d1
mov  [ecx],eax     ; c0 = EAX
xor  esp,esp       ; ESP = 0
xor  ebp,ebp       ; EBP = 0
mov  eax,[esi+4]   ; EAX = a1
mov  edx,[edi]     ; EDX = b0
mul  edx           ; EDX:EAX = a1*b0 = e1:e0
add  ebx,eax       ; EBX = d1 + e0
adc  ebp,edx       ; EBP = e1 + possible carry from d1+e0
adc  esp,0         ; Collect possible carry into c3.

```

```

mov eax,[esi]      ; EAX = a0
mov edx,[edx+4]   ; EDX = b1
mul edx           ; EDX:EAX = a0*b1 = f1:f0
add ebx,eax       ; EBX = d1 + e0 + f0
adc ebp,edx       ; EBP = e1 + f1 + carry
adc esp,0         ; Collect possible carry into c3.
mov [ecx+4],ebx   ; c1 = d1 + e0 + f0

mov eax,[esi+4]   ; EAX = a1
mov edx,[edi+4]   ; EDX = b1
mul edx           ; EDX:EAX = a1*b1 = g1:g0
add ebp,eax       ; EBP = e1 + f1 + g0 + carry
adc esp,edx       ; ESP = g1 + carry
mov [ecx+8],ebp   ; c2 = e1 + f1 + g0 + carry
mov [ecx+12],esp  ; c3 = g1 + carry
;=====
; Restore the register state.
mov esp, [tempESP]
pop ebp
pop edi
pop esi
pop ebx
mov esp, ebp
pop ebp
;=====
ret
_32bitalu_64x64 ENDP
_TEXT ENDS
END

```

To improve performance and substantially reduce code size, the following procedure performs the same 64-bit integer multiplication using 64-bit instead of 32-bit operations:

```

; 64bitalu_64x64(int *a, int *b, int *c);
;
; TO ASSEMBLE INTO *.obj DO THE FOLLOWING:
;     ml64.exe -c 64bitalu_64x64.asm
;
_TEXT SEGMENT
64bitalu_64x64 PROC NEAR
;=====
; Parameters passed into routine:
; rcx = ->a
; rdx = ->b
; r8 = ->c
;=====
mov rax, [rcx]    ; RAX = [a0]
mul [rdx]         ; Multiply [a0] by [b0] such that
                  ; RDX:RAX = [c1]:[c0].
mov [r8], rax     ; Store 128-bit product of a and b.
mov [r8+8], rdx
;=====
ret
64bitalu_64x64 ENDP
END

```

4.3 128-Bit Media Instructions and Floating-Point Operations

Optimization

Use 128-bit media instructions instead of x87 or 64-bit media instructions for floating-point operations.

Rationale

In 64-bit mode, the processor provides eight additional XMM registers (XMM8–XMM15) for a total of 16. These extra registers can substantially reduce register pressure in floating-point code written using 128-bit media instructions.

Although the processor fully supports the x87 and 64-bit media instructions, there are only eight registers available to these instructions (ST(0)–ST(7) or MMX0–MMX7, respectively). Additionally, the x87 and 64-bit media instructions require cumbersome register manipulation and mode switches, unlike SIMD instructions.

For further information, see Chapter 10, “Optimizing with SIMD Instructions” on page 167.

4.4 32-Bit Legacy GPRs and Small Unsigned Integers

Optimization

Use the 32-bit legacy general-purpose registers (EAX through ESI) instead of their 64-bit extensions to store unsigned integer values whose range never requires more than 32 bits, even if subsequent statements use the 32-bit value in a 64-bit operation. (For example, use ECX instead of RCX until you need to perform a 64-bit operation; then use RCX.)

Rationale

In 64-bit mode, the machine-language representation of many instructions that operate on unsigned 64-bit register operands requires a REX prefix byte, which increases the size of the code. However, instructions that operate on a 32-bit legacy register operand do not require the prefix and have the desirable side-effect of clearing the upper 32 bits of the extended register to zero. For example, using the AND instruction on ECX clears the upper half of RCX.

Caution

Because the assembler also uses a REX prefix byte to encode the 32-bit sizes of the eight new 64-bit general-purpose registers (R8D–R15D), you should only use one of the original eight general-purpose registers (EAX through ESI) to implement this technique.

Example

The following example illustrates the unnecessary use of 64-bit registers to calculate the number of bytes remaining to be copied by an aligned block-copy routine after copying the first few bytes having addresses not meeting the routine's 8-byte-alignment requirements. The first two statements, after the program comments, use the 64-bit R10 register—presumably, because this value is later used to adjust a 64-bit value in R8—even though it requires no more than four bits to represent the range of values stored in R10. Using R10 instead of a smaller register requires a REX prefix byte (in this case, 49), which increases the size of the machine-language code.

```
; Input:
; R10 = source address (src)
; R8 = number of bytes to copy (count)
49 F7 DA      neg r10          ; Subtract the source address from 2^64.
49 83 E2 07    and r10, 7          ; Determine how many bytes were copied separately.
4D 2B C2      sub r8, r10     ; Subtract the number of bytes already copied from
                               ; the number of bytes to copy.
```

To improve code density, the following rewritten code uses ECX until it is absolutely necessary to use RCX, eliminating two REX prefix bytes:

```
F7 D9      neg ecx          ; Subtract the source address from 2^32 (the processor
                               ; clears the high 32 bits of RCX).
83 E1 07    and ecx, 7          ; Determine how many bytes were copied separately.
4C 2B C1    sub r8, rcx       ; Subtract the number of bytes already copied from
                               ; the number of bytes to copy.
```

Chapter 5 Instruction-Decoding Optimizations

The optimizations in this chapter are designed to help maximize the number of instructions that the processor can decode at one time.

The AMD Family 15h processor instruction fetcher reads 32-byte packets from the L1 instruction cache. These packets are 32-byte aligned. The instruction bytes are then merged into a 32-byte pick window. On each cycle, the in-order front-end engine selects up to three AMD x86-64 instructions to decode from the pick window.

This chapter covers the following topics:

Topic	Page
Load-Execute Instructions for Floating-Point or Integer Operands	79
32/64-Bit vs. 16-Bit Forms of the LEA Instruction	82
Take Advantage of x86 and AMD64 Complex Addressing Modes	82
Short Instruction Encodings	82
Partial-Register Writes	84
Using LEAVE for Function Epilogues	89
Alternatives to SHLD Instruction	90
Code Padding with Operand-Size Override and Multibyte NOP	92

5.1 Load-Execute Instructions for Floating-Point or Integer Operands

A *load-execute instruction* is an instruction that loads a value from memory into a register and then performs an operation on that value. Many general purpose instructions, such as ADD, SUB, AND, etc., have load-execute forms:

```
ADD rax, QWORD PTR [foo]
```

This instruction loads the value `foo` from memory and then adds it to the value in the RAX register.

The work performed by a load-execute instruction can also be accomplished by using two discrete instructions—a load instruction followed by an execute instruction. The following example employs discrete load and execute stages:

```
MOV rbx, QWORD PTR [foo]
ADD rax, rbx
```

The first statement loads the value `foo` from memory into the RBX register. The second statement adds the value in RBX to the value in RAX.

The following optimizations govern the use of load-execute instructions:

- Load-Execute Integer Instructions in section 5.1.1.
- Load-Execute SIMD Instructions with Floating-Point or Integer Operands in section 5.1.2.
- 32/64-Bit vs. 16-Bit Forms of the LEA Instruction in section 5.2.

5.1.1 Load-Execute Integer Instructions

Optimization

❖ When performing integer computations, use load-execute instructions instead of discrete load and execute instructions. Use discrete load and execute instructions only under one or more of the following circumstances:

- to explicitly schedule load and execute operations
- to avoid scheduler stalls for longer executing instructions
- if the load target will be used multiple times in different instructions

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Most load-execute integer instructions are FastPath single-decodable and can be decoded at the rate of four per cycle. Splitting a load-execute integer instruction into two separate instructions reduces decoding bandwidth and increases register pressure, which can result in lower performance if the load cannot be scheduled to hide the latency.

Under certain conditions, 64-bit code using general purpose registers that uses discrete load and execute instructions can be more efficient than code that uses load-execute instructions. In 64-bit code, and particularly in 32-bit code, if the block containing the instructions has high register pressure, the optimization advocated in this section is recommended. otherwise it is not.

5.1.2 Load-Execute SIMD Instructions with Floating-Point or Integer Operands

Optimization

❖ When performing floating-point computation using floating-point or integer source operands, use load-execute instructions instead of discrete load and execute instructions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Using load-execute floating-point instructions that take floating-point or integer operands improves performance for the following reasons:

- Denser code allows more work to be held in the instruction cache.
- Denser code generates fewer internal macro-ops, allowing the floating-point scheduler to hold more work, which increases the chances of extracting parallelism from the code.

The use of load-execute packed SIMD instructions instead of distinct load and execute instructions improves performance in cases in which data might not be aligned on a 16-byte boundary. However, this requires setting the misaligned exception mask (MXCSR[17]). Setting this bit disables general protection exceptions for unaligned loads in SIMD load-execute instructions. See also 10.3 “Unaligned and Aligned Data Access” on page 169.

Code employing 64-bit general purpose registers can be more efficient without load-execute instructions. In 64-bit code, software can benefit from the use of load-execute instructions, if there is high register pressure in the block containing the instructions. When register pressure is high, 32-bit code benefits most from the use of load-execute instructions.

Example

Avoid code such as the following, which uses discrete load and execute SIMD instructions:

```
movss xmm0, [float_var1]
movss xmm12, [float_var2]
mulss xmm0, xmm12
```

Instead, use code such as the following, which uses a load-execute SIMD floating-point instruction:

```
movss xmm0, [float_var1]
mulss xmm0, [float_var2]
```

5.2 32/64-Bit vs. 16-Bit Forms of the LEA Instruction

Optimization

Use the 32-bit or 64-bit forms of the Load Effective Address (LEA) instruction rather than the 16-bit form.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The 32-bit and 64-bit LEA instructions are implemented as DirectPath operations with an execution latency of only two cycles. The 16-bit LEA instruction, however, is a VectorPath instruction, which lowers the decode bandwidth and has a longer execution latency.

5.3 Take Advantage of x86 and AMD64 Complex Addressing Modes

Optimization

When porting from other architectures, remember that the x86 architecture provides many complex addressing modes. By building the effective address in one instruction, the instruction count can sometimes be reduced, leading to better code density and greater decode bandwidth. Refer to the section on effective addresses in the *AMD64 Architecture Programmer's Manual, Volume 1: Application Programming*, order# 24592 for more detailed information on how effective addresses are formed.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Building the effective address sometimes seems to require numerous instructions when there is a base address (such as the base of an array), an index and a displacement (if applicable). However, the x86 architecture can often handle all of this information in one instruction. This can reduce code size and results in fewer instructions to decode. As always, attention should be paid to total instruction length, latencies and whether or not the instruction choices are DirectPath (fastest) or VectorPath (slower).

Example

The first instruction sequence of five instructions having a total latency of 8 can be replaced by one instruction.

Number of Bytes	Latency	Instruction
3	1	<code>mov r11d, r10d</code>
8	2	<code>lea rcx, 68E35h</code>
3	1	<code>add r11, rcx</code>
5	3	<code>mov cl, BYTE PTR [r11+r13]</code>
2	1	<code>cmp cl, al</code>

The following instruction replaces the functionality of the above sequence.

Number of Bytes	Latency	Instruction
8	4	<code>cmp BYTE PTR [r10+r13+68E35h], al</code>

Example

These two instructions:

```
mov r11, QWORD PTR ds:0x4c65a
mov r11, QWORD PTR [r11+r8*8]
```

can be replaced by one instruction:

```
mov r11, QWORD PTR [r8*8+0x4c65a]
```

5.4 Short Instruction Encodings

Optimization

Use instruction forms with shorter encodings rather than those with longer encodings. For example, use 8-bit displacements instead of 32-bit displacements, and use the single-byte form of simple integer instructions instead of the 2-byte opcode-ModR/M form.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Using shorter instructions increases the number of instructions that can fit into the L1 instruction cache and increases the average decode rate.

Using short, 8-bit sign-extended displacements for conditional branches improves code density with no negative effects on the processor. See 7.1 “Instruction Fetch” on page 123 for more details on optimizing branches.

Example

Avoid the use of instructions having longer encodings, such as the following:

```
81 C0 78 56 34 12  add eax, 12345678h ; 2-byte opcode form (with ModRM)
81 C3 FB FF FF FF  add ebx, -5       ; 32-bit immediate value
0F 84 05 00 00 00  jz  label1       ; 2-byte opcode, 32-bit immediate value
```

Instead, choose instructions having shorter encodings, such as:

```
05 78 56 34 12  add eax, 12345678h ; 1-byte opcode form
83 C3 FB        add ebx, -5       ; 8-bit sign-extended immediate value
74 05          jz  label1       ; 1-byte opcode, 8-bit immediate value
```

5.5 Partial-Register Writes

Optimization

When writing to a register for the purpose of initialization:

- Avoid instructions that write less than 32 bits of a general purpose integer register.

- Avoid instructions that write less than 128 bits of an XMM register when using legacy SIMD instructions.
- Avoid mixing legacy SIMD instructions with AVX instructions.

Legacy SIMD instructions behave quite differently from AVX instructions with regard to the treatment of the upper bits of YMM registers corresponding to source and destination registers. Legacy SIMD instructions do not affect the upper 128-bits of the YMM registers corresponding to either source or destination XMM registers. On the other hand, 128-bit AVX instructions do not affect the upper 128 bits of the YMM registers corresponding to source XMM operands, but clear the upper 128 bits of the YMM register corresponding to the XMM destination register. 256-bit AVX instructions affect the entire 256-bits of the destination YMM register.

For these reasons, if the use of legacy SIMD instructions is unavoidable,

- Restrict the use of legacy SIMD instructions to sections of code that do not mix legacy SIMD and AVX instructions,
- Clear the upper parts of the YMM registers to zeros before executing any legacy SIMD instructions,
- When leaving a legacy SIMD section of code, clear the upper parts of all corresponding YMM registers before executing any subsequent AVX instructions.

Clearing the upper 128 bits of YMM registers can be accomplished by means of the VZEROUPPER AVX instruction.

When writing to a register during the course of a *non-initializing* operation on the register,

- Avoid partial register writes.
- Schedule any prior operations on the target register well ahead of the point in the code where the partial write is to occur.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

In order to handle partial register writes, the processor's execution core implements a data merging scheme. In the execution unit, an instruction that writes part of a register merges the modified portion with the current state of the other part of the register. This creates a false dependency on the most recent instruction that writes to any part of the register.

When writing to a register for the purpose of register initialization, it is usually possible to avoid false dependencies by careful instruction selection. For example, rather than initializing a part of a floating-point 128-bit XMM register, initialize the entire 128-bit register.

When possible, choose instructions that write 0's to the unused portions of the full register. For AVX instructions that write to YMM, the full 256 bits are written. AVX instructions that write to XMM registers zero out the upper 128 bits of the underlying YMM register. Both of these situations avoid the partial register write false dependency.

When writing to a register during the course of a non-initializing legacy SIMD operation on the register, there is usually no additional performance loss due to partial register reads and writes. This is because, in the typical case, the partial register being written to is also a source operand to the operation.

For example, the following instruction does not suffer from merge dependencies:

```
addsd, xmm1, xmm2
```

However, in some cases of non-initializing operations on a register, it is preferable to avoid partial register writes and replace them with more efficient operations. In these cases, the partial register being written by the operation is not a source for the operation. Examples are provided below.

If it is not possible to avoid writing to a part of that register, you should schedule any such prior operation on any part of the register well ahead of the point where the partial write occurs. Such cases are also listed in the examples.

A general purpose integer register is viewed as a 64-bit register internal to the processor. A floating-point XMM register is viewed as one 128-bit register internal to the processor.

AMD Fam 15h processors implement an XMM register merge optimization.

The processor keeps track of XMM registers whose upper portions have been cleared to zeros. This information can be followed through multiple operations and register destinations until non-zero data is written into a register. For certain instructions, this information can be used to bypass the usual result merging for the upper parts of the register. For instance, SQRTSS does not change the upper 96 bits of the destination register. If some instruction clears the upper 96 bits of its destination register and any arbitrary following sequence of instructions fails to write non-zero data in these upper 96 bits, then the SQRTSS instruction can proceed without waiting for any instructions that wrote to that destination register.

The instructions that benefit from this merge optimization are:

CVTPI2PS	CVTSS2SS (32-/64-BIT)	MOVSS <i>xmm1, xmm2</i>	FRCZSD
CVTSD2SS	CVTSS2SD	MOVLPS <i>xmm1, xmm2</i>	FRCZSS
CVTSS2SD 32-/64-BIT)	MOVSD <i>xmm1, xmm2</i>	MOVLPS R,R	RCPSS
ROUNDSS	ROUNDSD	RSQRTSS	SQRTSD
	SQRTSS	CVTPI2PS <i>xmm1, xmm2</i>	

This optimization enables independent iterations of a loop to proceed in parallel by recognizing when unused register contents will be zero and preventing a write merge dependency.

Another optimization recognizes MOVLPD/MOVHPD pairs and internally converts the MOVLPD to a MOVSD *xmm, mem*. Since MOVSD *xmm, mem* writes a 0 to the upper part of the destination, there is no merge dependency on prior instructions that write to the specified destination register.

There are several instructions that initialize the lower 64 bits or 32 bits of an XMM register that also zero out the upper 64 or 96 bits and, thus, do not suffer from merge dependencies. For example, the following instructions do not have merge dependencies:

```
movsd xmm, [mem64]
movss xmm, [mem32]
```

Integer operations that write to the lower 32 bits of a general purpose integer register do not have a false merge dependency because they zero out the upper 32 bits. But operations that write to portions of a general purpose integer register narrower than 32 bits should be avoided.

Example 1

Avoid

```
MOV    al, bl
```

Preferred

```
MOVZX  eax, bl
```

Example 2

Avoid

```
MOV    al, [ebx]
```

Preferred

```
MOVZX  eax, byte ptr [ebx]
```

Example 3

Avoid

```
MOV    al, 01h
```

Preferred

```
MOV    eax, 00000001h
```

Example 4

Avoid

```
VMOVLPD xmm1, QWORD PTR [eax]
```

```
VMOVHPD xmm1, QWORD PTR [eax] ; Same memory location as used for MOVLPD.
```

Preferred

```
VMOVDDUP xmm1, QWORD PTR [eax];
```


Example 5

Avoid

```
VADDPD xmm1, xmm3
VMULPD xmm1, xmm2      ; Very bad!!! Legacy SSE instruction
                        ; following 128-bit AVX instruction.
```

Preferred

```
VADDPD xmm1, xmm3
VMULPD xmm1, xmm2      ; Good! 128-bit AVX instruction
                        ; following 128-bit AVX instruction.
```

The legacy SIMD instruction has a merge dependency on the upper part of the YMM destination. Also, mixing AVX and legacy SIMD instructions should be avoided.

5.6 Using LEAVE for Function Epilogues

Optimization

The recommended optimization for function epilogues depends on whether the function allocates local variables.

If the function

Allocates local variables.
Does not allocate local variables or does not have a frame-pointer.

Then

Replace the traditional function epilogue with the LEAVE instruction.
Do not use function prologues or epilogues. Access function arguments and local variables through rSP.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Functions That Allocate Local Variables

The LEAVE instruction is a single-byte instruction and saves 2 bytes of code space over the traditional epilogue. Replacing the traditional sequence with LEAVE also preserves decode bandwidth.

Functions That Do Not Allocate Local Variables

Accessing function arguments and local variables directly through ESP frees EBP for use as a general-purpose register.

Background

The function arguments and local variables inside a function are referenced through a so-called frame pointer. In AMD64 code, the base pointer register (rBP) is customarily used as a frame pointer. You set up the frame pointer at the beginning of the function using a function prologue:

```
push ebp                ; Save old frame pointer.
mov  ebp, esp           ; Initialize new frame pointer.
sub  esp, n             ; Allocate space for local variables (only if the
                        ; function allocates local variables).
```

Function arguments on the stack can now be accessed at positive offsets relative to rBP, and local variables are accessible at negative offsets relative to rBP.

Example

The traditional function epilogue looks like this:

```
mov esp, ebp           ; Deallocate local variables (only if space was allocated).
pop ebp                ; Restore old frame pointer.
```

Replace the traditional function epilogue with a single LEAVE instruction:

```
leave
```

5.7 Alternatives to SHLD Instruction

Optimization

Where register pressure is low, replace the SHLD instruction with alternative code using ADD and ADC, or SHR and LEA.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Using alternative code in place of SHLD achieves lower overall latency and requires fewer execution resources. The 32-bit and 64-bit forms of ADD, ADC, SHR, and LEA (except 16-bit form) are DirectPath instructions, while SHLD is a VectorPath instruction. Use of the replacement code optimizes decode bandwidth because it potentially enables the simultaneous decoding of a third DirectPath instruction. However, the replacement code may increase register pressure because it destroys the contents of one register (*reg2* in the following examples) whereas the register is preserved by SHLD.

Example 1

Replace this instruction:

```
shld reg32a, reg32b, 1 ; Operands are 32-bit registers.
```

with this code sequence:

```
add reg32b, reg32b ; Operands are 32-bit registers
adc reg32a, reg32a
```

Example 2

Replace this instruction:

```
shld reg1, reg2, 2
```

with this code sequence:

```
shr reg2, 30
lea reg1, [reg1*4+reg2]
```

Example 3

Replace this instruction:

```
shld reg1, reg2, 3
```

with this code sequence:

```
shr reg2, 29
lea reg1, [reg1*8+reg2]
```

5.8 Code Padding with Operand-Size Override and Multibyte NOP

Optimization

Use the multibyte NOP instruction (0F 1Fh) to align code and space out branches..

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Occasionally it is necessary to insert neutral code fillers into the code stream (for example, for code-alignment purposes or to space out branches). Because this filler code is executable, it should take up as few execution resources as possible, should not diminish decode density, and should not modify any processor state other than to advance the instruction pointer (rIP). Although there are several possible multibyte NOP-equivalent instructions that do not change the processor state (other than rIP), combinations of the operand-size override and the multibyte NOP instruction are more efficient. These NOP instructions are only available on AMD Athlon™ and later processors. For processors older than Athlon, use the standard NOP (opcode 090h) in combination with up to three operand size override prefixes (opcode 66h). The use of more than three legacy prefixes limits decoder performance.

Example

Assign code-padding sequences like these and use them to align code and space out branches. These sequences are suitable for both 32-bit and 64-bit code, and you can use them on the AMD Family 15h processors:

```

NOP1_OVERRIDE_NOP TEXTEQU <DB 090h>
NOP2_OVERRIDE_NOP TEXTEQU <DB 066h, 090h>
NOP3_OVERRIDE_NOP TEXTEQU <DB 00fh, 01fh, 000h>
NOP4_OVERRIDE_NOP TEXTEQU <DB 00fh, 01fh, 040h, 000h>
NOP5_OVERRIDE_NOP TEXTEQU <DB 00fh, 01fh, 044h, 000h, 000h>
NOP6_OVERRIDE_NOP TEXTEQU <DB 066h, 00fh, 01fh, 044h, 000h, 000h>
NOP7_OVERRIDE_NOP TEXTEQU <DB 00fh, 01fh, 080h, 000h, 000h, 000h, 000h>
NOP8_OVERRIDE_NOP TEXTEQU <DB 00fh, 01fh, 084h, 000h, 000h, 000h, 000h, 000h>
NOP9_OVERRIDE_NOP TEXTEQU <DB 066h, 00fh, 01fh, 084h, 000h, 000h, 000h, 000h,
000h>
NOP10_OVERRIDE_NOP TEXTEQU <DB 066h, 066h, 00fh, 01fh, 084h, 000h, 000h, 000h,
000h, 000h>
NOP11_OVERRIDE_NOP TEXTEQU <DB 066h, 066h, 066h, 00fh, 01fh, 084h, 000h, 000h,
000h, 000h, 000h>

```

In certain rare situations, padding of up to 31 bytes can improve performance by aligning “hot” branch targets. For example, run-time profile information may reveal that a forward branch is very often taken. In these cases, generate padding by combining a minimum number of the large NOP instructions used in the above code-padding sequences. Note that NOP instructions which contain more than three prefix bytes degrade performance; in this case, use two NOPs to implement the alignment.

Chapter 6 Cache and Memory Optimizations

The optimizations in this chapter take advantage of the 16KB per cluster L1 caches and 2MB shared L2 caches of AMD Family 15h processors.

This chapter covers the following topics:

Topic	Page
Memory-Size Mismatches	95
Natural Alignment of Data Objects	97
Store-to-Load Forwarding Restrictions	98
Prefetch and Streaming Instructions	105
Write-Combining	114
L1 Data Cache Bank Conflicts	114
Placing Code and Data in the Same 64-Byte Cache Line	116
Memory and String Routines	117
Stack Considerations	119
Cache Issues When Writing Instruction Bytes to Memory	120
Interleave Loads and Stores	121

6.1 Memory-Size Mismatches

Optimization

❖ Avoid memory-size mismatches when different instructions operate on the same data. When one instruction stores and another instruction subsequently loads the same data, align instruction operands and keep the loads/stores of each operand the same size.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Examples—Store-to-Load-Forwarding Stalls

The following code examples result in a store-to-load-forwarding stall:

Avoid (64-bit)

```
foo DQ ? ; Assume foo is 8-byte aligned.
```

```

...
mov DWORD PTR foo, eax           ; Store a DWORD to foo.
mov DWORD PTR foo+4, ebx        ; Now store to foo+4.
mov rcx, QWORD PTR foo         ; Load a QWORD from foo.

```

Avoid (32-bit)

```

foo DQ ?                          ; Assume foo is 4-byte aligned.
...
mov    foo, eax                    ; Store a DWORD in foo.
mov    foo+4, edx                  ; Store a DWORD in foo+4.
...
vmovq xmm0, foo                   ; Load a QWORD from foo.

```

Preferred

```

mov    foo, eax
mov    foo+4, edx
...
vmovd  xmm0, foo
vunpcklps mm0, foo+4

```

Preferred If Stores Are Close to the Load

```

vmovd  xmm0, eax
mov    foo+4, edx
vunpcklps xmm0, xmm0, foo+4

```

Examples—Large-to-Small Mismatches

Avoid large-to-small mismatches, as shown in the following code:

Avoid (64-bit)

```

foo DQ ?                          ; Assume foo is 8-byte aligned.
...
mov QWORD PTR foo, rax            ; Store a QWORD to foo.
mov eax, DWORD PTR foo          ; Load a DWORD from foo.
mov edx, DWORD PTR foo+4        ; Load a DWORD from foo+4.

```

Avoid

```

movq foo, xmm0
...
mov  eax, foo
mov  edx, foo+4

```


Preferred

```
vmovd    foo, xmm0
vandpd   xmm0, xmm1, xmm0
vmovd    foo+4, xmm0
vandpd   xmm0, xmm2, xmm0
...
mov      eax, foo
mov      edx, foo+4
```

Preferred If the Contents of XMM0 are No Longer Needed

```
vmovd    foo, xmm0
vunpckhps xmm0, xmm0, xmm0
vmovd    foo+4, xmm0
...
mov      eax, foo
mov      edx, foo+4
```

Preferred If the Stores and Loads are Close Together, Option 1

```
vmovd    eax, xmm0
vandpd   xmm0, xmm1, xmm0
vmovd    edx, xmm0
vandpd   xmm0, xmm2, xmm0
```

Preferred If the Stores and Loads are Close Together, Option 2

```
vmovd    eax, xmm0
vunpckhps xmm0, xmm0, xmm0
vmovd    edx, xmm0
```

6.2 Natural Alignment of Data Objects

Optimization

❖ Make sure data objects are *naturally aligned*. An object is naturally aligned if it is located at an address that is a multiple of its size.

Locate this type of object	At an address evenly divisible by
Word	2
Doubleword	4
Quadword	8
Ten-byte (for example, TBYTE or REAL10)	8 (instead of 10)
Double quadword	16

Note: If 32-byte or 16-byte loads are performed from addresses which have lesser alignment there is one extra cycle of latency with some effects on load throughput in the pipeline.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

In AMD Family 15h and family 15h processors, any load that spans multiple 128 byte aligned blocks of memory is defined as a misaligned load. Misaligned loads are further classified as follows:

- Misaligned Load: A generic 128b misaligned load
- Cache Line Crosser: A misaligned load that spans two cache lines
- Page Crosser: A misaligned load that spans multiple pages

A misaligned store or load operation suffers a minimum one-cycle penalty in the processor's load-store pipeline. Also, using misaligned loads and stores increase the likelihood of encountering a store-to-load forwarding pitfall, especially when operating in long mode (64-bit software). (For a more detailed discussion of store-to-load forwarding issues, see "Store-to-Load Forwarding Restrictions" on page 98.)

In addition, if the Alignment Mask bit is set in Control Register 0 (CR0), an unaligned memory reference may cause an alignment check exception. For more information on this topic, see the *AMD64 Architecture Programmer's Manual, Volume 2*, order# 24593.

6.3 Store-to-Load Forwarding Restrictions

Optimization

Maintain consistent operand sizes across all loads and stores. Preferably use doubleword, quadword, or 128-bit operand sizes. Avoid store-to-load forwarding pitfalls, such as

- narrow-to-wide forwarding cases.
- mismatched addresses for stores and loads.
- misaligned data references.
- loading data from anywhere in the same doubleword of memory other than the identical start addresses of the stores when using word or byte stores

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Store-to-load forwarding refers to the process of a load reading (forwarding) data from the store buffer. Where this is possible, it can lead to a performance improvement because the load does not have to wait for the recently written (stored) data to be written to cache and then read back out again.

There are circumstances under which AMD Family 15h processor load-store (LS) architecture does not allow data to be read from a store in the store buffer. Store forwarding only occurs when the load virtual address exactly matches the store virtual address and the store size is greater than or equal to the load size. In such cases, it is impossible to load the needed data into a register until the store has retired out of the store buffer and written to the data cache. A store-buffer entry cannot retire and write to the data cache until *every* instruction before the store has completed and retired from the reorder buffer. The implication of this restriction is that all instructions in the reorder buffer, up to and including the store, must complete and retire out of the reorder buffer before the load can complete. Effectively, the load has a false dependency on every instruction up to the store.

Due to the significant depth of the LS buffer of AMD Family 15h processors, any load that is dependent on a store that cannot bypass data through the LS buffer may experience significant delays of up to tens of clock cycles, where the exact delay is a function of pipeline conditions.

The following sections describe store-to-load forwarding examples.

Store-to-Load Forwarding Pitfalls—True Dependencies

A load is *not* allowed to read data from the store-buffer entry if any of the following conditions occur:

- The start address of the load does not match the start address of the store.
- The load operand size is greater than the store operand size.
- Either the load or the store is misaligned. See Section 6.2, “Natural Alignment of Data Objects” on page 97 for additional information on alignment recommendations.
- A high byte (or word) store and a low byte (or word) store in the same aligned doubleword are followed by either a low or high byte (or word) load.

The following sections describe common-case scenarios to avoid. In these scenarios, a load has a true dependency on an LS2-buffered store, but cannot read (forward) data from a store-buffer entry.

Load Operand Size Greater than the Store Operand Size

If the following conditions are present, there is a narrow-to-wide store-buffer data-forwarding restriction:

- The operand size of the store data is smaller than the operand size of the load data.
- The range of addresses spanned by the store data covers some subrange of the addresses spanned by the load data.

Examples

Avoid

```
mov rax, 10h
mov WORD PTR [rax], bx      ; Word store
...
mov ecx, DWORD PTR [rax]   ; Doubleword load--cannot forward upper byte
                           ; from store buffer
```

Avoid

```
mov rax, 10h
mov BYTE PTR [rax+3], bl   ; Byte store
...
mov ecx, DWORD PTR [rax]  ; Doubleword load--cannot forward upper byte
                           ; from store buffer
```

Avoid

```
mov rax, 10h
movvsd QWORD PTR [rax], xmm0 ; Quadword store
movvsd QWORD PTR [rax+8], xmm1 ; Quadword store
...
movapd xmm2, XMMWORD PTR [rax] ; Octal Word load--cannot forward upper and
                                ; lower Quadwords from store buffer.
```

Preferred

```
mov rax, 10h
movapd xmm3, xmm0        ; Assumes XMM3 is available and will not be detrimental
                           ; to register pressure
vshufpd xmm3, xmm3, xmm1, 0
```

```

vmovapd xmmword ptr [rax], xmm3
....
vmovapd xmm2, xmmword ptr [rax] ; Octal Word load--can forward from
                                ; Octal word store from store buffer

```

Mismatched Store and Load Addresses

A data-forwarding restriction exists if the start address of the store data does not match the start address of the load data.

In general, wide stores can forward data to narrow loads if the start address of the load matches that of the store and neither store nor load is misaligned. However, store-to-load forwarding cannot occur if the start addresses of the load and store do not match, with the exception that stores to an aligned 128-bit location can forward to loads of 64 bits or less, starting at its upper 64-bit quadword.

Examples

Avoid

```

foo DQ ? ; Assume foo is 8-byte aligned.
vmovq [foo], xmm1 ; Store upper and lower half.
...
add eax, [foo] ; Fine
add edx, [foo+4] ; Not good!

```

Preferred

```

vmovd [foo], xmm1 ; Store lower half.
vunpckhps xmm1, xmm1, xmm1 ; Copy upper half into lower half.
vmovd [foo+4], xmm1 ; Store lower half.
...
add eax, [foo] ; Fine
add edx, [foo+4] ; Fine

```

Acceptable

```

mov rax, 10h
vmovaps XMMWORD PTR [rax], xmm0 ; Store upper and lower half.
...
vmovsd xmm1, QWORD PTR [rax] ; Fine
vmovsd xmm2, QWORD PTR [rax+8] ; Load of upper 64 bits, OK.

```

Misaligned Store-Buffer Data-Forwarding Restriction

If the following condition is present, there is a misaligned store-buffer data-forwarding restriction:

- The store or load address is misaligned. For example, a quadword store is not aligned to a quadword boundary.

A common case of misaligned store-data forwarding involves the passing of misaligned quadword floating-point data on the doubleword-aligned integer stack. Avoid the type of code shown in the following example:

```
mov  rsp, 24h
vmovups QWORD PTR [rsp]      , xmm0      ; RSP = 24h
...                          ; Store occurs to quadword misaligned address.
vmovups xmm1, QWORD PTR [rsp] ; Quadword load cannot forward from quadword
                              ; misaligned 'FSTP[ESP]' store operation.
```

Forwarding Restriction from Distinct Stores on Distinct Bytes (or Words) to a Subsequent Load on One of the Same Bytes (or Words) within the Same Aligned Doubleword Location

When there are two or more distinct stores to distinct bytes (or words) inside the same aligned doubleword memory location, it may not be possible to forward the data from the stores to a subsequent load from one of the same byte (or word) locations. Therefore, it is recommended to use doubleword, quadword or 128-bit operand sizes to allow store-to-load forwarding.

However, when there is only a single store and a single load to a byte (or word) inside an aligned doubleword location, store-to-load forwarding is allowed to occur as long as all of the other conditions listed previously in this section for store-to-load forwarding are satisfied.

Examples

In all of the examples below, the operations in between the store and the load indicated by the ... are assumed not to write to any part of the aligned doubleword under consideration.

Allowed

```
mov  rax, 10h
mov  BYTE PTR [rax], bl      ; Low-byte store to an aligned doubleword
...
mov  dl, BYTE PTR[rax]      ; Low byte load CAN forward from low byte store
```

Allowed

```
mov  rax, 10h
mov  WORD PTR [rax], bx     ; Low-word store to an aligned doubleword
...
mov  dl, BYTE PTR[rax]     ; Low byte load CAN forward from low word store
```

Allowed

```

mov rax, 10h
mov DWORD PTR [rax], ebx      ; Doubleword store to an aligned doubleword
...
mov dx, WORD PTR[rax]        ; Low word load CAN forward from doubleword store

```

Avoid

```

mov rax, 10h
mov BYTE PTR[rax], bl        ; Low-byte store to an aligned doubleword
mov BYTE PTR[rax+1], bh      ; High-byte store to an aligned doubleword
...
mov dl, BYTE PTR [rax]       ; low byte load cannot forward from low byte store

```

Preferred

```

mov rax, 10h
mov WORD PTR [rax], bx ;
...
mov dl, BYTE PTR [rax] ;

```

Avoid

```

mov rax, 10h
mov BYTE PTR[rax], bl        ; Low-byte store to an aligned doubleword
mov BYTE PTR[rax+1], bh      ; High-byte store to an aligned doubleword
...
mov dl, BYTE PTR [rax]       ; Low byte load cannot forward from low byte store
mov dh, BYTE PTR [rax+1]     ; High byte load CAN forward from high byte store

```

Preferred

```

mov rax, 10h
mov WORD PTR [rax], bx ;
...
mov dx, WORD PTR [rax] ;

```

Summary of Store-to-Load-Forwarding Pitfalls to Avoid

The following list summarizes the situations that require care to handle store-to-load forwarding cases:

- Avoid narrow-to-wide forwarding cases.
- Avoid mismatched addresses for stores and loads.
- Avoid misaligned data references.
- When using word or byte stores, avoid having two or more distinct stores to distinct bytes (or words) inside the same aligned doubleword memory location followed by a subsequent load from one of the same byte (or word) locations.

- Maintain consistent operand sizes across all loads and stores. Preferably use doubleword, quadword, or 128-bit operand sizes.

6.4 Good Practices for Avoiding False Store-to-Load Forwarding

Optimization

Choose linear addresses for the source and destination operands of REP MOVS/CMPS that are not an exact multiple of 4K pages away from each other.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

As mentioned in the previous section, store-to-load forwarding occurs when the store address matches the load address. This address match is split into two stages. In the first stage, bits 4:11 of the store and the load addresses are matched. In addition the double word mask of the store and load addresses is matched. The double word mask indicates whether the load/store pair is accessing the same double word in a 16-byte bank. If both these parameters match, then a store-to-load forward is initiated. In the second stage the remaining bits 12:47 of the store and load addresses is matched. If the remaining bits match, then the STLF is considered as a true STLF and is allowed to proceed. Otherwise it is considered as a false STLF and the load is cancelled and retried.

The previous section deals with true STLF and describes the practices to follow to promote it. This section deals with the cases of *false* STLF and what the developer needs to do to avoid these from occurring in the first place, thereby avoiding the later penalty of STLF cancellation.

Example

For REP MOVS/CMPS, choose linear addresses that avoid conflicts.

REP stands for the repeat function. This function repeats or iterates its associated string instruction as many times as specified in the counter register (rCX) and terminates the repetition when the value in rCX reaches 0. For example, REP MOVS moves a string from a source address to a destination address a specified number of times. In the event that bits 4:11 of the linear address of the store address in the first iteration match the load address in the second iteration, a store-to-load forward may be initiated.

When the destination address of an iteration is located at an exact multiple of 4K pages away from the source address of the next iteration, an STLF will be initiated. When the remaining address bits are found to be mismatched later, the STLF is cancelled and the load has to be retried. This results in a significant penalty of wasted DC bandwidth due to having to retry loads multiple times.

For example, a REP MOVS instruction suffers from these inefficiencies if RSI is 0x1ffeee000000 and RDI is 0x1ffeee401000.

6.5 Prefetch and Streaming Instructions

Optimization

❖ Where appropriate, use one of the prefetch instructions to increase the effective bandwidth of AMD Family 15h processors.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Prefetch instructions take advantage of the high bus bandwidth of AMD Family 15h processors to hide latencies when fetching data from system memory. A prefetch instruction initiates a read request of a specified address and reads the entire cache line that contains that address.

AMD Family 15h processors perform three types of prefetches:

Prefetch type	Description
Load	Reads the data into the L1 data cache; the data is later evicted to the L2 cache. The following instructions perform load prefetches: PREFETCH, PREFETCHT0, PREFETCHT1, and PREFETCHT2.
Store	Reads the data into the L1 data cache and marks the data as modified; the data is later evicted to the L2 cache. The PREFETCHW instruction performs a store prefetch.
Nontemporal	The PREFETCHNTA instruction performs a nontemporal prefetch. The data is read into the L1 data cache; to avoid cache pollution, when a PREFETCHNTA misses in the L2 cache and reads from memory, the data is never evicted to the L2 cache. When a PREFETCHNTA hits in the L2 cache, the data is evicted back to the L2 cache.

The prefetch instructions can be used anywhere, in any type of code. The use of prefetch instructions is not affected by the values of Control Register 0 (CR0) bits, such as CR0.EM and CR0.TS.

Prefetching versus Preloading

In code that makes irregular memory accesses rather than sequential accesses, an ordinary MOV instruction is the best way to load data. But in situations where sequential addresses are read, prefetch instructions can improve performance. Prefetch instructions only update the L1 data cache and do not update an architectural register.

Unit-Stride Access

Large data sets typically require unit-stride access to ensure that all data pulled in by a prefetch instruction are actually used. Large data sets make use of all data that are read from memory, rather than using only a sparse subset of the memory. If necessary, you should reorganize algorithms or data structures to allow unit-stride access. For a definition of unit-stride access, see “Definitions” on page 112.

Hardware Prefetcher Optimizations

In AMD Family 15h processors, the data hardware prefetches data into either the L1 or L2 cache. The L1 hardware prefetcher in AMD Family 15h processors is a stride prefetcher that is triggered by L1 cache misses and received training data from the L2 prefetcher. Stride of up to ± 504 bytes accessed from the same RIP will be detected and this data is used to start the stride prefetcher. Stride prefetching will propagate as long as requests continue at the same stride, or until a 4K page boundary is encountered. Up to 12 different stride patterns can be active at one time.

Region Prefetcher loads into L2

The L2 prefetcher provides high-performance prefetch for both strided and non-strided access patterns. The L2 strided prefetcher works in cooperation with the L1 strided prefetcher to prefetch a large number of streams and to prefetch further ahead in the data stream. The L2 non-strided prefetcher captures correlated data accesses that don't have fixed stride relationship. The L2 prefetcher can capture up to 4096 independent streams or correlated data patterns. Loads, stores, L1 hardware prefetch, and software prefetch other than NT prefetches or loads, to/from addresses currently not cached, are all tracked by the L2 prefetchers.

Contraindications for Prefetching

There are situations in which careless software prefetching can hurt performance.

- Thrashing—This is potentially the worst scenario. Thrashing occurs if more than two arrays are prefetched in parallel and the addresses are separated by whole multiples of 4K bytes (the 16-KB 4-way L1 cache size divided by the associativity). When this occurs, some of the prefetched data evicts other prefetched data before it can be used. This is inefficient even without prefetching—which simply makes the situation worse. Thrashing can be particularly bad if PREFETCHNTA is used.

- **Cache pollution**—This is a problem when the code prefetches a large amount of unused data, such as when the data is used conditionally or consists of many short sequences and the prefetches extend beyond the ends of the ranges of addresses that are actually desired.
- **Prefetch from unmapped pages**—This occurs when there is a prefetch in a loop, and the prefetch address is simply the data address plus some offset. Normally you should make the offset large enough so the data is fetched before the loop catches up to it, but this means there will be some over-run at the end of the loop. An over-run in an unmapped page can result in a significant delay. This is not so important if the over-run falls at the end of a very long stream of useful data.

In general, prefetching is useful where the program is neither totally memory-bound nor totally compute-bound, and the pattern of data access is fairly predictable within the code. The ideal fetch-ahead distance depends on the code, on the DRAM latency, and on how the data is laid out in address space.

The following table summarizes which prefetch instructions to use based on data size and data type.

Table 3. Prefetching Guidelines

Data	Less Than ½ L1 Size	Less Than ½ L2 Size or Unknown Size		Greater Than ½ L2 Size
		Reused	Not Reused	
Read-Only	PREFETCH ¹ or PREFETCHNTA	PREFETCH ¹	PREFETCHNTA	PREFETCHNTA
Sequential Read-Only	Prefetcher + PREFETCH ^{1,3}	Prefetcher + PREFETCH ^{1,3}	PREFETCHNTA	PREFETCHNTA
Read-Write	PREFETCHW	PREFETCHW	PREFETCHNTA	PREFETCHNTA
Sequential Read-Write	PREFETCHW	PREFETCHW	PREFETCHNTA	PREFETCHNTA
Write-Only	PREFETCHW	PREFETCHW	MOVNT ^{2,5}	MOVNT ^{2,5}
Sequential Write-Only	Prefetcher + PREFETCHW ⁴	Prefetcher + PREFETCHW ⁴	MOVNT ^{2,5}	MOVNT ^{2,5}

Notes:

1. *PREFETCH* is a place-holder for any of *PREFETCH*, *PREFETCHT0*, *PREFETCHT1* or *PREFETCHT2*.
2. *MOVNT* is a place-holder for any of *MOVNTI*, *MOVNTQ*, *MOVNTDQ*, *MOVNTPD*, *MOVNTPS*, *MOVNTSD*, *MOVNTSS*, *MASKMOVQ* or *MASKMOVDQU*.
3. Use *PREFETCH*¹ twice before iterations to jump-start the prefetcher, if advantageous. Otherwise, do not use *PREFETCH*¹.
4. Use *PREFETCHW* twice before iterations to jump-start the prefetcher, if advantageous. Otherwise, do not use *PREFETCHW*.
5. If no suitable *MOVNT2* instruction is available, use *PREFETCHNTA*.
6. Use *MOVNTDQA*. *NT* loads, if executed ahead of prefetch, guarantee data comes into core with *NT* type, and is written only to L1 data cache, not to L2 cache.
7. Use *NT* store, which writes directly to memory. Normal store would write line to L2. L2 does not find that line and then must fetch it from memory.

For guidance on when to use software prefetching for memory and string routines, see Section 6.9, “Memory and String Routines” on page 117.

PREFETCH/W versus PREFETCHNTA/T0/T1/T2

PREFETCHNTA, PREFETCHT0, PREFETCHT1, and PREFETCHT2 are SIMD instructions and are processor-implementation dependent. For AMD Family 15h processors, data that is prefetched with the PREFETCHNTA instruction is not placed into the L2 cache when it is evicted unless it was originally in L2 when prefetched.

PREFETCHNTA is intended for non-temporal data that will not be needed again soon. PREFETCHNTA should also be used when reading arrays that are so large that they are larger than the L2 cache. Because of their size, such large arrays will not be available in L2 even if they are needed again, and by feeding them through the L2 cache, other possibly useful data will also be evicted from L2.

Note: The sizes of the L1 and L2 caches of the processor can be determined by using the CPUID instruction.

Note: DC misses on PREFETCHNTA trigger the hardware prefetcher on AMD Family 15h processors, but those prefetch streams are marked “NT”, so that they are not evicted back to L2 or L3.

Note: PREFETCHNTA should not be used for large arrays that are only being written, not read. In such cases, write-combining stores should be used. (See “Write-Combining” on page 114, Appendix A, “Implementation of Write-Combining.”, and “Write-Combining” in the AMD64 Architecture Programmer’s Manual, Volume 2, order# 24593.)

AMD Family 15h processors implement the PREFETCHT0, PREFETCHT1, and PREFETCHT2 instructions in exactly the same way as the PREFETCH instruction. That is, the data is brought into the L1 data cache. This functionality could change in future implementations of the AMD Family 15h processor.

PREFETCHW versus PREFETCH

Code intended to modify the cache line that is brought in through prefetching should use the PREFETCHW instruction. PREFETCHW provides a hint to the AMD Family 15h processor of an intent to modify the cache line. The AMD Family 15h processor marks the cache line being read by PREFETCHW as *modified*. Using PREFETCHW can save additional cycles compared to PREFETCH, and avoid the subsequent cache state change caused by a write to the prefetched cache line. Only use PREFETCHW if there is a write to the same cache line afterwards.

Use of Streaming Instructions

Use streaming instructions instead of PREFETCHW in situations where all of the following conditions are true:

- The code will overwrite one or more complete cache lines with new data.

- The new data will not be used again soon.

Streaming instructions include the non-temporal stores MOVNTDQ, MOVNTI, MOVNTPS, MOVNTPD, MOVNTSD, MOVNTSS and the MMX instruction MOVNTQ. However, unlike regular stores, non-temporal stores are weakly ordered relative to other loads and stores. If strong ordering of stores is required, an SFENCE instruction should be used between the non-temporal stores and any succeeding normal stores. See Section 11.4, “Memory Barrier Operations” on page 207 for further recommendations on memory barrier instructions.

Streaming instructions can dramatically improve memory-write performance. They write data directly to memory through write-combining buffers, bypassing the cache. This is faster than PREFETCHW because data does not need to be initially read from memory to fill the cache lines, only to be completely overwritten shortly thereafter. The new data is simply written to memory, replacing the old data in memory, so no memory read is performed.

One application where streaming is useful, often in conjunction with prefetch instructions, is in copying large blocks of memory.

***Note:** The streaming instructions are not recommended or necessary for write-combined memory regions since the processor automatically combines writes for those regions. Write-combine memory types are indicated through the MTRRs and the page-attribute table (PAT).*

***Note:** For best performance, do not mix streaming instructions on a cache line with non-streaming store instructions.*

The following performance caveats apply when using streaming stores on AMD Family 15h cores.

- When writing out a single stream of data sequentially, performance of AMD Family 15h processors is comparable to previous generations of AMD processors.
- When writing out two streams of data, AMD Family 15h version 1 processors can be up to three times slower than previous-generation AMD processors. AMD Family 15h version 2 processor performance is approximately 1.5 times slower than previous AMD processors.
- When writing out four non-temporal streams, AMD Family 15h version 1 can be up to three times slower than previous AMD processors. AMD Family 15h version 2 processor performance is comparable to previous AMD processors.
- Using non-temporal stores but not writing out an entire cacheline may cause performance to be up to six times slower than previous AMD processors.

For more information on write-combining, see Appendix A, “Implementation of Write-Combining.”

Multiple Prefetches

Programmers can initiate multiple outstanding prefetches on AMD Family 15h processors. These processors can have a theoretical maximum of 16 outstanding cache misses, including prefetches. When all resources are filled by various memory read requests, the processor waits until resources become free before processing the next request. Multiple prefetch requests are essentially handled in order, prefetching data in the order that it is needed.

The following example shows how to initiate multiple prefetches when traversing more than one array.

Example—Multiple Prefetches

```
; TO ASSEMBLE INTO *.obj DO THE FOLLOWING:
; ml64.exe -c array_multiply_prf.asm
;
;
; Original C code:
;
; void array_multiply_prf(float *a, float *b, float *c, int LARGE_NUM)
; {
;   double array_a[LARGE_NUM];
;   double array_b[LARGE_NUM];
;   double array_c[LARGE_NUM];
;   int i;
;
;   for (i = 0; i < LARGE_NUM; i++) {
;     a[i] = b[i] * c[i];
;   }
; }
;
TEXT SEGMENT page 'CODE'
PUBLIC array_multiply_prefetch
array_multiply_prefetch: proc frame
;
; =====
; Parameters passed into routine according to the Microsoft AMD64 ABI:
; rcx = ->a
; rdx = ->b
; r8  = ->c
; r9  = N
; =====
    xor rbx,rbx

ALIGN 32                ; Align address of code loop to a 32-byte
boundary.

loop:
    prefetchw [rcx+rbx*8+512]           ; Eight cache lines ahead
    prefetch [rdx+rbx*8+512]           ; Eight cache lines ahead
    prefetch [r8+rbx*8+512]           ; Eight cache lines ahead
    vmovaps YMMWORD PTR [rdx+rbx*8],xmm0 ; b[i,i+1,i+2,...,i+7]
    vmulps  YMMWORD PTR [r8+rbx*8],xmm0 ; b[i] * c[i]
    vmovaps xmm0,XMMWORD PTR [rcx+rdx*8] ; a[i] = b[i] * c[i]
    vextractf128 ymm0, YMMWORD PTR [rcx+rbx*8+16],1
    vmovaps YMMWORD PTR [rdx+rbx*8+32],xmm0 ; b[i+8]
    vmulps  YMMWORD PTR [r8+rbx*8+32],xmm0 ; b[i+8] * c[i+8]
    vmovaps xmm0,XMMWORD PTR [rcx+rdx*8+32] ; a[i+8] = b[i+8] * c[i+8]
    vextractf128 ymm0, YMMWORD PTR [rcx+rbx*8+48],1

    add rbx, 8                        ; Compute next 16 products
    sub 16,r9                          ;
    jnz loop                          ; until none left.
```

```
    ret
array_multiply_prefetch ENDP
TEXT ENDS
END
```

The following optimization rules are applied to this example:

- Partially unroll loops to ensure that the data stride per loop iteration is equal to the length of a cache line. This avoids overlapping PREFETCH instructions and makes optimal use of the available number of outstanding prefetches.
- Because the array `array_a` is written rather than read, use PREFETCHW instead of PREFETCH to avoid overhead for switching cache lines to the correct state.
- Avoid the use of a microcoded 256-bit store by using `vextractf128` to store the upper half of the result operand.

Determining Prefetch Distance

When determining how far ahead to prefetch, the basic guideline is to initiate the prefetch early enough so that the data is in the cache by the time it is needed.

To determine the optimal prefetch distance, use empirical benchmarking when possible. Prefetching four to eight cache lines ahead (256 to 512 bytes) is a good starting point. Trying to prefetch either too far ahead or too soon impairs performance.

Memory-Limited versus Processor-Limited Code

Software prefetching can help to hide the memory latency, but it cannot increase the total memory bandwidth. Many loops are limited by memory bandwidth rather than processor speed, as shown in Figure 5. In these cases, the best that software prefetching can do is to ensure that enough memory requests are “in flight” to keep the memory system busy all of the time. AMD Family 15h processors support a maximum of eight concurrent memory requests to different cache lines. Multiple requests to the same cache line count as only one towards this limit of eight.

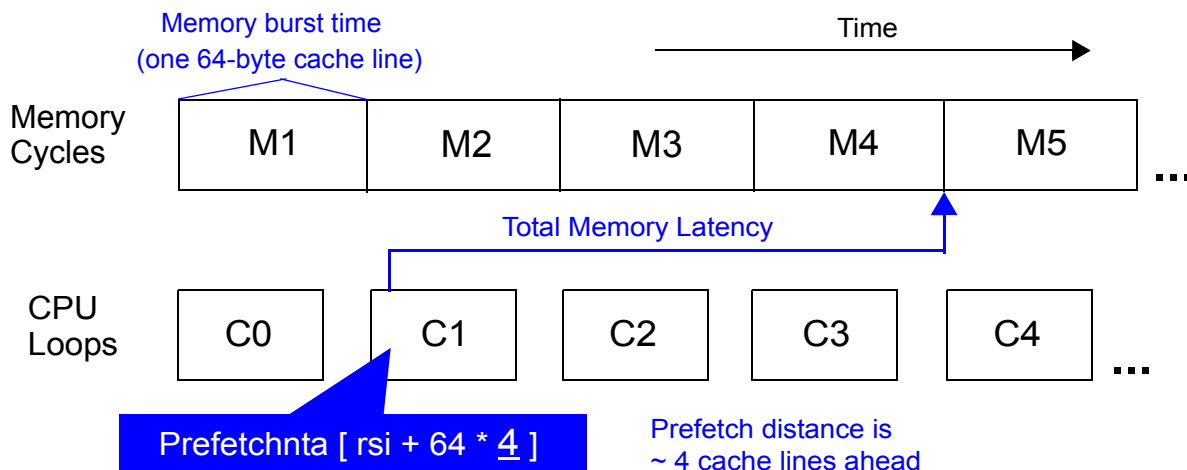


Figure 5. Memory-Limited Code

Code that performs many computations on each cache line is limited by processor speed rather than memory bandwidth, as shown in Figure 6. In this case, the goal of software prefetching is just to ensure that the memory data is available when the processor needs it. As the processor speed increases, optimal prefetch distance increases until memory bandwidth becomes the limiting factor.

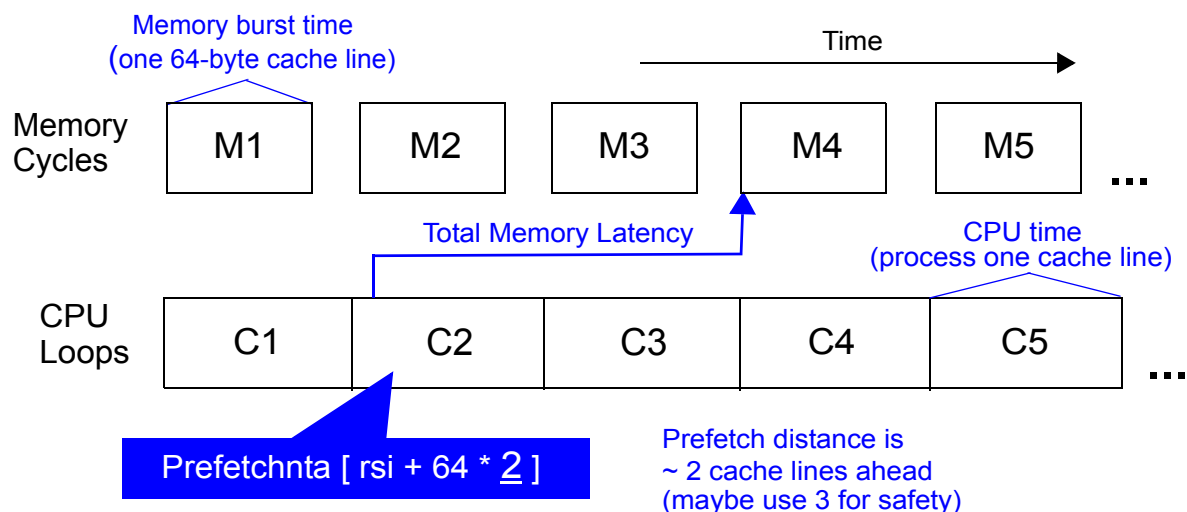


Figure 6. Processor-Limited Code

Definitions

Unit-stride access refers to a memory access pattern where consecutive memory accesses are made to consecutive array elements, in ascending or descending order. If the arrays are made of elemental types, then they imply adjacent memory locations as well. For example:

```
char j, k[MAX];
```



```

for (i = 0; i < MAX; i++) {
...
j += k[i];    // Every byte is used.
...
}
double x, y[MAX];
for (i = 0; i < MAX; i++) {
...
x += y[i];    // Every byte is used.
...
}

```

Exception to Unit Stride

The unit-stride concept works well when stepping through arrays of elementary data types. In some instances, unit stride alone may not be sufficient to determine how to use the PREFETCH instruction properly. For example, assume that there is a vertex structure of 256 bytes and the code steps through the vertices in unit stride, but using only the x, y, z, w components, each being of type `float` (for example, the first 16 bytes of each vertex). In this case, the prefetch distance obviously should be some function of the data size structure (for a properly chosen n):

```

prefetch [rax+n*structure_size]
...
add     rax, structure_size

```

You should experiment to find the optimal prefetch distance; there is no formula that works for all situations.

Data Stride per Loop Iteration

Assuming unit-stride access to a single array, the data stride of a loop (the *loop stride*) refers to the number of bytes accessed in the array per loop iteration. For example:

```

vxorpd
add_loop:
vaddsd xmm1, xmm1, QWORD PTR [rbx*8+base_address]
dec    rbx
jnz   add_loop

```

The data stride of the above loop is eight bytes. In general, for optimal use of prefetching, the data stride per iteration is the length of a cache line (64 bytes in AMD Family 15h processors). If the loop stride is smaller, unroll the loop enough to use a whole cache line of data per iteration. However, unrolling the loop may not be feasible if the original loop stride is very small (for example, only two bytes).

Prefetch at Least 64 Bytes Away from Surrounding Stores

The prefetch instructions can be affected by false dependencies on stores. If there is a store to an address that matches a request, that request (the prefetch instruction) may be blocked until the store is written to the cache. Therefore, code should prefetch data that is located at least 64 bytes away from any surrounding store's data address.

6.6 Write-Combining

Optimization

❖ Operating-system, device-driver, and BIOS programmers should take advantage of the write-combining capabilities of AMD Family 15h processors.

For details, see Appendix A, “Implementation of Write-Combining.” For more information on write-combining, see “Write-Combining” in the *AMD64 Architecture Programmer’s Manual, Volume 2*, order# 24593.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

In order to improve system performance, AMD Family 15h processors aggressively combine multiple memory-write cycles (of any data size) that address locations within a 64-byte cache-line-aligned write buffer.

6.7 L1 Data Cache Bank Conflicts

Optimization

Utilize pair loads that do not have a bank conflict in the L1 data cache to improve load throughput.


Application

This optimization applies to:

- 32-bit software
- 64-bit software

Fields Used to Address the Multibank L1 Data Cache

The L1 data cache is a multibank design consisting of sixteen banks total, where each bank is 16 bytes wide. To address the L1 data cache, the processor uses fields within the address as shown in the following diagram:



How to Know If a Bank Conflict Exists

The existence of a bank conflict between two neighboring loads depends on their bank and index values:

When the bank is	And the index is	Then a bank conflict
Different	Either the same or different	Does not exist
The same	The same	Does not exist
The same	Different	Exists

In other words, with common data types, consecutive array elements cannot have a bank conflict. If the array elements are 8 bytes or less, the two loads are to the same index and the same bank, and no conflict occurs. If the array elements are 16 bytes, the loads are to the same index but different banks, so a bank conflict does not occur either.

Rationale

Loads are served by the L1 data cache in program order, but the number of loads that the processor can perform in one cycle depends on whether a bank conflict exists between the loads:

When a bank conflict	Then the number of loads the processor can perform per cycle is
Exists	1
Does not exist	2

Therefore, pairing loads that do not have a bank conflict helps maximize load throughput.

Example

Avoid code like this, where two loads without a bank conflict are separated by other instructions:

```
vmovsd xmm0,qword ptr [rax]
vmulsd xmm0,xmm0,qword ptr [rbx]
vaddsd xmm3,xmm3,xmm0
vmovsd xmm0, qword ptr [rax+16]
vmulsd xmm0,xmm0,qword ptr [rbx+16]
vaddsd xmm2, xmm2,xmm0
```

Instead, rearrange the two loads so they appear as a pair:

```
vmovsd xmm0, qword ptr [rax]
vmovsd xmm1, qword ptr [rax+16]
vmulsd xmm0, xmm0, qword ptr [rbx]
vmulsd xmm1, xmm1, qword ptr [rbx+16]
vaddsd xmm3, xmm3, xmm0
vaddsd xmm2, xmm2, xmm1
```

6.8 Placing Code and Data in the Same 64-Byte Cache Line

Optimization

❖ Avoid placing code and data together within a cache line, especially if the data becomes modified.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Sharing code and data in the same 64-byte cache line may cause the L1 caches to thrash (unnecessarily cast out code or data) in order to maintain coherency between the separate instruction and data caches. AMD Family 15h processors have a cache-line size of 64 bytes.

For example, consider the case of a memory-indirect JMP instruction that accesses data in a jump table that resides in the same 64-byte cache line as the JMP instruction. This mixing of code and data in the same cache line degrades performance.

Do not place critical code at the border between 32-byte-aligned code segments and data segments. Code at the beginning or end of a data segment should be executed as infrequently as possible or padded.

In summary, avoid self-modifying code and storing data in code segments.

6.9 Memory and String Routines

Optimization

❖ Use the memory and string routines provided in the run-time libraries, rather than creating new custom versions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

To achieve optimal performance, it is necessary to use different memory and string manipulation algorithms to handle different block sizes and alignments. These algorithms must consider system configuration as well as the cache and memory subsystems.

The run-time libraries have optimized routines that combine several algorithms. However, if it is necessary to create fast, specific-purpose memory or string routines or to build routines that

complement the run-time library, the following pseudo-code can be used as a guide to write new routines combining different algorithms:

```
if (block size is less than 8 bytes for 32 bits or less than 16 bytes for 64
bits):
    perform operations in byte, word and doubleword for 32 bits and additionally
    in quad-word for 64 bits, starting with the widest operation;

if (block size is between 8 bytes for 32 bits or 16 bytes for 64 bits and L1
    cache-line size):
    perform operations in the natural word-size in a simple loop;

align the destination or a source block to the natural word-size;

if (block size is between cache-line size and smallest page-size):
    perform operations in the natural word-size in an unrolled loop,
    one cache-line size per iteration;

if (block size is between smallest page-size and half of L1 cache-size):
    if (there is a suitable repeated string instruction)
        use repeated string instruction;
    else
        perform operations in the natural word-size in an unrolled loop, one
        cache-line size per iteration;

if (block size is between half of L1 cache-size and half of L2 cache-size)
    perform operations in the natural word-size using temporal prefetching in an
    unrolled loop, one cache-line size per iteration;

if (block size is between half of L2 cache-size and one-fourth of a core's share
    of L3 cache-size)
    perform operations in the natural word-size using non-temporal prefetching in
    an unrolled loop, one cache-line size per iteration;
```

This pseudo-code makes the following assumptions:

- Some thresholds are specified as half of a cache level because some routines have either two sources (e.g., `strcmp()`) or a source and a destination (e.g., `memcpy()`). A routine that has a single source or destination (e.g., `strlen()` or `memset()`), could use all of a cache level for its work. However, while there is usually no drawback in using all of the L1 or even L2 caches, using all of the L3 cache can hurt the performance of other processes on a system. Using only up to a core's share of the L3 cache (e.g., on a four-core processor, up to 1/4 of the L3 cache) is recommended.
- The natural word-size is a doubleword for 32 bits and a quadword for 64 bits.
- The block size thresholds between one algorithm and the other assume that the block size is unknown at the beginning. Therefore, if the block size is known beforehand to be within a certain range, experimentation may lead to different thresholds.
- Memory routines are almost completely memory bandwidth-limited; operations within loops being limited to data movement and pointers and counter maintenance. However, string routines may additionally require some computation to find the terminating null character or to ignore

character case; this computation can dominate memory bandwidth. Therefore, some string routines may require many fewer algorithms than memory routines.

- Each core on a processor has access to exclusive L1 data cache and both cores on a compute unit share the L2 data cache and to a shared L3 cache.
- Instead of using the L2 cache-size as a threshold, particular needs and experimentation may favor using the L3 cache-size as a threshold.
- When software prefetching is used, the distance is typically eight cache lines, but experimentation may lead to a different distance.

Refer to Section 2.5, “AMD Family 15h Processor Cache Operations” on page 32 for a list and descriptions of cache operations.

See also Section 6.5, “Prefetch and Streaming Instructions” on page 105, and Section 9.3, “Repeated String Instructions” on page 150.

6.10 Stack Considerations

Optimization

Make sure the stack is suitably aligned for the local variable with the largest base type. Then, using the technique described in Section 3.16, “Sorting and Padding C and C++ Structures” on page 63, all variables can be properly aligned with no padding.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Aligning the Stack for Local Variables

A calling convention requires a certain stack alignment on function entry. For example, the Win32 32-bit ABI arranges for 32-bit stack alignment.

If a function has no local variables with a base type larger than the guaranteed stack alignment, no further work is necessary. If the function has local variables whose base type is larger than a doubleword, insert additional code to ensure proper alignment of the stack. For example, SIMD packed data requires 16-byte alignment. The following code achieves double quadword (16-byte) alignment:

```
prologue:
    push rbp
    mov  rbp, rsp
    sub  rsp, SIZE_OF_LOCALS    ; Size of local variables
```

```
    and  rsp, -16
    ...
; Push registers that need to be preserved.

epilogue:
; Pop register that needed to be preserved.
    leave
    ret
```

For functions which have local variables that need 8-byte alignment, change the above code to use:

```
    and  rsp, -8
```

With this technique, function arguments can be accessed through EBP, and local variables can be accessed through ESP. Save and restore EBP between the prologue and the epilogue to keep it free for general use.

6.11 Cache Issues When Writing Instruction Bytes to Memory

Optimization

When writing data consisting of instructions for future execution to memory use streaming store (write-combining) instructions such as MOVNTDQ and MOVNTI.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

This optimization pertains to software that writes executable instructions to memory for subsequent execution, such as might be done by a just-in-time compiler. If normal store instructions are used to write the code to memory, then the L2 cache lines will be in a modified state. When the processor eventually tries to execute the code, it will miss in the instruction cache. Because the instruction cache cannot contain cache lines that are in a modified state, the data must be flushed to memory before it can be fetched into the instruction cache. This unnecessarily evicts possibly useful information from the caches. By using write-combining instructions, the contents of the cache is preserved with no performance penalty, and this possibly provides a performance improvement.

6.12 Interleave Loads and Stores

Optimization

When loading and storing data as in a copy routine, the organization of the sequence of loads and stores can affect performance.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

When using SIMD instructions to perform loads and stores, it is best to interleave them in the following pattern—Load, Store, Load, Store, Load, Store, etc. This enables the processor to maximize the load/store bandwidth.

Example

The following example illustrates a sequence of 128-bit loads and stores:

```
vmovdqa    xmm0, [rdx+r8*8]           ; Load
vmovntdq   [rcx+r8*8], xmm0           ; Store
vmovdqa    xmm1, [rdx+r8*8+16]       ; Load
vmovntdq   [rcx+r8*8+16], xmm1       ; Store
```


Chapter 7 Branch Optimizations

The optimizations in this chapter help improve branch prediction and minimize branch penalties.

This chapter covers the following topics:

Topic	Page
Instruction Fetch	123
Branch Fusion	124
Branches That Depend on Random Data	125
Pairing CALL and RETURN	126
Nonzero Code-Segment Base Values	128
Replacing Branches	128
Avoiding the LOOP Instruction	130
Far Control-Transfer Instructions	130
Branches Not-Taken Preferable to Branches Taken	131

7.1 Instruction Fetch

As described in Chapter 2, “Microarchitecture of AMD Family 15h Processors” on page 29, each pair of integer execution units and one floating-point unit shares one instruction fetch, branch predictor, decode and dispatch unit, also known as the *shared frontend*. Code layout and alignment optimizations relative to these shared resources can improve performance. The function of this shared front end and related code optimizations are discussed below.

7.1.1 Instruction Fetch

Optimization

When possible, align branch targets to a 32-byte boundary and limit the number of branches in a 32-byte boundary to one.

Rationale

The shared front end carries out an instruction fetch in a 32-byte window. Therefore, it is recommended that tight loops not cross 32-byte boundaries, when possible. The use of NOP instructions or the assembler `.align 32` directive are recommended to achieve this alignment. The NOPs utilized should have no more than three prefix bytes prefixed to the NOP opcode. If the NOP has more than three prefix bytes, decode throughput for this instruction is reduced by more than an order of magnitude. Furthermore, to reduce branch-not-taken bubbles, it is also recommended to limit

branches to one per 32-byte fetch window. This recommendation is more relevant in the case of conditional branches.

7.1.2 Reduce Instruction Size

Optimization

Reduce the size of instructions when possible.

Rationale

Using smaller instruction sizes improves instruction fetch throughput. Specific examples include the following:

- In SIMD code, use the single-precision (PS) form of instructions instead of the double-precision (PD) form. For example, for register to register moves, MOVAPS achieves the same result as MOVAPD, but uses one less byte to encode the instruction and has no prefix byte. Other examples in which single-precision forms can be substituted for double-precision forms include MOVUPS, MOVNTPS, XORPS, ORPS, ANDPS, and SHUFPS.
- Reduce the size of displacements to a single byte by using complex addressing modes where possible.
- When shuffling the contents of a single register, use the PSHUFD, PSHUFW, and PSHUFLW instructions instead of other shuffles. These instructions do not use the source register as a destination, and thus can avoid the additional micro-op required to copy the register contents.

7.2 Branch Fusion

Optimization

Place a comparison or test instruction and its associated branch instruction sequentially adjacent in the code.

Rationale

AMD Family 15 processors introduce a new feature where, in some cases, a comparison or test instruction and its associated branch instruction can be "fused" into a single micro-operation. In order to take advantage of this optimization, the comparison and associated branch instructions must be adjacent in the code with no intervening instructions between them. This branch fusion will not occur if the comparison/test instruction is the fourth and final instruction of a dispatch group. However, once a comparison and branch are fused, it only counts as a single micro-operation in the current dispatch group.

In general, associated comparison associated and branch instructions should be adjacent in the code. However, branch fusion is not possible unless the branch target address is RIP relative. Branch fusion will not occur if the comparison instruction uses a RIP-relative addressing mode, the comparison and jump instructions both have immediate operands and displacements, or either instruction uses SIB address mode and utilizes a base which is not specified by a register but by an immediate value.

7.3 Branches That Depend on Random Data

Optimization

❖ Avoid conditional branches that depend on random data, as these branches are difficult to predict.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Suppose a piece of code receives a random stream of characters “A” through “Z” and branches if the character is before “M” in the collating sequence. Data-dependent branches acting upon basically random data cause the branch-prediction logic to mispredict the branch about 50% of the time.

If possible, design branch-free alternative code sequences that result in shorter average execution time. This technique is especially important if the branch body is small.

Examples

The following examples illustrate this concept using the CMOVxx instruction.

Signed Integer ABS Function ($x = \text{labs}(x)$)

```
mov    ecx, [x]    ; Load value.
mov    ebx, ecx    ; Save value.
neg    ecx         ; Negate value.
cmovs ecx, ebx    ; If negated value is negative, select value.
mov    [x], ecx    ; Save labs result.
```

Unsigned Integer min Function ($z = x < y ? x : y$)

```
mov    eax, [x]    ; Load x value.
mov    ebx, [y]    ; Load y value.
cmp    eax, ebx    ; EBX <= EAX ? CF = 0 : CF = 1
cmovnc eax, ebx    ; EAX = (EBX <= EAX) ? EBX : EAX
```

```
mov    [z], eax    ; Save min(X,Y) .
```

Conditional Write

```
// C code:
```

```
int a, b, i, dummy, c[BUFSIZE];
```

```
if (a < b) {
    c[i++] = a;
}
```

```
;-----
; Assembly code:
```

```
lea esi, [dummy]    ; &dummy
xor ecx, ecx        ; i = 0
...
lea  edi, [c+ecx*4]  ; &c[i]
lea  edx, [ecx+1]    ; i++
cmp   eax, ebx       ; a < b ?
cmovge edi, esi     ; ptr = (a >= b) ? &dummy : &c[i]
cmovl  ecx, edx      ; a < b ? i : i + 1
mov   [edi], eax     ; *ptr = a
```

7.4 Pairing CALL and RETURN

Optimization

For each CALL to a subroutine, use a RET instruction to return to the caller.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

As explained in section 2.6, “Branch-Prediction” on page 34, the Return Address Stack (RAS) can predict a limited number of branches by the RET instruction. CALL instructions push the next RIP on the return address stack. The corresponding RET instruction uses this address for its target prediction. If the RAS overflows, then the oldest return address is lost and the corresponding RET will likely be mispredicted, considerably lengthening the latency of the RET instruction.

When a CALL instruction is not paired with a RET instruction, the RAS can get out of sync, lengthening the latency of other RET instructions whose return addresses remain in the RAS. However, there is an important special case, shown in the following example, commonly used to get the value in the EIP register into a general-purpose register in 32-bit software:

```
CALL 0h  
POP EAX ; EAX contains the value of EIP
```

When the CALL instruction is used with a displacement of zero, it is recognized and treated specially; the RAS remains consistent even if there is not a corresponding RET instruction.

To get the value in the RIP register into a general-purpose register in 64-bit software, you can use RIP-relative addressing, as in the following example:

```
LEA RAX, [RIP+0] ; RAX contains the value of RIP.
```

7.5 Nonzero Code-Segment Base Values

Optimization

In 32-bit threads, avoid using a nonzero code-segment (CS) base value. (In 64-bit mode, segmentation is disabled and the segment base value is ignored and treated as zero.)

Application

This optimization applies to:

- 32-bit software

Rationale

A nonzero CS base value causes an additional two cycles of branch-misprediction penalty when compared with a CS base value of zero.

7.6 Replacing Branches

Optimization

Use muxing constructs to simulate conditional moves in SIMD, AVX, and XOP code.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Branches can negatively impact the performance of code. In SIMD, AVX, and XOP code, if the body of the branch is small, you can achieve higher performance instead computing both paths of the branch and using muxing constructs to construct the result. This simulates predicated execution or conditional moves. There are many SIMD AVX, and XOP instructions that can be useful for accomplishing this. The principal instructions are as follows:

ANDPS	ORPS	VCMPPS	VPMINPS	VPMINSW,
ANDPD	PAND	VCMPSD	VMINPD	(V)PMINUW,
ANDNPS	PANDN	VCMPS	VPMINSD	(V)PMINUD,
ANDNPD	PCMPEQB	VMAXPD	VPMINSS	VPMINUB,
ANDNPD	PCMPEQD	VMAXPS	VPMINPS	(V)PMAXS
CMPPD	PCMPEQW	VMAXSD	XORPD	(V)PMAXS
CMPPS	PCMPGTB	VMAXSS	VORPS	VPMAXSW,
XORPS	PCMPGTD	VPANDN	VPAND	(V)PMA XUW,
CMPSD	PCMPGTW	VPCMOV	VPANDN,	(V)PMA XU
CMPSS	PMA XS	(V)PCMPEQQ	(V)PCMPEQQ,	VPMAXUB,
MAXPD	PMA XUB	VPCOMB	VPCMPEQB,	VPOR, VPXOR,
MAXPS	PMINSW	VPCOMUB	VPCMPEQW,	VXORPS,
MAXSD	PMINUB	VPCOMUD	VPCMPEQD,	VXORPD
MAXSS	POR	VPCOMUQ	(V)PCMPGTQ,	
MINPD	PXOR	VPCOMUW	VPCMPGTB,	
MINPS	VANDNPD	VPCOMD	VPCMPGTW,	
MINSD	VANDPD	VPCOMQ	VPCMPGTD,(V)	
MINSS	VANDPS	VPCOMW	PMINSD,	
ORPD	VCMPPD	VPPERM	(V)PMINSB,	

Muxing Constructs

The most important construct to use in avoiding branches in SIMD code is a two-way muxing construct that is equivalent to the ternary operator (`?:`) in C and C++.

Examples

SIMD Solution (Preferred)

```
; r = (x < y) ? a : b
;
; In:  YMM0 = a
;     YMM1 = b
;     YMM2 = x
;     YMM3 = y
; Out: YMM0 = r
```

```
vcmpps ymm2, ymm2, ymm3, 1 ; x<y?)xffffffff:0 (Create selector in ymm2.)
vpcmov ymm2, ymm0, ymm1, ymm2 ; r=(x<y)?a:b
; (Use selector in ymm2; store result in ymm2.)
```

7.7 Avoiding the LOOP Instruction

Optimization

Avoid using the LOOP instruction.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The LOOP instruction has a latency of 7 cycles in 32-bit protected mode and 8 cycles in 64-bit protected mode.

Example

Avoid code like this, which uses the LOOP instruction:

```
label:
    ...
    loop label           ;Latency is 7/8 cycles, depending upon whether we
                        ; are in 32-bit or 64-bit protected mode.
```

Instead, replace the loop instruction with a DEC and a JNZ:

```
label:
    ...
    dec rcx             ;Latency of 1 cycle for register operand form of DEC.
    jnz label          ;Latency of 1 cycle.
```

7.8 Far Control-Transfer Instructions

Optimization

Use far control-transfer instructions only when necessary. (Far control-transfer instructions include the far forms of JMP, CALL, and RET, as well as the INT, INTO, and IRET instructions.)

Application

This optimization applies to:

- 32-bit software

- 64-bit software

Rationale

The processor's branch-prediction unit does not predict far branches.

7.9 Branches Not-Taken Preferable to Branches Taken

Optimization

Whenever possible, use branches that are biased toward being not-taken over branches that are biased toward being taken.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Correctly-predicted taken branches have at least one prediction-based bubble while not-taken branches do not. In addition, taken branches consume more branch prediction resources.

Chapter 8 Scheduling Optimizations

The optimizations discussed in this chapter help improve scheduling in the processor.

This chapter covers the following topics:

Topic	Page
Instruction Scheduling by Latency	133
Loop Unrolling	133
Inline Functions	138
MOVZX and MOVSX	139
Pointer Arithmetic in Loops	139
Pushing Memory Data Directly onto the Stack	141

8.1 Instruction Scheduling by Latency

Optimization

In general, select instructions with shorter latencies that are FastPath Single—not FastPath Double—instructions. For a list of instruction latencies and classifications, see Appendix B, “Instruction Latencies.”

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

AMD Family 15h processors can execute up to three AMD64 instructions per cycle, with each instruction possibly having a different latency. AMD Family 15h processors have flexible scheduling, but for absolute maximum performance, schedule instructions according to their latencies and data dependencies. The goal is to reduce the overall length of dependency chains.

8.2 Loop Unrolling

Optimization

Use loop unrolling where appropriate to increase instruction-level parallelism:

If all of these conditions are true	Then use
<ul style="list-style-type: none"> • The loop is in a frequently executed piece of code. • The number of loop iterations is known at compile time. • The loop body includes fewer than 10 instructions. 	Complete loop unrolling
<ul style="list-style-type: none"> • Spare registers are available (for example, when operating in 64-bit mode, where additional registers are available). • The loop body is small, so that loop overhead is significant. • The number of loop iterations is likely greater than 10. 	Partial loop unrolling

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Loop Unrolling

Loop unrolling is a technique that duplicates the body of a loop one or more times in order to increase the number of instructions relative to the branch and allow operations from different loop iterations to execute in parallel.

There are two types of loop unrolling:

- Complete loop unrolling
- Partial loop unrolling

Complete Loop Unrolling

Complete loop unrolling eliminates the loop overhead completely by replacing the loop with copies of the loop body.

Because complete loop unrolling removes the loop counter, it also reduces register pressure. However, completely unrolling very large loops can result in the inefficient use of the L1 instruction cache.

Example—Complete Loop Unrolling

In the following C code, the number of loop iterations is known at compile time and the loop body is less than 100 instructions:

```
#define ARRAY_LENGTH 3
int sum, i, a[ARRAY_LENGTH];
```

```

...
sum = 0;
for (i = 0; i < ARRAY_LENGTH; i++) {
    sum = sum + a[i];
}

```

To completely unroll an n -iteration loop, remove the loop control and replicate the loop body n times:

```

sum = 0;
sum = sum + a[0];
sum = sum + a[1];
sum = sum + a[2];

```

Partial Loop Unrolling

Partial loop unrolling reduces the loop overhead by duplicating the loop body several times, changing the increment in the loop, and adding cleanup code to execute any leftover iterations of the loop. The number of times the loop body is duplicated is known as the *unroll factor*.

However, partial loop unrolling may increase register pressure. Below, the calculations for computing the floating-point addition operations per cycle (FADDs/cycle) are shown for the loop before and after loop unrolling to illustrate why an unroll factor of two is chosen.

Example—Partial Loop Unrolling

In the following C code, each element of one array is added to the corresponding element of another array:

```

double a[MAX_LENGTH], b[MAX_LENGTH];

for (i = 0; i < MAX_LENGTH; i++) {
    a[i] = a[i] + b[i];
}

```

Without loop unrolling, this is the equivalent assembly-language code:

```

mov rcx, MAX_LENGTH    ; Initialize counter.
mov rax, OFFSET a      ; Load address of array a into RAX.
mov rbx, OFFSET b      ; Load address of array b into RBX.

```

```

add_loop:
    vmovsd  xmm0, QWORD PTR [rax] ; Load double pointed to by RAX
    vaddsd  xmm0, QWORD PTR [rbx] ; Add double pointed to by RBX
    vmovsd  QWORD PTR [rax], xmm0 ; Store double result.
    add    rax, 8                ; Point to next element of array a.
    add    rbx, 8                ; Point to next element of array b.
    dec    rcx                  ; Decrement counter.
    jnz   add_loop              ; If elements remain, then jump.

```

The rolled loop consists of seven instructions. AMD Family 15h processors can decode and retire up to four instructions per cycle. This code cannot execute faster than seven instructions in two cycles.

cyc	#instrs	instrs	note
1	3	movsd, addsd, movsd	Only one load/store pair per dispatch
2	4	add, add, dec, jnz	max of 4

With the pipelined floating-point adder allowing one FADD every cycle [still to confirm], and one FADD in each iteration, the FADDS/cycle for this rolled loop is 7/14 as calculated here:

$$\frac{7 \text{ instrs}}{2 \text{ cycles}} \times \frac{\text{iters}}{7 \text{ instrs}} \times \frac{1 \text{ FADD}}{\text{iters}} = \frac{7 \text{ FADDS}}{14 \text{ cycles}} = 0.5 \text{ FADDS/cycle}$$

After partial loop unrolling at an unroll factor of two, the new code creates a potential end case that must be handled outside the loop.

The unrolled loop consists of 10 instructions. This code can now go no faster than ten instructions per three cycles:

cyc	#instrs	instrs	note
1	3	movsd, addsd, movsd	Only one load/store pair per dispatch
2	4	movsd, addsd, movsd, add	max of 4
3	3	add, dec, jnz	max of 4

For the partially unrolled loop the FADDS/cycle is now 20/30 which is 1.333 times as fast as the original loop:

$$\frac{10 \text{ instrs}}{3 \text{ cycles}} \times \frac{\text{iters}}{10 \text{ instrs}} \times \frac{2 \text{ FADD}}{\text{iters}} = \frac{20 \text{ FADDS}}{30 \text{ cycles}} = 0.666 \text{ FADDS/cycle}$$

cyc	#instrs	instrs	note
1	3	movsd, addsd, movsd	Only one load/store pair per dispatch
2	4	add, add, dec, jnz	max of 4

With the pipelined floating-point adder allowing one FADD every cycle and one FADD in each iteration, the FADDS/cycle for this rolled loop is 7/14 as calculated here:

$$\frac{7 \text{ instrs}}{2 \text{ cycles}} \times \frac{\text{iters}}{7 \text{ instrs}} \times \frac{1 \text{ FADD}}{\text{iters}} = \frac{7 \text{ FADDS}}{14 \text{ cycles}} = 0.5 \text{ FADDS/cycle}$$

After partial loop unrolling at an unroll factor of two, the new code creates a potential end case that must be handled outside the loop:

The unrolled loop consists of 10 instructions. This code can now go no faster than ten instructions per three cycles:

cyc	#instrs	instrs	note
1	3	movsd, addsd, movsd	Only one load/store pair per dispatch
2	4	movsd, addsd, movsd, add	max of 4
3	3	add, dec, jnz	max of 4

For the partially unrolled loop the FADDS/cycle is now 20/30 which is 1.333 times as fast as the original loop:

$$\frac{10 \text{ instrs}}{3 \text{ cycles}} \times \frac{\text{iters}}{10 \text{ instrs}} \times \frac{2 \text{ FADD}}{\text{iters}} = \frac{20 \text{ FADDs}}{30 \text{ cycles}} = 0.666 \text{ FADDs/cycle}$$

Deriving the Loop Control for Partially Unrolled Loops

A frequently used loop construct is a counting loop. In a typical case, the loop count starts at some lower bound (*low*), increases by some fixed, positive increment (*inc*) for each iteration of the loop, and may not exceed some upper bound (*high*):

```
for (k = low; k <= high; k += inc) {
    x[k] = ...
}
```

The following code shows how to partially unroll such a loop by an unroll factor (*factor*) and how to derive the loop control for the partially unrolled version of the loop:

```
for (k = low; k <= (high - (factor - 1) * inc); k += factor * inc) {
    // Begin the series of unrolled statements.
    x[k + 0 * inc] = ...
    // Continue the series if the unrolling factor is greater than 2.
    x[k + 1 * inc] = ...
    x[k + 2 * inc] = ...
    ...
    // End the series.
    x[k + (factor - 1) * inc] = ...
}

// Handle the end cases.
for (k = k; k <= high; k += inc) {
    x[k] = ...
}
```

Related Information

For information on loop unrolling at the C-source level, see 3.4 “Unrolling Small Loops” on page 47.

8.3 Inline Functions

Optimization

Use function inlining when:

- A function is called from just one site in the code. (For the C language, determination of this characteristic is made easier if functions are explicitly declared `static` unless they require external linkage.)
- A function—once inlined—contains fewer than 25 machine instructions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

There are advantages and disadvantages to function inlining. On the one hand, function inlining eliminates function-call overhead and allows better register allocation and instruction scheduling at the site of the function call. The disadvantage of function inlining is decreased code reference locality, which can increase execution time due to instruction cache misses.

For functions that create fewer than 25 machine instructions once inlined, it is likely that the function-call overhead is close to, or more than, the time spent executing the function body. In these cases, function inlining is recommended.

Function-call overhead on the AMD Family 15h processors can be low because calls and returns are executed very quickly due to the use of prediction mechanisms. However, there is still overhead due to passing function arguments through memory, which creates store-to-load-forwarding dependencies. (In 64-bit mode, this overhead is typically avoided by passing more arguments in registers, as specified in the *AMD64 Application Binary Interface [ABI]* for the operating system.)

For longer functions, inlining yields diminishing returns. A function that results in the insertion of more than 500 machine instructions at the call site should probably not be inlined. Some larger functions might consist of multiple, relatively short paths. The execution time of the body of such a function may be relatively short compared to the function overhead, in which case inlining can improve performance. Profiling information is the best guide in determining whether to inline such large functions.

Additional Recommendations for Compiler Writers

In general, function inlining works best if the compiler utilizes feedback from a profiler to identify the function calls most frequently executed. If such data is not available, a reasonable approach is to

concentrate on function calls inside loops. Do not consider as candidates for inlining any functions that are directly recursive. However, if they are end-recursive, the compiler should convert them to an iterative equivalent to avoid potential overflow of the processor's return-prediction mechanism (return stack) during deep recursion. For best results, a compiler should support function inlining across multiple source files. In addition, a compiler should provide intrinsic functions for commonly used library routines, such as `sin`, `strcmp`, or `memcpy`.

8.4 MOVZX and MOVSX

Optimization

Use the MOVZX and MOVSX instructions to zero-extend or sign-extend, respectively, an operand to a larger size.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Typical code for zero extension that replaces MOVZX uses more decode and execution resources than MOVZX. It also has higher latency due to the superset dependency between the XOR and the MOV, which requires a merge operation.

Example

When zero-extending an operand (in this case, a byte), avoid code such as the following:

```
xor rax, rax
mov al, mem
```

Instead, use the MOVZX instruction:

```
movzx rax, BYTE PTR mem
```

8.5 Pointer Arithmetic in Loops

Optimization

Minimize pointer arithmetic in loops, especially if the loop bodies are small. Take advantage of scaled-index addressing modes to utilize the loop counter as an index into memory arrays.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

In small loops, pointer arithmetic causes significant overhead. Using scaled-index addressing modes has no negative impact on execution speed, but the reduced number of instructions preserves decode bandwidth.

Example

Consider the following C code, which adds the elements of two arrays and stores them in a third array:

```
int a[MAXSIZE], b[MAXSIZE], c[MAXSIZE], i;

for (i = 0; i < MAXSIZE; i++) {
    c[i] = a[i] + b[i];
}
```

Avoid an assembly-language equivalent like this, which uses base and displacement components (for example, `[esi+a]`) to compute array-element addresses, requiring additional pointer arithmetic to increment the offsets into the forward-traversed arrays:

```
    mov ecx, MAXSIZE    ; Initialize loop counter.
    xor esi, esi        ; Initialize offset into array a.
    xor edi, edi        ; Initialize offset into array b.
    xor ebx, ebx        ; Initialize offset into array c.

add_loop:
    mov eax, [esi+a]    ; Get element from a.
    mov edx, [edi+b]    ; Get element from b.
    add eax, edx        ; a[i] + b[i]
    mov [ebx+c], eax    ; Write result to c.
    add esi, 4          ; Increment offset into a.
    add edi, 4          ; Increment offset into b.
    add ebx, 4          ; Increment offset into c.
    dec ecx             ; Decrement loop count
    jnz add_loop        ; until loop count is 0.
```

Instead, traverse the arrays in a downward direction (from higher to lower addresses), in order to take advantage of scaled-index addressing (for example, `[ecx*4+a]`), which minimizes pointer arithmetic within the loop:

```
    mov ecx, MAXSIZE - 1 ; Initialize index.
```

```

add_loop:
    mov eax, [ecx*4+a]    ; Get element from a.
    mov edx, [ecx*4+b]    ; Get element from b.
    add eax, edx         ; a[i] + b[i]
    mov [ecx*4+c], eax   ; Write result to c.
    dec ecx              ; Decrement index
    jns add_loop        ; until index is negative.

```

A change in the direction of traversal is possible only if each loop iteration is completely independent of the others. If you cannot change the direction of traversal for a given array, it is still possible to minimize pointer arithmetic by using as a base address a displacement that points to the byte past the end of the array, and using an index that starts with a negative value and reaches zero when the loop expires:

```

    mov ecx, (-MAXSIZE) ; Initialize index.

add_loop:
    mov eax, [ecx*4+a+MAXSIZE*4] ; Get element from a.
    mov edx, [ecx*4+b+MAXSIZE*4] ; Get element from b.
    add eax, edx                 ; a[i] + b[i]
    mov [ecx*4+c+MAXSIZE*4], eax ; Write result to c.
    inc ecx                      ; Increment index
    jnz add_loop                 ; until index is 0.

```

If the base addresses of the arrays are held in registers (for example, when the base addresses are passed as the arguments of a function), biasing the base addresses requires additional instructions to perform the biasing at run time, and a small amount of additional overhead is incurred.

8.6 Pushing Memory Data Directly onto the Stack

Optimization

Push memory data directly onto the stack instead of loading it into a register first.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Pushing memory data directly onto the stack reduces register pressure and eliminates data dependencies.

Example

Avoid code that first loads the memory data into a register and then pushes it onto the stack:

```
mov rax, mem  
push rax
```

Instead, push the memory data directly onto the stack:

```
push mem
```

Chapter 9 Integer Optimizations

The optimizations in this chapter help improve integer performance.

This chapter covers the following topics:

Topic	Page
Replacing Division with Multiplication	143
Alternative Code for Multiplying by a Constant	147
Repeated String Instructions	150
Using XOR to Clear Integer Registers	151
Efficient 64-Bit Integer Arithmetic in 32-Bit Mode	152
Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants	159
Optimizing Integer Division	165
Efficient Implementation of Population Count and Leading-Zero Count	166

9.1 Replacing Division with Multiplication

Optimization

Replace integer division by constants with multiplication by the reciprocal.

Rationale

AMD Family 15h processors have very fast integer multiplication instructions (IMUL, MUL) whereas the integer division instructions (IDIV and DIV) are vector instructions having a variable latency that depends on the number of bits in the divisor. (For exact latencies, see “Optimizing Integer Division” on page 165 and Appendix B, “Instruction Latencies.”)

For this reason division by a constant should be replaced by multiplication by the reciprocal of the constant. The exact code to use for multiplication by the reciprocal of the constant can be found either in the examples later in this section or by using the utilities in “Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants” on page 159.

Multiplication by Reciprocal (Division) Utility

The code for the utilities is shown in “Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants” on page 159. The utilities provided in this document are for reference only and are not supported by AMD.

Signed Division Utility

The `sdiv.exe` utility finds the fastest code for *signed* division by a constant. The utility displays the code after the user enters a signed constant divisor. To redirect the code to a file, type the following command:

```
sdiv > example.out
```

Unsigned Division Utility

The `udiv.exe` utility finds the fastest code for *unsigned* division by a constant. The utility displays the code after the user enters an unsigned constant divisor. To redirect the code to a file, type the following command:

```
udiv > example.out
```

Unsigned Division by Multiplication of Constant

Algorithm: Divisors $1 \leq d < 2^{31}$, Odd d

The following code shows an unsigned division using a constant value multiplier.

```
; a = algorithm
; m = multiplier
; s = shift factor

; a == 0
mov eax, m
mul dividend
shr edx, s ; EDX = quotient

; a == 1
mov eax, m
mul dividend
add eax, m
adc edx, 0
shr edx, s ; EDX = quotient
```

Code for determining the algorithm (a), multiplier (m), and shift factor (s) from the divisor (d) is found in the section “Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants” on page 159.

Algorithm: Divisors $2^{31} \leq d < 2^{32}$

For divisors $2^{31} \leq d < 2^{32}$, the possible quotient values are either 0 or 1. For this reason, it is easy to establish the quotient by simple comparison of the dividend and divisor. When the dividend needs to be preserved, consider using code like the following:

```
; In: EAX = dividend
; Out: EDX = quotient
```



```
xor edx, edx    ; 0
cmp eax, d      ; CF = (dividend < divisor) ? 1 : 0
sbb edx, -1     ; quotient = 0 + 1 - CF = (dividend < divisor) ? 0 : 1
```

When it is not necessary to preserve the dividend, the division can be accomplished without the use of an additional register, thus reducing register pressure, as shown in the following example:

```
; In:  EAX = dividend
; Out: EDX = quotient

cmp edx, d      ; CF = (dividend < divisor) ? 1 : 0
mov eax, 0      ; 0
sbb eax, -1     ; quotient = 0 + 1 - CF = (dividend < divisor) ? 0 : 1
```

Simpler Code for Restricted Dividend

Integer division by a constant can be accelerated by limiting the range of the dividend, which removes a shift associated with most divisors. For example, for a divide-by-10 operation, use the following code, if the dividend is less than 4000_0005h:

```
mov eax, dividend
mov edx, 01999999Ah
mul edx
mov quotient, edx
```

Signed Division by Multiplication of Constant

Algorithm: Divisors $2 \leq d < 2^{31}$

The following algorithms work if the divisor is positive. If the divisor is negative, use `ABS(d)` instead of `d`, and append a `NEG edx` instruction to the code. These changes make use of the fact that $n/-d = -(n/d)$.

```
; a is the algorithm to select between two sets of code
;   sequences depending on the calculation of multiplier.
; m is the multiplier, the constant used with the multiply instruction.
; s is the amount of right shifting to accomplish the division after the
;   multiplication of a constant.

; a == 0
mov  eax, m
imul dividend
mov  eax, dividend
shr  eax, 31
sar  edx, s
add  edx, eax          ; Quotient in EDX

; a == 1
mov  eax, m
imul dividend
mov  eax, dividend
add  edx, eax
shr  eax, 31
```

```
sar edx, s
add edx, eax ; Quotient in EDX
```

Code for determining the algorithm (a), multiplier (m), and shift factor (s) is shown in “Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants” on page 159.

Signed Division by 2

```
; In: EAX = dividend
; Out: EAX = quotient

cmp eax, 80000000h ; CF = 1 if dividend >= 0.
sbb eax, -1 ; Increment dividend if it is < 0.
sar eax, 1 ; Perform right shift.
```

Signed Division by 2^n

```
; In: EAX = dividend
; Out: EAX = quotient

cdq ; Sign extend into EDX.
and edx, (2^n - 1) ; Mask correction (use divisor - 1)
add eax, edx ; Apply correction if necessary.
sar eax, (n) ; Perform right shift by log2(divisor).
```

Signed Division by -2

```
; In: EAX = dividend
; Out: EAX = quotient

cmp eax, 80000000h ; CF = 1 if dividend >= 0.
sbb eax, -1 ; Increment dividend if it is < 0.
sar eax, 1 ; Perform right shift.
neg eax ; Use (x / -2) == -(x / 2).
```

Signed Division by $-(2^n)$

```
; In: EAX = dividend
; Out: EAX = quotient

cdq ; Sign extend into EDX.
and edx, (2^n - 1) ; Mask correction (-divisor - 1).
add eax, edx ; Apply correction if necessary.
sar eax, (n) ; Right shift by log2(-divisor).
neg eax ; Use (x / -(2^n)) == -(x / 2^n).
```

Remainder of Signed Division by 2 or -2

```
; In: EAX = dividend
; Out: EAX = remainder

cdq ; Sign extend into EDX.
and eax, 1 ; Compute remainder.
xor eax, edx ; Negate remainder if
sub eax, edx ; dividend was < 0.
```

Remainder of Signed Division by 2^n or $-(2^n)$

```

; In:  EAX = dividend
; Out: EAX = remainder

cdq                ; Sign extend into EDX.
and edx, (2^n - 1) ; Mask correction (abs(divisor) - 1)
add eax, edx       ; Apply pre-correction.
and eax, (2^n - 1) ; Mask out remainder (abs(divisor) - 1)
sub eax, edx       ; Apply pre-correction if necessary.

```

9.2 Alternative Code for Multiplying by a Constant**Optimization**

Devise instruction sequences with lower latency to accomplish multiplication by certain constant multipliers.

Rationale

A 32-bit integer multiplied by a constant has a latency of 3 cycles; a 64-bit integer multiplied by a constant has a latency of 4 cycles. For certain constant multipliers, instruction sequences can be devised that accomplish the multiplication with lower latency. Because AMD Family 15h processors contain only one integer multiplier but three integer execution units, the replacement code can provide better throughput as well.

Most replacement sequences require the use of an additional temporary register, thus increasing register pressure. If register pressure in a piece of code that performs integer multiplication with a constant is already high, it could be better for the overall performance of that code to use the IMUL instruction instead of the replacement code. Similarly, replacement sequences with low latency but containing many instructions may negatively influence decode bandwidth as compared to the IMUL instruction. In general, replacement sequences containing more than four instructions are not recommended.

The following code samples are designed for the original source to receive the final result. Other sequences are possible if the result is in a different register. Sequences that do not require a temporary register are favored over those requiring a temporary register, even if the latency is higher. To keep code size small, arithmetic-logic-unit operations are preferred over shifts. Similarly, both arithmetic-logic-unit operations and shifts are favored over the LEA instruction.

There are improvements in the AMD Family 15h processors' multiplier over that of previous x86 processors. For this reason, when doing 32-bit multiplication, only use the alternative sequence if the alternative sequence has a latency that is less than or equal to 2 cycles. For 64-bit multiplication, only use the alternative sequence if the alternative sequence has a latency that is less than or equal to 3 cycles.

Examples

```
by 2:  add reg1, reg1

by 3:  lea reg1, [reg1+reg1*2]

by 4:  shl reg1, 2

by 5:  lea reg1, [reg1+reg1*4]

by 6:  lea reg1, [reg1+reg1*2]
      add reg1, reg1

by 7:  mov reg2, reg1
      shl reg1, 3
      sub reg1, reg2

by 8:  shl reg1, 3

by 9:  lea reg1, [reg1+reg1*8]

by 10: lea reg1, [reg1+reg1*4]
      add reg1, reg1

by 11: lea reg2, [reg1+reg1*8]
      add reg1, reg1
      add reg1, reg2

by 12: lea reg1, [reg1+reg1*2]
      shl reg1, 2

by 13: lea reg2, [reg1+reg1*2]
      shl reg1, 4
      sub reg1, reg2

by 14: lea reg2, [reg1+reg1]
      shl reg1, 4
      sub reg1, reg2

by 15: mov reg2, reg1
      shl reg1, 4
      sub reg1, reg2

by 16: shl reg1, 4

by 17: mov reg2, reg1
      shl reg1, 4
      add reg1, reg2

by 18: lea reg1, [reg1+reg1*8]
      add reg1, reg1

by 19: lea reg2, [reg1+reg1*2]
      shl reg1, 4
      add reg1, reg2
```

```
by 20: lea reg1, [reg1+reg1*4]
      shl reg1, 2

by 21: lea reg2, [reg1+reg1*4]
      shl reg1, 4
      add reg1, reg2

by 22: imul reg1, 22           ; Use the IMUL instruction.

by 23: lea reg2, [reg1+reg1*8]
      shl reg1, 5
      sub reg1, reg2

by 24: lea reg1, [reg1+reg1*2]
      shl reg1, 3

by 25: lea reg2, [reg1+reg1*8]
      shl reg1, 4
      add reg1, reg2

by 26: imul reg1, 26           ; Use the IMUL instruction.

by 27: lea reg2, [reg1+reg1*4]
      shl reg1, 5
      sub reg1, reg2

by 28: lea reg2, [REG1*4]
      shl reg1, 5
      sub reg1, reg2

by 29: lea reg2, [reg1+reg1*2]
      shl reg1, 5
      sub reg1, reg2

by 30: lea reg2, [reg1+reg1]
      shl reg1, 5
      sub reg1, reg2

by 31: mov reg2, reg1
      shl reg1, 5
      sub reg1, reg2

by 32: shl reg1, 5
```

9.3 Repeated String Instructions

Optimization

Use the REP prefix judiciously when performing string operations.

Rationale

In general, using the REP prefix to repeatedly perform string instructions is less efficient than other methods, especially when copying blocks of memory. Even though using the REP prefix may seem attractive due to its small code size, a loop may yield better performance due to its minimal overhead, compared to the setup overhead of using the REP prefix. However, certain string operations can benefit from using the REP prefix when the increased throughput compared to that of a loop makes up for its setup overhead for any specific repeat count.

Guidelines for Repeated String Instructions

The following sections contain guidelines for the careful scheduling of VectorPath repeated string instructions.

Use the Largest Possible Operand Size

Always move data using the largest operand size possible. For example, in 32-bit applications, use `REP MOVSD` rather than `REP MOVSW`, and `REP MOVSW` rather than `REP MOVSB`. Use `REP STOSD` rather than `REP STOSW`, and `REP STOSW` rather than `REP STOSB`.

In 64-bit mode, a quadword data size is available and offers better performance (for example, `REP MOVSQ` and `REP STOSQ`).

Make Sure that DF is 0 (Increment)

Some string instructions with `DF = 1` (decrement) may be slower.

Align Source and Destination with Operand Size

Make sure that accesses are aligned and handle the end case separately, if necessary. If there are both a source (read from) and a destination (written to) and only one can be aligned, align the destination and leave the source misaligned in order to optimize internal resources usage.

Inline REP String with Constant Small Counts

If the repeat count is constant and low (less than eight), expand REP string instructions into equivalent sequences of simple AMD64 instructions. For example, use an inline sequence of loads and stores to emulate REP MOVS or use a sequence of stores to emulate REP STOS. This technique eliminates the setup overhead of REP instructions and increases instruction throughput.

Use REP String with Constant Large Counts

If the repeat count is constant and large (in the hundreds), use REP string instructions up to approximately the data cache size. Above this limit, other techniques must be used to achieve optimal performance.

Use a Loop for REP String with Low Variable Counts

If the repeat count is variable, but is (likely) less than eight, use a simple loop to move or store the data. Otherwise, use an unrolled loop to move or store the data. These techniques avoid the overhead of REP MOVS and REP STOS.

Use a Loop for REP MOVS/CMPS If There Can Be Conflicts

The REP MOVS and REP CMPS instructions both issue two data cache operations per iteration. If certain bits of the linear addresses match, the load-store unit might have to cancel an operation and retry. To avoid this behavior, make sure the following bits in the linear address do not match:

- [6:4]—if these bits match, a cache bank conflict will occur
- [11:3]—if these bits match, a store-to-load forwarding mismatch will occur

For details, see 6.3 “Store-to-Load Forwarding Restrictions” on page 98 and 6.7 “L1 Data Cache Bank Conflicts” on page 114.

All Other Cases

For all other cases, it is best to call the appropriate routines in the run-time library, assuming that optimized routines are available. For more details on writing routines using repeated string instructions, see 6.9 “Memory and String Routines” on page 117.

9.4 Using XOR to Clear Integer Registers

Optimization

To clear an integer register to all zeros, use the XOR instruction to exclusive OR the register with itself, as shown below.

Rationale

AMD Family 15h processors are able to avoid the false read dependency on the XOR instruction.

Examples

Acceptable

```
mov reg, 0
```

Preferred

```
xor reg, reg
```

9.5 Efficient 64-Bit Integer Arithmetic in 32-Bit Mode

Optimization

The following section contains a collection of code snippets and subroutines showing the efficient implementation of 64-bit arithmetic in 32-bit mode. Note that these are 32-bit recommendations, in 64-bit mode it is important to use 64-bit integer instructions for best performance.

Addition, subtraction, negation, and shifting are best handled by inline code. Multiplication, division, and the computation of remainders are less common operations and are usually implemented as subroutines. If these subroutines are used often, the programmer should consider inlining them. Except for division and remainder calculations, the following code works for both signed and unsigned integers. The division and remainder code shown works for unsigned integers, but can easily be extended to handle signed integers.

64-Bit Addition

```
; Add ECX:EBX to EDX:EAX, and place sum in EDX:EAX.
add eax, ebx
adc edx, ecx
```

64-Bit Subtraction

```
; Subtract ECX:EBX from EDX:EAX and place difference in EDX:EAX.
sub eax, ebx
sbb edx, ecx
```

64-Bit Negation

```
; Negate EDX:EAX.
not edx
neg eax
sbb edx, -1 ; Fix: Increment high word if low word was 0.
```

64-Bit Left Shift

```
; Shift EDX:EAX left, ??shift count in ECX (count
; applied modulo 64).
shld edx, eax, cl ; First apply shift count.
shl eax, cl ; ??mod 32 to EDX:EAX
test ecx, 32 ; Need to shift by another 32?
jz lshift_done ; No, done.
mov edx, eax ; Left shift EDX:EAX
xor eax, eax ; by 32 bits
```

```
lshift_done:
```

64-Bit Right Shift

```
shrd eax, edx, cl ; First apply shift count.
shr edx, cl ; ??mod 32 to EDX:EAX
```



```

    test ecx, 32          ; Need to shift by another 32?
    jz  rshift_done      ; No, done.
    mov  eax, edx         ; Left shift EDX:EAX
    xor  edx, edx         ; by 32 bits.

```

```
rshift_done:
```

64-Bit Multiplication

```
; _llmul computes the low-order half of the product of its
; arguments, two 64-bit integers.
;
```

```
; In:      [ESP+8]:[ESP+4] = multiplicand
;          [ESP+16]:[ESP+12] = multiplier
; Out:     EDX:EAX = (multiplicand * multiplier) % 2^64
; Destroys: EAX, ECX, EDX, EFlags

```

```
_llmul PROC
    mov  edx, [esp+8]     ; multiplicand_hi
    mov  ecx, [esp+16]    ; multiplier_hi
    or   edx, ecx        ; One operand >= 2^32?
    mov  edx, [esp+12]    ; multiplier_lo
    mov  eax, [esp+4]     ; multiplicand_lo
    jnz  twomul          ; Yes, need two multiplies.
    mul  edx              ; multiplicand_lo * multiplier_lo
    ret                  ; Done, return to caller.

```

```
twomul:
    imul edx, [esp+8]     ; p3_lo = multiplicand_hi * multiplier_lo
    imul ecx, eax         ; p2_lo = multiplier_hi * multiplicand_lo
    add  ecx, edx         ; p2_lo + p3_lo
    mul  dword ptr [esp+12] ; p1 = multiplicand_lo * multiplier_lo
    add  edx, ecx         ; p1 + p2_lo + p3_lo = result in EDX:EAX
    ret                  ; Done, return to caller.

```

```
_llmul ENDP
```

64-Bit Unsigned Division

```
; _ulldiv divides two unsigned 64-bit integers and returns the quotient.
;
```

```
; In:      [ESP+8]:[ESP+4] = dividend
;          [ESP+16]:[ESP+12] = divisor
; Out:     EDX:EAX = quotient of division
; Destroys: EAX, ECX, EDX, EFlags

```

```
_ulldiv PROC
    push ebx             ; Save EBX as per calling convention.
    mov  ecx, [esp+20]   ; divisor_hi
    mov  ebx, [esp+16]   ; divisor_lo
    mov  edx, [esp+12]   ; dividend_hi
    mov  eax, [esp+8]    ; dividend_lo
    test ecx, ecx        ; divisor > (2^32 - 1)?
    jnz  big_divisor     ; Yes, divisor > 2^32 - 1.
    cmp  edx, ebx        ; Only one division needed (ECX = 0)?
    jae  two_divs        ; Need two divisions.
    div  ebx              ; EAX = quotient_lo

```

```

    mov  edx, ecx          ; EDX = quotient_hi = 0 (quotient in EDX:EAX)
    pop  ebx              ; Restore EBX as per calling convention.
    ret                   ; Done, return to caller.

two_divs:
    mov  ecx, eax         ; Save dividend_lo in ECX.
    mov  eax, edx         ; Get dividend_hi.
    xor  edx, edx        ; Zero-extend it into EDX:EAX.
    div  ebx              ; quotient_hi in EAX
    xchg eax, ecx        ; ECX = quotient_hi, EAX = dividend_lo
    div  ebx              ; EAX = quotient_lo
    mov  edx, ecx        ; EDX = quotient_hi (quotient in EDX:EAX)
    pop  ebx              ; Restore EBX as per calling convention.
    ret                   ; Done, return to caller.

big_divisor:
    push edi              ; Save EDI as per calling convention.
    mov  edi, ecx         ; Save divisor_hi.
    shr  edx, 1          ; Shift both divisor and dividend right
    rcr  eax, 1           ; by 1 bit.
    ror  edi, 1
    rcr  ebx, 1
    bsr  ecx, ecx        ; ECX = number of remaining shifts
    shrd ebx, edi, cl    ; Scale down divisor and dividend
    shrd eax, edx, cl    ; such that divisor is less than
    shr  edx, cl         ; 2^32 (that is, it fits in EBX).
    rol  edi, 1          ; Restore original divisor_hi.
    div  ebx              ; Compute quotient.
    mov  ebx, [esp+12]   ; dividend_lo
    mov  ecx, eax        ; Save quotient.
    imul edi, eax        ; quotient * divisor high word (??low only)
    mul  dword ptr [esp+20] ; quotient * divisor low word
    add  edx, edi        ; EDX:EAX = quotient * divisor
    sub  ebx, eax        ; dividend_lo - (quot.*divisor)_lo
    mov  eax, ecx        ; Get quotient.
    mov  ecx, [esp+16]   ; dividend_hi
    sbb  ecx, edx        ; Subtract (divisor * quot.) from dividend.
    sbb  eax, 0          ; Adjust quotient if remainder negative.
    xor  edx, edx        ; Clear high word of quot. (EAX<=FFFFFFFFh).
    pop  edi              ; Restore EDI as per calling convention.
    pop  ebx              ; Restore EBX as per calling convention.
    ret                   ; Done, return to caller.

_udiv  ENDP

```

64-Bit Signed Division

```
; _lldiv divides two signed 64-bit numbers and delivers the quotient
;
; In:      [ESP+8]:[ESP+4] = dividend
;         [ESP+16]:[ESP+12] = divisor
; Out:    EDX:EAX = quotient of division
; Destroys: EAX, ECX, EDX, EFlags
```

```
_lldiv PROC
    push ebx      ; Save EBX as per calling convention.
    push esi      ; Save ESI as per calling convention.
    push edi      ; Save EDI as per calling convention.
    mov ecx, [esp+28] ; divisor_hi
    mov ebx, [esp+24] ; divisor_lo
    mov edx, [esp+20] ; dividend_hi
    mov eax, [esp+16] ; dividend_lo
    mov esi, ecx    ; divisor_hi
    xor esi, edx    ; divisor_hi ^ dividend_hi
    sar esi, 31    ; (quotient < 0) ? -1 : 0
    mov edi, edx    ; dividend_hi
    sar edi, 31    ; (dividend < 0) ? -1 : 0
    xor eax, edi    ; If (dividend < 0),
    xor edx, edi    ; compute 1's complement of dividend.
    sub eax, edi    ; If (dividend < 0),
    sbb edx, edi    ; compute 2's complement of dividend.
    mov edi, ecx    ; divisor_hi
    sar edi, 31    ; (divisor < 0) ? -1 : 0
    xor ebx, edi    ; If (divisor < 0),
    xor ecx, edi    ; compute 1's complement of divisor.
    sub ebx, edi    ; If (divisor < 0),
    sbb ecx, edi    ; compute 2's complement of divisor.
    jnz big_divisor ; divisor > 2^32 - 1
    cmp edx, ebx    ; Only one division needed (ECX = 0)?
    jae two_divs   ; Need two divisions.
    div ebx        ; EAX = quotient_lo
    mov edx, ecx   ; EDX = quotient_hi = 0 (quotient in EDX:EAX)
    xor eax, esi   ; If (quotient < 0),
    xor edx, esi   ; compute 1's complement of result.
    sub eax, esi   ; If (quotient < 0),
    sbb edx, esi   ; compute 2's complement of result.
    pop edi        ; Restore EDI as per calling convention.
    pop esi        ; Restore ESI as per calling convention.
    pop ebx        ; Restore EBX as per calling convention.
    ret            ; Done, return to caller.

two_divs:
    mov ecx, eax   ; Save dividend_lo in ECX.
    mov eax, edx   ; Get dividend_hi.
    xor edx, edx   ; Zero-extend it into EDX:EAX.
    div ebx        ; quotient_hi in EAX
    xchg eax, ecx  ; ECX = quotient_hi, EAX = dividend_lo
    div ebx        ; EAX = quotient_lo
    mov edx, ecx   ; EDX = quotient_hi (quotient in EDX:EAX)
    jmp make_sign  ; Make quotient signed.
```

```

big_divisor:
    sub    esp, 12                ; Create three local variables.
    mov    [esp], eax             ; dividend_lo
    mov    [esp+4], ebx           ; divisor_lo
    mov    [esp+8], edx           ; dividend_hi
    mov    edi, ecx               ; Save divisor_hi.
    shr    edx, 1                 ; Shift both
    rcr    eax, 1                 ; divisor and
    ror    edi, 1                 ; and dividend
    rcr    ebx, 1                 ; right by 1 bit.
    bsr    ecx, ecx               ; ECX = number of remaining shifts
    shrd   ebx, edi, cl           ; Scale down divisor and
    shrd   eax, edx, cl           ; dividend such that divisor is
    shr    edx, cl                ; less than 2^32 (that is, fits in EBX).
    rol    edi, 1                 ; Restore original divisor_hi.
    div    ebx                    ; Compute quotient.
    mov    ebx, [esp]             ; dividend_lo
    mov    ecx, eax               ; Save quotient.
    imul  edi, eax                ; quotient * divisor high word (??low only)
    mul   DWORD PTR [esp+4]       ; quotient * divisor low word
    add   edx, edi                ; EDX:EAX = quotient * divisor
    sub   ebx, eax                ; dividend_lo - (quot.*divisor)_lo
    mov   eax, ecx                ; Get quotient.
    mov   ecx, [esp+8]           ; dividend_hi
    sbb   ecx, edx                ; Subtract (divisor * quot.) from dividend
    sbb   eax, 0                  ; Adjust quotient if remainder is negative.
    xor   edx, edx                ; Clear high word of quotient.
    add   esp, 12                 ; Remove local variables.

make_sign:
    xor   eax, esi                ; If (quotient < 0),
    xor   edx, esi                ; compute 1's complement of result.
    sub   eax, esi                ; If (quotient < 0),
    sbb   edx, esi                ; compute 2's complement of result.
    pop   edi                     ; Restore EDI as per calling convention.
    pop   esi                     ; Restore ESI as per calling convention.
    pop   ebx                     ; Restore EBX as per calling convention.
    ret                                ; Done, return to caller.
_lldiv ENDP

```

64-Bit Unsigned Remainder Computation

```

; _ullrem divides two unsigned 64-bit integers and returns the remainder.
;
; In:      [ESP+8]:[ESP+4] = dividend
;          [ESP+16]:[ESP+12] = divisor
;
; Out:     EDX:EAX = remainder of division
;
; Destroys: EAX, ECX, EDX, EFlags

_ullrem PROC
    push ebx                      ; Save EBX as per calling convention.
    mov  ecx, [esp+20]             ; divisor_hi
    mov  ebx, [esp+16]             ; divisor_lo
    mov  edx, [esp+12]             ; dividend_hi
    mov  eax, [esp+8]              ; dividend_lo

```

```

    test ecx, ecx           ; divisor > 2^32 - 1?
    jnz r_big_divisor     ; Yes, divisor > 32^32 - 1.
    cmp edx, ebx         ; Only one division needed (ECX = 0)?
    jae r_two_divs       ; Need two divisions.
    div ebx              ; EAX = quotient_lo
    mov eax, edx         ; EAX = remainder_lo
    mov edx, ecx         ; EDX = remainder_hi = 0
    pop ebx             ; Restore EBX per calling convention.
    ret                 ; Done, return to caller.

r_two_divs:
    mov ecx, eax        ; Save dividend_lo in ECX.
    mov eax, edx        ; Get dividend_hi.
    xor edx, edx        ; Zero-extend it into EDX:EAX.
    div ebx            ; EAX = quotient_hi, EDX = intermediate remainder
    mov eax, ecx        ; EAX = dividend_lo
    div ebx            ; EAX = quotient_lo
    mov eax, edx        ; EAX = remainder_lo
    xor edx, edx        ; EDX = remainder_hi = 0
    pop ebx            ; Restore EBX as per calling convention.
    ret                 ; Done, return to caller.

r_big_divisor:
    push edi            ; Save EDI as per calling convention.
    mov edi, ecx        ; Save divisor_hi.
    shr edx, 1          ; Shift both divisor and dividend right
    rcr eax, 1          ; by 1 bit.
    ror edi, 1
    rcr ebx, 1
    bsr ecx, ecx        ; ECX = number of remaining shifts
    shrd ebx, edi, cl   ; Scale down divisor and dividend such
    shrd eax, edx, cl   ; that divisor is less than 2^32
    shr edx, cl         ; (that is, it fits in EBX).
    rol edi, 1          ; Restore original divisor (EDI:ESI).
    div ebx             ; Compute quotient.
    mov ebx, [esp+12]   ; dividend low word
    mov ecx, eax        ; Save quotient.
    imul edi, eax       ; quotient * divisor high word (??low only)
    mul DWORD PTR [esp+20] ; quotient * divisor low word
    add edx, edi        ; EDX:EAX = quotient * divisor
    sub ebx, eax        ; dividend_lo - (quot.*divisor)_lo
    mov ecx, [esp+16]   ; dividend_hi
    mov eax, [esp+20]   ; divisor_lo
    sbb ecx, edx        ; Subtract divisor * quot. from dividend.
    sbb edx, edx        ; (remainder < 0) ? 0xFFFFFFFF : 0
    and eax, edx        ; (remainder < 0) ? divisor_lo : 0
    and edx, [esp+24]   ; (remainder < 0) ? divisor_hi : 0
    add eax, ebx        ; remainder += (remainder < 0) ? divisor : 0
    pop edi             ; Restore EDI as per calling convention.
    pop ebx            ; Restore EBX as per calling convention.
    ret                 ; Done, return to caller.

_ullrem ENDP

```

64-Bit Signed Remainder Computation

```

; _llrem divides two signed 64-bit numbers and returns the remainder.
;
; In:      [ESP+8]:[ESP+4] = dividend
;         [ESP+16]:[ESP+12] = divisor
;
; Out:     EDX:EAX = remainder of division
;
; Destroys: EAX, ECX, EDX, EFlags

    push ebx                ; Save EBX as per calling convention.
    push esi                ; Save ESI as per calling convention.
    push edi                ; Save EDI as per calling convention.
    mov ecx, [esp+28]       ; divisor-hi
    mov ebx, [esp+24]       ; divisor-lo
    mov edx, [esp+20]       ; dividend-hi
    mov eax, [esp+16]       ; dividend-lo
    mov esi, edx            ; sign(remainder) == sign(dividend)
    sar esi, 31            ; (remainder < 0) ? -1 : 0
    mov edi, edx            ; dividend-hi
    sar edi, 31            ; (dividend < 0) ? -1 : 0
    xor eax, edi            ; If (dividend < 0),
    xor edx, edi            ; compute 1's complement of dividend.
    sub eax, edi            ; If (dividend < 0),
    sbb edx, edi            ; compute 2's complement of dividend.
    mov edi, ecx           ; divisor-hi
    sar edi, 31            ; (divisor < 0) ? -1 : 0
    xor ebx, edi            ; If (divisor < 0),
    xor ecx, edi            ; compute 1's complement of divisor.
    sub ebx, edi            ; If (divisor < 0),
    sbb ecx, edi            ; compute 2's complement of divisor.
    jnz sr_big_divisor     ; divisor > 2^32 - 1
    cmp edx, ebx           ; Only one division needed (ECX = 0)?
    jae sr_two_divs        ; No, need two divisions.
    div ebx                ; EAX = quotient_lo
    mov eax, edx           ; EAX = remainder_lo
    mov edx, ecx           ; EDX = remainder_lo = 0
    xor eax, esi           ; If (remainder < 0),
    xor edx, esi           ; compute 1's complement of result.
    sub eax, esi           ; If (remainder < 0),
    sbb edx, esi           ; compute 2's complement of result.
    pop edi                ; Restore EDI as per calling convention.
    pop esi                ; Restore ESI as per calling convention.
    pop ebx                ; Restore EBX as per calling convention.
    ret                    ; Done, return to caller.

sr_two_divs:
    mov ecx, eax           ; Save dividend_lo in ECX.
    mov eax, edx           ; Get dividend_hi.
    xor edx, edx           ; Zero-extend it into EDX:EAX.
    div ebx                ; EAX = quotient_hi, EDX = intermediate remainder
    mov eax, ecx           ; EAX = dividend_lo
    div ebx                ; EAX = quotient_lo
    mov eax, edx           ; remainder_lo
    xor edx, edx           ; remainder_hi = 0
    jmp sr_makesign       ; Make remainder signed.

```

```

sr_big_divisor:
    sub    esp, 16                ; Create three local variables.
    mov    [esp], eax             ; dividend_lo
    mov    [esp+4], ebx           ; divisor_lo
    mov    [esp+8], edx           ; dividend_hi
    mov    [esp+12], ecx          ; divisor_hi
    mov    edi, ecx               ; Save divisor_hi.
    shr    edx, 1                 ; Shift both
    rcr    eax, 1                 ; divisor and
    ror    edi, 1                 ; and dividend
    rcr    ebx, 1                 ; right by 1 bit.
    bsr    ecx, ecx               ; ECX = number of remaining shifts
    shrd   ebx, edi, cl           ; Scale down divisor and
    shrd   eax, edx, cl           ; dividend such that divisor is
    shr    edx, cl                ; less than 2^32 (that is, fits in EBX).
    rol    edi, 1                 ; Restore original divisor_hi.
    div    ebx                    ; Compute quotient.
    mov    ebx, [esp]             ; dividend_lo
    mov    ecx, eax               ; Save quotient.
    imul  edi, eax                ; quotient * divisor high word (??low only)
    mul   DWORD PTR [esp+4]       ; quotient * divisor low word
    add   edx, edi                ; EDX:EAX = quotient * divisor
    sub   ebx, eax                ; dividend_lo - (quot.*divisor)_lo
    mov   ecx, [esp+8]            ; dividend_hi
    sbb   ecx, edx                ; Subtract divisor * quot. from dividend.
    sbb   eax, eax                ; remainder < 0 ? 0xffffffff : 0
    mov   edx, [esp+12]           ; divisor_hi
    and   edx, eax                ; remainder < 0 ? divisor_hi : 0
    and   eax, [esp+4]            ; remainder < 0 ? divisor_lo : 0
    add   eax, ebx                ; remainder_lo
    add   edx, ecx                ; remainder_hi
    add   esp, 16                ; Remove local variables.

sr_makesign:
    xor   eax, esi                ; If (remainder < 0),
    xor   edx, esi                ; compute 1's complement of result.
    sub   eax, esi                ; If (remainder < 0),
    sbb   edx, esi                ; compute 2's complement of result.
    pop   edi                     ; Restore EDI as per calling convention.
    pop   esi                     ; Restore ESI as per calling convention.
    pop   ebx                     ; Restore EBX as per calling convention.
    ret

```

9.6 Derivation of Algorithm, Multiplier, and Shift Factor for Integer Division by Constants

The following examples illustrate the derivation of algorithm, multiplier and shift factor for signed and unsigned integer division.

Unsigned Integer Division

The utility `udiv.exe` was compiled from the code shown in this section. The utilities provided in this document are for reference only and are not supported by AMD.

The following code derives the multiplier value used when performing integer division by constants. The code works for unsigned integer division and for odd divisors between 1 and $2^{31} - 1$, inclusive. For divisors of the form $d = d' * 2^n$, the multiplier is the same as for d' and the shift factor is $s + n$.

Example

```

/* This program determines the algorithm (a), multiplier (m), and
   shift factor (s) to be used to accomplish *unsigned* division by
   a constant divisor. Compile with MSVC.
*/

#include <stdio.h>

typedef unsigned __int64 U64;
typedef unsigned long U32;

U32 log2(U32 i)
{
    U32 t = 0;
    i = i >> 1;
    while (i) {
        i = i >> 1;
        t++;
    }
    return(t);
}

U32 res1, res2;
U32 d, l, s, m, a, r, n, t;
U64 m_low, m_high, j, k;

int main (void)
{
    fprintf(stderr, "\n");
    fprintf(stderr, "Unsigned division by constant\n");
    fprintf(stderr, "=====\n\n");
    fprintf(stderr, "enter divisor: ");
    scanf("%lu", &d);
    printf("\n");
    if (d == 0) goto printed_code;

    if (d >= 0x80000000UL) {
        printf("; dividend: register or memory location\n");
        printf("\n");
        printf("CMP    dividend, 0%08lxh\n", d);
        printf("MOV    EDX, 0\n");
        printf("SBB    EDX, -1\n");
        printf("\n");
        printf("; quotient now in EDX\n");
    }
}

```



```

    goto printed_code;
}

/* Reduce divisor until it becomes odd. */

n = 0;
t = d;
while (!(t & 1)) {
    t >>= 1;
    n++;
}

if (t == 1) {
    if (n == 0) {
        printf("; dividend: register or memory location\n");
        printf("\n");
        printf("MOV    EDX, dividend\n", n);
        printf("\n");
        printf("; quotient now in EDX\n");
    }
    else {
        printf("; dividend: register or memory location\n");
        printf("\n");
        printf("SHR    dividend, %d\n", n);
        printf("\n");
        printf("; quotient replaced dividend\n");
    }
    goto printed_code;
}

/* Generate m, s for algorithm 0. Based on: Granlund, T.; Montgomery,
P.L.: "Division by Invariant Integers using Multiplication."
SIGPLAN Notices, Vol. 29, June 1994, page 61.
*/

l = log2(t) + 1;
j = (((U64)(0xffffffff)) % ((U64)(t)));
k = (((U64)(1)) << (32 + 1)) / ((U64)(0xffffffff - j));
m_low = (((U64)(1)) << (32 + 1)) / t;
m_high = (((U64)(1)) << (32 + 1)) + k) / t;
while (((m_low >> 1) < (m_high >> 1)) && (l > 0)) {
    m_low = m_low >> 1;
    m_high = m_high >> 1;
    l = l - 1;
}
if ((m_high >> 32) == 0) {
    m = ((U32)(m_high));
    s = 1;
    a = 0;
}

/* Generate m and s for algorithm 1. Based on: Magenheimer, D.J.; et al:
"Integer Multiplication and Division on the HP Precision Architecture."
IEEE Transactions on Computers, Vol. 37, No. 8, August 1988, page 980.*/
else {
    s = log2(t);
    m_low = (((U64)(1)) << (32 + s)) / ((U64)(t));
}

```

```

    r = ((U32) (((U64) (1)) << (32 + s)) % ((U64) (t))));
    m = (r < ((t >> 1) + 1)) ? ((U32) (m_low)) : ((U32) (m_low)) + 1;
    a = 1;
}

/* Reduce multiplier for either algorithm to smallest possible.*/
while (!(m & 1)) {
    m = m >> 1;
    s--;
}

/* Adjust multiplier for reduction of even divisors. */

s += n;

if (a) {
    printf("; dividend: register other than EAX or memory location\n");
    printf("\n");
    printf("MOV    EAX, 0%08lXh\n", m);
    printf("MUL    dividend\n");
    printf("ADD    EAX, 0%08lXh\n", m);
    printf("ADC    EDX, 0\n");
    if (s) printf("SHR    EDX, %d\n", s);
    printf("\n");
    printf("; quotient now in EDX\n");
}
else {
    printf("; dividend: register other than EAX or memory location\n");
    printf("\n");
    printf("MOV    EAX, 0%08lXh\n", m);
    printf("MUL    dividend\n");
    if (s) printf("SHR    EDX, %d\n", s);
    printf("\n");
    printf("; quotient now in EDX\n");
}

printed_code:

    fprintf(stderr, "\n");
    exit(0);

    return(0);
}

```

Signed Integer Division

The utility `sdiv.exe` was compiled using the following code. The utilities provided in this document are for reference only and are not supported by AMD.

Example

```

/* This program determines the algorithm (a), multiplier (m), and
   shift factor (s) to be used to accomplish *signed* division by
   a constant divisor. Compile with MSVC.
*/

```

```

#include <stdio.h>

typedef unsigned __int64 U64;
typedef unsigned long   U32;

U32 log2(U32 i)
{
    U32 t = 0;
    i = i >> 1;
    while (i) {
        i = i >> 1;
        t++;
    }
    return(t);
}

long e;
U32 res1, res2;
U32 oa, os, om;
U32 d, l, s, m, a, r, t;
U64 m_low, m_high, j, k;

int main(void)
{
    fprintf(stderr, "\n");
    fprintf(stderr, "Signed division by constant\n");
    fprintf(stderr, "=====\n\n");

    fprintf(stderr, "enter divisor: ");
    scanf("%ld", &d);
    fprintf(stderr, "\n");

    e = d;
    d = labs(d);

    if (d == 0) goto printed_code;

    if (e == (-1)) {
        printf("; dividend: register or memory location\n");
        printf("\n");
        printf("NEG    dividend\n");
        printf("\n");
        printf("; quotient replaced dividend\n");
        goto printed_code;
    }
    if (d == 2) {
        printf("; dividend expected in EAX\n");
        printf("\n");
        printf("CMP    EAX, 080000000h\n");
        printf("SBB    EAX, -1\n");
        printf("SAR    EAX, 1\n");
        if (e < 0) printf("NEG    EAX\n");
        printf("\n");
        printf("; quotient now in EAX\n");
    }
}

```

```

    goto printed_code;
}

if (!(d & (d - 1))) {
    printf("; dividend expected in EAX\n");
    printf("\n");
    printf("CDQ\n");
    printf("AND    EDX, 0%081Xh\n", (d-1));
    printf("ADD    EAX, EDX\n");
    if (log2(d)) printf("SAR    EAX, %d\n", log2(d));
    if (e < 0) printf("NEG    EAX\n");
    printf("\n");
    printf("; quotient now in EAX\n");
    goto printed_code;
}

/* Determine algorithm (a), multiplier (m), and shift factor (s) for 32-bit
   signed integer division. Based on: Granlund, T.; Montgomery, P.L.:
   "Division by Invariant Integers using Multiplication". SIGPLAN Notices,
   Vol. 29, June 1994, page 61.
*/

l = log2(d);
j = (((U64)(0x80000000)) % ((U64)(d)));
k = (((U64)(1)) << (32 + 1)) / ((U64)(0x80000000 - j));
m_low = (((U64)(1)) << (32 + 1)) / d;
m_high = (((U64)(1)) << (32 + 1)) + k) / d;

while (((m_low >> 1) < (m_high >> 1)) && (l > 0)) {
    m_low = m_low >> 1;
    m_high = m_high >> 1;
    l = l - 1;
}
m = ((U32)(m_high));
s = 1;
a = (m_high >> 31) ? 1 : 0;

if (a) {
    printf("; dividend: memory location or register other than EAX or EDX\n");
    printf("\n");
    printf("MOV    EAX, 0%08LXh\n", m);
    printf("IMUL   dividend\n");
    printf("MOV    EAX, dividend\n");
    printf("ADD    EDX, EAX\n");
    if (s) printf("SAR    EDX, %d\n", s);
    printf("SHR    EAX, 31\n");
    printf("ADD    EDX, EAX\n");
    if (e < 0) printf("NEG    EDX\n");
    printf("\n");
    printf("; quotient now in EDX\n");
}
else {
    printf("; dividend: memory location of register other than EAX or EDX\n");

```

```
    printf("\n");
    printf("MOV    EAX, 0%08LXh\n", m);
    printf("IMUL   dividend\n");
    printf("MOV    EAX, dividend\n");
    if (s) printf("SAR    EDX, %d\n", s);
    printf("SHR    EAX, 31\n");
    printf("ADD    EDX, EAX\n");
    if (e < 0) printf("NEG    EDX\n");
    printf("\n");
    printf("; quotient now in EDX\n");
}

printed_code:

    fprintf(stderr, "\n");
    exit(0);
}
```

9.7 Optimizing Integer Division

Optimization

For all data types, except in 8-bit division, making the absolute value of the most significant word (in DX/EDX/RDX) of the dividend all 0s for the DIV instruction or all 0s or all 1s for the IDIV instruction lowers the latency of integer division. If this is not possible, then use a smaller data type for integer division.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Integer division latency is dependent on the operand size. These latency numbers could go down even lower, depending on the number of leading zero bits in the absolute value of the dividend. Table 4 provides details about the latency of any particular instance of a DIV/IDIV instruction.

When integer division constitutes a substantial computational load, it may be beneficial to check whether the most significant word of the absolute value of the dividend in DX/EDX/RDX can be set to all 0s for DIV or to all 0s or all 1s for IDIV. If that is not possible, then using a smaller division size will help to lower the latency.

In any case, assembly language output generated by high-level language compilers should be verified that the desired code is generated. When dividing by a constant, if possible, substitute the division with a multiplication. (See “Replacing Division with Multiplication” on page 143 for more details.)

Table 4. DIV/IDIV Latencies

Divisor	Absolute Value of Dividend	Latency	
		DIV	IDIV
8 Bits	Reg	NA	NA
	Mem	NA	NA
16, 32, 64 Bits	0	NA	NA
16 Bits	> 0 and $< 2^{16}$	20 (MSB in bit 0, 1, or 2); or 16 + bit position of the MSB of the dividend (MSB \geq bit 3)	24 + bit position of the MSB of the absolute value of the dividend
32 Bits	> 0 and $< 2^{32}$		
64 Bits	> 0 and $< 2^{64}$		
16 Bits	$\geq 2^{16}$	NA	NA
32 Bits	$\geq 2^{32}$	NA	NA
64 Bits	$\geq 2^{64}$	NA	NA

Note: MSB—Most significant bit.

9.8 Efficient Implementation of Population Count and Leading-Zero Count

Optimization

Use the POPCNT instruction to implement a population count and use LZCNT to perform a leading-zero count operation.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

A population count determines the number of set bits in a bit string. The POPCNT instruction, a new instruction for AMD Family 15h processors, is the preferred way to implement a population count.

A leading-zero count is an operation that counts the number of leading bits in the input operand that are cleared to zero. Counting starts downward from the most significant bit and stops at the highest bit which is one or when the least significant bit is encountered. LZCNT is a new instruction for AMD Family 15h processors that implement this function.

The POPCNT and LZCNT instructions can count the bits in a 32-bit operand in 32-bit mode or a 64-bit operand in 64-bit mode.

Chapter 10 Optimizing with SIMD Instructions

The 64-bit, 128-bit and 256-bit SIMD instructions should be used to encode floating-point and packed integer operations.

- The SIMD instructions use a flat register file rather than the stack register file used by x87 floating-point instructions. This allows arbitrary sequences of operations to map more efficiently to the instruction set.
- AMD Family 15h processors with 128-bit multipliers and adders achieve better throughput using SIMD instructions. (Double precision throughput is 2× and single precision is 4× the throughput of x87.)
- SIMD instructions work well in both 32-bit and 64-bit threads.
- In 64-bit mode, there are twice as many XMM registers available as in 32-bit mode, however, the number of x87 registers is the same in both 32-bit mode and 64-bit mode.

The SIMD instructions provide a theoretical single-precision peak throughput of four additions and four multiplications per clock cycle, whereas x87 instructions can only sustain one addition and one multiplication per clock cycle. The double-precision peak throughput of the SIMD instructions is two additions and two multiplications per clock cycle.

This chapter covers the following topics:

Topic	Page
Ensure All Packed Floating-Point Data are Aligned	168
Explicit Load Instructions	168
Unaligned and Aligned Data Access	169
Use SIMD Instructions to Construct Fast Block-Copy Routines	170
Using SIMD Instructions for Fast Square Roots and Divisions	171
Use XOR Operations to Negate Operands of SIMD Instructions	174
Clearing SIMD Registers with XOR Instructions	175
Finding the Floating-Point Absolute Value of Operands of SIMD Instructions	176
Accumulating Single-Precision Floating-Point Numbers Using SIMD Instructions	176
Complex-Number Arithmetic Using AVX Instructions	178
Optimized 4 X 4 Matrix Multiplication on 4 X 1 Column Vector Routines	181
Floating-Point-to-Integer Conversion	183
Reuse of Dead Registers	184
Floating-Point Scalar Conversions	185
Move/Compute Optimization	186
Using SIMD Instructions for Rounding	188
Using SIMD Instructions for Floating-Point Comparisons	189

10.1 Ensure All Packed Floating-Point Data are Aligned

Optimization

Align all packed floating-point data on 16-byte boundaries.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Misaligned memory accesses reduce the available memory bandwidth and SIMD instructions have shorter latencies when operating on aligned memory operands.

Aligning data on 16-byte boundaries reduces the possibility of stalling floating-point addition and multiplication instructions that are dependent on the load data. See also section 10.3, “Unaligned and Aligned Data Access” on page 169.

10.2 Explicit Load Instructions

Optimization

Use `VMOVSD xmm1, mem64` when loading a scalar floating-point double-precision value from memory.

Use `VMOVSS xmm1, mem32` when loading a scalar floating-point single-precision value from memory.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The `VMOVSD xmm1, mem64` instruction is more efficient than `VMOVLPD xmm1, mem64` on an AMD Family 15h processor, since it modifies the entire XMM register, thus breaking the dependency chain on the high-order bits of the register.

The `VMOVSS xmm1, mem32` instruction zeroes the unaffected remaining bits of the XMM register and breaks any dependency chain. It also assures that the upper half of the XMM register contains a normal floating-point single-precision value.

10.3 Unaligned and Aligned Data Access

Optimization

When data alignment cannot be guaranteed, use `VMOVUPx` or `VMOVDQU` for loads and the `VMOVLPx/VMOVHPx` pair for stores on AMD Family 15h processors.

Otherwise, when data alignment is guaranteed, always use `VMOVAPx` or `VMOVDQA`.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

On AMD Family 15h processors, the `VMOVUPx` and `VMOVDQU` instructions are DirectPath for loads, but VectorPath for stores, therefore the `VMOVLPx / VMOVHPx` pair should be used for stores. `VMOVUPx` or `VMOVDQU` loads can be as fast as `VMOVAPx` or `VMOVDQA` loads when the memory location is 16-byte aligned. The `VMOVUPx` and `VMOVDQU` instructions break dependency chains by changing the entire XMM register when loading data from memory. On the other hand, because both `VMOVLPx` and `VMOVHPx` loads change one half of the XMM register, there is a dependency between each of them and any previous instructions that change any part of the same XMM register.

10.4 Moving Data Between General-Purpose and XMM/YMM Registers

Optimization

When moving data from a GPR to an XMM register, use separate store and load instructions to move the data first from the source register to a temporary location in memory and then from memory into the destination register, taking the memory latency into account when scheduling both stages of the load-store sequence.

When moving data from an XMM register to a general-purpose register, use the `VMOVD` instruction.

Whenever possible, use loads and stores of the same data length. (See 6.3, “Store-to-Load Forwarding Restrictions” on page 98 for more information.)

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

When a GPR is the *source* to the VMOVD instruction, VMOVD is a higher-latency DirectPath Double instruction; compared to the low-latency DirectPath Single instructions used first to store the contents of the GPR to memory and then to load this value into an XMM register.

When a GPR is the *destination* of the VMOVD instruction, VMOVD is a DirectPath Single instruction.

10.5 Use SIMD Instructions to Construct Fast Block-Copy Routines

Optimization

Use XMM registers instead of general purpose registers to copy blocks of data that reside in cache.

Application

This optimization applies to:

- 64-bit software

Rationale

SIMD loads and stores can read and write 16 bytes in a single clock cycle, while a SIMD store can write 16 bytes in two cycles. VMOVDQU can safely access 16-byte data regardless of alignment, with performance equal to VMOVDQA when data is actually 16-byte aligned, so use VMOVDQU and align the destination and/or the source to 16-byte boundaries when possible.

Example

The following code illustrates an implementation of an optimized memory block copy using 128 bit XMM registers. This code uses a partially unrolled loop with an unroll factor of two to hide the execution latencies of the pointer/counter arithmetic and branch instructions.

```
; rsi = ptr to destination, must be 16-byte aligned
; rdi = ptr to source, must be 16-byte aligned
; rdx = count, make sure it's at least 32 bytes and
;           is a function of sizeof(dest type) * len
```

```

    shrq rdx, 5 ; we move 32 bytes per loop
    jz SSE_done
    align 16 ; align loop top for best performance
SSE_loop:
    vmovdqa xmm0, [rdi]
    vmovdqa xmm1, [rdi + 16]
    add rdi, 32
    vmovdqa [rsi], xmm0,
    vmovdqa [rsi + 16], xmm1,
    add rsi, 32
    dec rdx
    jnz SSE_loop
SSE_done:
    ; (move any residual bytes)

```

10.6 Using SIMD Instructions for Fast Square Roots and Divisions

Optimization

Use SIMD vectorized square root (VSQRTSS/VSQRTPS) and reciprocal (VRCPSS/VRCPPS) instructions to calculate square roots and divisions of single-precision numbers.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The calculation of reciprocal square root and reciprocation of single-precision numbers are often used in multimedia applications. These SIMD instructions can be used for such operations when a slight inaccuracy is acceptable.

Although these instructions return their results with a maximum error of 2^{-11} , they can be used with the Newton-Raphson method to obtain more accurate results.

For square roots accurate to 2.5 ULPs, the following algorithm is obtained after one Newton-Raphson iteration:

$$y = 0.5 * a * x * (3.0 - a * x * x)$$

Where x is the initial approximation of the reciprocal of the square root of a and y , the square root of a .

For divisions accurate to 1.5 ULPs, the following algorithm is obtained after one Newton-Raphson iteration:

$$y = a * x * (2.0 - b * x)$$

Where x is the initial approximation of the reciprocal of b and y , the quotient of a divided by b .

Although more Newton-Raphson iterations could be used to increase accuracy, the execution time would be longer than the equivalent instructions. This implementation of the Newton-Raphson technique is not 100% compliant to the IEEE-754 specification, but its results are acceptable in most applications.

Example

The following functions calculate the square root:

```
#include <xmmintrin.h>
/* nr_sqrtf: return scalar square root accurate to 2.5ulps.

   This approximation assumes finite math; never returns denormals, but zero;
   does not return the expected values after C89;
   is not compliant with IEEE754 semantics.

   Note: AVX code can be generated using compiler flags without modification to
   this example. */

float nr_sqrtf (float a)
{
    __m128 x0, x1, x2, x3, x4, x5, m;
    float y;

    m = _mm_cmpneq_ss (_mm_set_ss (a), _mm_setzero_ps ()); // m = (a != 0.0? T: F)

    x0 = _mm_rsqrt_ss (_mm_set_ss (a)); // x0 = initial estimate
    x1 = _mm_and_ps (m, x0); // x1 = m & x0
    x2 = _mm_mul_ss (_mm_set_ss (a), x1); // x2 = a * x1
    x3 = _mm_mul_ss (_mm_set_ss (0.5F), x2); // x3 = 0.5 * x2
    x4 = _mm_mul_ss (x1, x2); // x4 = x1 * x2
    x5 = _mm_sub_ss (_mm_set_ss (3.0F), x4); // x5 = 3.0 - x4

    _mm_store_ss (&y, _mm_mul_ss (x3, x5)); // y = x3 * x5
    return (y); // y = sqrtf (a)
}

/* nr_sqrtvf: return vector square root accurate to 2.5ulps.

   This approximation assumes finite math; never returns denormals, but zero;
   does not return the expected values after C89;
   is not compliant with IEEE754 semantics.

   Note: AVX code can be generated using compiler flags without modification to
   this example. */

__m128 nr_sqrtvf (__m128 a)
```

```

{
  __m128 x0, x1, x2, x3, x4, x5, m, y;

  m = _mm_cmpneq_ps (a, _mm_setzero_ps ()); // m = (a != 0.0? T: F)

  x0 = _mm_rsqrt_ps (a); // x0 = initial estimate
  x1 = _mm_and_ps (m, x0); // x1 = m & x0
  x2 = _mm_mul_ps (a, x1); // x2 = a * x1
  x3 = _mm_mul_ps (_mm_set1_ps (0.5F), x2); // x3 = 0.5 * x2
  x4 = _mm_mul_ps (x1, x2); // x4 = x1 * x2
  x5 = _mm_sub_ps (_mm_set1_ps (3.0F), x4); // x5 = 3.0 - x4

  y = _mm_mul_ps (x3, x5); // y = x3 * x5
  return (y); // y = sqrtf (a)
}

```

These functions return the quotient:

```
#include <xmmintrin.h>
```

```
/* nr_divf: return scalar quotient accurate to 1.5ulps.
```

```

    This approximation assumes finite math; never returns denormals, but zero;
    does not return the expected values after C89;
    is not compliant with IEEE754 semantics.

```

```

    Note: AVX code can be generated using compiler flags without modification to
    this example. */

```

```

float nr_divf (float a, float b)
{
  __m128 x0, x1, x2, x3;
  float y;

  x0 = _mm_rcp_ss (_mm_set_ss (b)); // x0 = initial estimate
  x1 = _mm_mul_ss (_mm_set_ss (a), x0); // x1 = a * x0
  x2 = _mm_mul_ss (_mm_set_ss (b), x0); // x2 = b * x0
  x3 = _mm_sub_ss (_mm_set_ss (2.0F), x2); // x3 = 2 - x2

  _mm_store_ss (&y, _mm_mul_ss (x1, x3)); // y = x1 * x3
  return (y); // y = a / b
}

```

```
/* nr_divvf: return vector quotient accurate to 1.5ulps.
```

```

    This approximation assumes finite math; never returns denormals, but zero;
    does not return the expected values after C89;
    is not compliant with IEEE754 semantics.

```

```

    Note: AVX code can be generated using compiler flags without modification to
    this example. */

```

```
__m128 nr_divf (__m128 a, __m128 b)
```

```

{
  __m128 x0, x1, x2, x3, y;

  x0 = _mm_rcp_ps (b);           // x0 = initial estimate
  x1 = _mm_mul_ps (a, x0);       // x1 = a * x0
  x2 = _mm_mul_ps (b, x0);       // x2 = b * x0
  x3 = _mm_sub_ps (_mm_set1_ps (2.0F), x2); // x3 = 2 - x2

  y = _mm_mul_ps (x1, x3);       // y = x1 * x3
  return (y);                    // y = a / b
}

```

10.7 Use XOR Operations to Negate Operands of SIMD Instructions

Optimization

For AMD Family 15h processors, use instructions that perform XOR operations (VXORPS, and VXORPD) instead of multiplication instructions to change the sign bits of operands of SIMD instructions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

On the AMD Family 15h processors, using XOR-type instructions allows for more parallelism, since these instructions can execute in either FPMA pips 2 or 3 of the floating-point unit. Also, the latency of the VMULPS or VMULPD instruction is longer than the latency of VXORPS or VXORPD (see Appendix B, “Instruction Latencies”).

Single Precision

This example shows how to toggle the sign bit of four floating-point values using single-precision SIMD instructions:

```

signmask_DQ 8000000080000000h,8000000080000000h
vxorps xmm0, xmm0, [signmask] ; Toggle sign bits of all four floats.

```

Double Precision

The following example shows how to toggle the sign bit of two doubles using double-precision AVX instructions:

```
signmask DQ 8000000000000000h,8000000000000000h
v xorpd xmm0, xmm0, [signmask] ; Toggle sign bit of both doubles.
```

10.8 Clearing SIMD Registers with XOR Instructions

Optimization

Use instructions that perform XOR operations (VPXOR, VXORPS, and VXORPD) to clear all the bits in XMM/YMM registers.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The VXORPS and VXORPD instructions are more efficient than loading a zero value into an XMM register from memory and then storing it (see Appendix B, “Instruction Latencies”). In addition, the processor “knows” that the VXORPS and VXORPD instructions that use the same register for both source and destination do not have a real dependency on the previous contents of the register, and thus, do not have to wait before completing.

Examples

The following examples illustrate how to clear the bits in a register using the different exclusive-OR instructions:

```
; AVX packed single precision:
v xorps xmm0, xmm0, xmm0 ; Clear the XMM0 register.

; AVX packed double precision:
v xorpd xmm0, xmm0, xmm0 ; Clear the XMM0 register.
```

10.9 Finding the Floating-Point Absolute Value of Operands of SIMD Instructions

Optimization

Use instructions that perform AND operations ((V)ANDPS, and (V)ANDPD) to determine the absolute value of floating-point operands of SIMD instructions.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Examples

The following examples illustrate how to clear the sign bits. See Appendix C for latencies of the ANDPS and ANDPD instructions:

```
; SSE
absmask DQ 7FFFFFFF7FFFFFFFh,7FFFFFFF7FFFFFFFh
andps xmm0, [absmask] ; Clear the sign bits of all four floats in XMM0.
; SSE2
absmask DQ 7FFFFFFFFFFFFFFFh,7FFFFFFFFFFFFFFFh
andpd xmm0, [absmask] ; Clear the sign bits of both doubles in XMM0.
```

10.10 Accumulating Single-Precision Floating-Point Numbers Using SIMD Instructions

Optimization

Careful selection of SIMD instructions based on efficient data organization can lead to more economical code.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

SIMD vectorized multiplication and addition instructions are useful for carrying out such operations as complex-number multiplication, 4 x 4 matrix multiplication, and dot products.

Examples

The following example uses SIMD instructions. Four floating-point values are loaded into four XMM registers, XMM4–XMM7. These values are then rearranged and added, so as to accumulate the sum of each XMM register into a float in XMM1.

```

;-----
; The instructions below take the 4 floats in each XMM register below:
; xmm4 = [d,c,b,a]
; xmm5 = [D,C,B,A]
; xmm6 = [h,g,f,e]
; xmm7 = [H,G,F,E]
;
; and arranges them to look like:
; xmm4 = [E,e,A,a]
; xmm1 = [F,f,B,b]
; xmm2 = [G,g,C,c]
; xmm3 = [H,h,D,d]
vmovaps xmm3, xmm4 ; xmm3 | [d,c,b,a]
vmovaps xmm0, xmm5 ; xmm0 | [D,C,B,A]
vunpcklps xmm4, xmm4, xmm6 ; xmm4 | [f,b,e,a]
vunpckhps xmm3, xmm3, xmm6 ; xmm3 | [h,d,g,c]
vmovaps xmm1, xmm4 ; xmm1 | [f,b,e,a]
vmovaps xmm2, xmm3 ; xmm2 | [h,d,g,c]
vunpcklps xmm5, xmm5, xmm7 ; xmm5 | [F,B,E,A]
vunpckhps xmm0, xmm0, xmm7 ; xmm0 | [H,D,G,C]
vunpcklps xmm4, xmm4, xmm5 ; xmm4 | [E,e,A,a]
vunpckhps xmm1, xmm1, xmm5 ; xmm1 | [F,f,B,b]
vunpcklps xmm3, xmm3, xmm0 ; xmm3 | [G,g,C,c]
vunpckhps xmm2, xmm2, xmm0 ; xmm2 | [H,h,D,d]
; Now if we compute the sum of these registers, we get the dot-product
; of the first row of A with vector X:
;
; a+b+c+d
;
; in the lower DWORD of the resultant XMM register. The dot-product of the
; second row is stored in the second DWORD and so on, such that:
;
; xmm1 = [V+X+Y+Z,v+x+y+z,A+B+C+D,a+b+c+d]
vaddps xmm1, xmm1, xmm4 ; xmm1 | [E+F,e+f,A+B,a+b]
vaddps xmm3, xmm3, xmm2 ; xmm3 | [G+H,g+h,C+D,c+d]
vaddps xmm1, xmm1, xmm3 ; xmm1 | [E+F+G+H,e+f+g+h,A+B+C+D,a+b+c+d]

```

10.11 Complex-Number Arithmetic Using AVX Instructions

Optimization

Use vectorizing AVX instructions to perform complex number calculations.

Application

This optimization applies to:

- 64-bit software

Rationale

Complex numbers have a “real” part and an “imaginary” part (where the imaginary part is denoted by the letter i). For example, the complex number $z1$ might have a real part equal to 4 and an imaginary part equal to 3, written as $4 + 3i$. Multiplying and adding complex numbers is an integral part of many areas of mathematics. Complex number addition is illustrated here using two complex numbers, $z1$ ($4 + 3i$) and $z2$ ($5 + 2i$):

$$z1 + z2 = (4 + 3i) + (5 + 2i) = [4+5] + [3+2]i = 9 + 5i$$

or:

```
sum.real = z1.real + z2.real
sum.imag = z1.imag + z2.imag
```

Complex number multiplication is illustrated below using the same two complex numbers:

$$z1 \times z2 = (4 + 3i)(5 + 2i) = [4 \times 5 - 3 \times 2] + [3 \times 5 + 4 \times 2]i = 14 + 23i$$

or:

```
product.real = z1.real \times z2.real - z1.imag \times z2.imag
product.imag = z1.real \times z2.imag + z1.imag \times z2.real
```

Complex numbers can be stored as streams of two-element vectors, the two elements being the real and imaginary parts of the complex numbers. Addition of complex numbers can be achieved using vectorizing SIMD instructions, such as VADDPS and VADDPD. Multiplication of complex numbers is more involved, but AVX and FMA4 instructions are available to perform exactly the operations required.

From the formulas for multiplication, the real and imaginary parts of one of the numbers must be interchanged, and, additionally, the products must be positively or negatively accumulated depending upon whether you are computing the imaginary or real portion of the product.

The following functions use AVX and FMA4 instructions to illustrate complex multiplication of streams of complex numbers $x[]$ and $y[]$ stored in a product stream $prod[]$. For these examples,

assume that the sizes of $x[]$ and $y[]$ are even multiples of eight. If this assumption were not true, extra code would be required to compute the remaining products. Another assumption is that $x[]$ and $y[]$ are aligned on 16 byte boundaries, as discussed in 6.2, “Natural Alignment of Data Objects” on page 97 and 10.3, “Unaligned and Aligned Data Access” on page 169.

Example

Complex Multiplication of Streams of Complex Numbers using AVX and FMA4 Instructions

```
; void cmplx_multiply_avx(float *x, float *y, int num_cmplx_elem, float *prod);
;
; TO ASSEMBLE INTO *.obj DO THE FOLLOWING:
; ml64.exe -c cmplx_multiply_avx.asm
;
;
; define local variable storage offsets
save_rdi equ 00h ;qword
save_rsi equ 08h ;qword
stack_size equ 018h
TEXT SEGMENT page 'CODE'
PUBLIC cmplx_multiply_avx
cmplx_multiply_avx: proc frame
;=====
; INSTRUCTIONS BELOW SAVE THE REGISTER STATE WITH WHICH THIS ROUTINE WAS ENTERED
; REGISTERS RSI, and RSI ARE CONSIDERED VOLATILE AND ASSUMED TO BE CHANGED
; WHILE THE REGISTERS BELOW MUST BE PRESERVED IF THE USER IS CHANGING THEM
;
sub rsp,stack_size
mov QWORD PTR [rsp+save_rdi],rdi ; save rdi
mov QWORD PTR [rsp+save_rsi],rsi ; save rsi
;=====
; Parameters passed into routine according to the Microsoft AMD64 ABI:
; rcx = ->x
; rdx = ->y
; r8d = num_cmplx_elem
; r9 = ->prod
;=====
mov rsi, rcx ; rsi = ->x
mov rdi, rdx ; rdi = ->y
mov rcx, r8d ; rcx = num_cmplx_elem (zero extends the destination register)
```

```

;=====
; THE 6 ASM LINES BELOW OFFSET THE ADDRESS TO THE ARRAYS x[] AND y[] SUCH
; THAT THEY CAN BE ACCESSED IN THE MOST EFFICIENT MANNER AS ILLUSTRATED
; BELOW IN THE LOOP mult8cplxnum_loop WITH THE MINIMUM NUMBER OF
; ADDRESS INCREMENTS
;=====
mov r8, rcx                ; rdx = num_cplx_elem
neg rcx                   ; rcx = -num_cplx_elem
imul r8, 8                ; r8 = 8 * num_cplx_elem = # bytes in x[]
                           ; and y[] to multiply
add rsi, r8               ; rsi = -> to last element of x[] to multiply
add rdi, r8               ; rdi = -> to last element of y[] to multiply
add r9, r8                ; r9 = -> end of prod[] to calculate
;=====
; THIS LOOP MULTIPLIES 4 COMPLEX #s FROM "x[]" UPON 4 COMPLEX #s FROM "y[]"
; AND RETURNS THE PRODUCT IN "prod[]".
;=====
ALIGN 32                   ; Align loop top address to a 32-byte
                           ; boundary.

four_cplx_prod_loop:
vmovaps ymm0, YMMWORD PTR [rsi+rcx*8] ; ymm0=[x3i,x3r,x2i,x2r,x1i,x1r,x0i,x0r]
vmovaps ymm1, YMMWORD PTR [rdi+rcx*8] ; ymm1=[y3i,y3r,y2i,y2r,y1i,y1r,y0i,y0r]
vmovshdup ymm2, ymm0        ; ymm2=[x3i,x3i,x2i,x2i,x1i,x1i,x0i,x0i]
vmovsldup ymm0, ymm0       ; ymm0=[x3r,x3r,x2r,x2r,x1r,x1r,x0r,x0r]
vshufps ymm3, ymm1, ymm1, 0b1h ; ymm3=[y3r,y3i,y2r,y2i,y1r,y1i,y0r,y0i]
vmulps ymm2, ymm3, ymm2    ; ymm2=[x3i*y3r,x3i*y3i,x2i*y2r,x2i*y2i,
; x1i*y1r,x1i*y1i,x0i*y0r,x0i*y0i]
vfmaddsubps ymm0, ymm0, ymm1, ymm2 ; ymm0=[x3r*y3i+x3i*y3r,x3r*y3r-x3i*y3i,
; x2r*y2i+x2i*y2r,x2r*y2r-x2i*y2i,
; x1r*y1i+x1i*y1r,x1r*y1r-x1i*y1i,
; x0r*y0i+x0i*y0r,x0r*y0r-x0i*y0i]

vmovntps YMMWORD PTR [r9+rcx*8], ymm0 ; Stream ymm0 to destination
add rcx, 4 ; RCX = RCX +4
jnz four_cplx_prod_loop
;=====
; INSTRUCTIONS BELOW RESTORE THE REGISTER STATE WITH WHICH THIS ROUTINE WAS
; ENTERED
; REGISTERS RDI, RSI ARE CONSIDERED VOLATILE AND ASSUMED TO BE CHANGED
; WHILE THE REGISTERS BELOW MUST BE PRESERVED IF THE USER IS CHANGING THEM
mov rdi, QWORD PTR [rsp+save_rdi] ; restore rdi
mov rsi, QWORD PTR [rsp+save_rsi] ; restore rsi
add rsp, stack_size
ret
cplx_multiply_avx endp
TEXT ENDS
END

```

This example takes advantage of the VFMADDSSUBPS FMA4 instruction to perform one of the multiplication steps, combining the subtraction for the real terms ($x.r*y.r - x.i*y.i$) and the addition for imaginary terms ($x.r*y.i + x.i*y.r$).

The example also uses MOVNTPS instructions—nontemporal writes to memory that stream data to main memory. These instructions increase throughput to memory and make more efficient use of the bandwidth provided by the processor and memory controller. Nontemporal writes, such as MOVNTPS, and MOVNTDQ, should only be used on data that is not going to be accessed again in

the near future. This is described in more detail in 6.5, “Prefetch and Streaming Instructions” on page 105.

10.12 Optimized 4 X 4 Matrix Multiplication on 4 X 1 Column Vector Routines

Optimization

Transpose the rotation matrix to eliminate the need to accumulate floating-point values in an XMM register.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

The multiplication of a 4×4 matrix with a 4×1 vector is commonly used in 3-D graphics for geometric transformation (translating, scaling, rotating, and applying perspective to 3-D points represented in homogeneous coordinates). Efficiency in single-precision matrix multiplication can be enhanced by use of SIMD instructions to increase throughput, but there are other general optimizations that can be implemented to further increase performance. The first optimization is the transposition of the rotation matrix such that column n of the matrix becomes row n and row m becomes column m . There are no SIMD instructions that accumulate the floats and doubles in a single XMM register; for this reason, the matrix must be transposed. If the rotation matrix is not transposed, then the dot-product of a row of the matrix with a column vector necessitates the accumulation of the four floating-point values in an XMM register. The multiplication on the column vector is illustrated here

$$\text{tr}(R) \times v = \text{tr} \begin{vmatrix} r_{00} & r_{01} & r_{02} & r_{03} \\ r_{10} & r_{11} & r_{12} & r_{13} \\ r_{20} & r_{21} & r_{22} & r_{23} \\ r_{30} & r_{31} & r_{32} & r_{33} \end{vmatrix} \times v = \begin{vmatrix} r_{00} & r_{10} & r_{20} & r_{30} \\ r_{01} & r_{11} & r_{21} & r_{31} \\ r_{02} & r_{12} & r_{22} & r_{32} \\ r_{03} & r_{13} & r_{23} & r_{33} \end{vmatrix} \times \begin{vmatrix} v_0 \\ v_1 \\ v_2 \\ v_3 \end{vmatrix} = \begin{vmatrix} v'_0 \\ v'_1 \\ v'_2 \\ v'_3 \end{vmatrix}$$

$$\begin{array}{l} \begin{array}{cccc} & \text{Step 0} & \text{Step 1} & \text{Step 2} & \text{Step 3} \\ |v'_0| & |r_{00} \times v_0| & |r_{01} \times v_1| + |r_{02} \times v_2| + |r_{03} \times v_3| \\ |v'_1| = & |r_{10} \times v_0| + |r_{11} \times v_1| + |r_{12} \times v_2| + |r_{13} \times v_3| \\ |v'_2| & |r_{20} \times v_0| + |r_{21} \times v_1| + |r_{22} \times v_2| + |r_{23} \times v_3| \\ |v'_3| & |r_{30} \times v_0| + |r_{31} \times v_1| + |r_{32} \times v_2| + |r_{33} \times v_3| \end{array} \end{array}$$

In each step above, the elements of the rotation matrix can be loaded into an XMM register with the MOVAPS instruction, assuming the rotation matrix begins at a 16-byte-aligned memory location. Transposition of the rotation matrix eliminates the need to accumulate the floating-point values in an

XMM register, but it does require the duplication of the elements of the 4×1 column vector V in all four floating-point values of the XMM register in each step above. The following example shows a SIMD function that performs 4×4 matrix multiplication upon a stream of `num_vertices_to_rotate` vertices.

Example

4 X 4 Matrix Multiplication (SIMD)

```
include listing.inc

INCLUDELIB LIBCMT
INCLUDELIB OLDNAMES

PUBLIC  _matrix_x_vector_simd

; Function compile flags: /Odtp

_TEXT SEGMENT
_matrix_x_vector_simd PROC

; File d:\storage\amd64\bd_swog\code\chapter9.13\test\chapter9.13_test.c
;=====
; Parameters passed into routine:
; rcx = ->trR
; rdx = ->v
; r8 = num_vertices_to_rotate
; r9 = ->rotv
;=====
;
;=====
; THE 4 ASM LINES BELOW LOAD THE FUNCTION'S ARGUMENTS INTO GENERAL-PURPOSE
; REGISTERS (GPRS)
; rcx = address of Transposed Rotation Matrix
; rdx = address of vertices to rotate
; r8 = # of vertices to rotate
; r9 = address of rotated vertices
;=====
mov r10, r8                ; R10 = num_vertices_to_rotate
shl r10, 4                 ; R10 = 16*num_vertices_to_rotate
shl r8, 1                  ; R8 = # quadwords of vertices to rotate
add rdx, r10               ; RDX = -> end of "v"
add r9, r10                ; R9 = -> end of "rotv"
neg r8                     ; R8 = -# quadwords of vertices to rotate
;=====
; THE 4 ASM LINES BELOW LOAD THE TRANSPOSED ROTATION MATRIX "R" INTO XMM0-XMM3
; IN THE FOLLOWING MANNER:
; xmm0 = column 0 of "R" or row 0 of "R" transpose
; xmm1 = column 1 of "R" or row 1 of "R" transpose
; xmm2 = column 2 of "R" or row 2 of "R" transpose
; xmm3 = column 3 of "R" or row 3 of "R" transpose
;=====
vzeroupper
vmovaps xmm0, [rcx]        ; XMM0 = [R30,R20,R10,R00]
vmovaps xmm1, [rcx+16]    ; XMM1 = [R31,R21,R11,R01]
```

```

vmovaps xmm2, [rcx+32]          ; XMM2 = [R32,R22,R12,R02]
vmovaps xmm3, [rcx+48]          ; XMM3 = [R33,R23,R13,R03]
;=====
; THIS LOOP ROTATES "num_vertices_to_rotate" VERTICES BY THE TRANSPOSED
; ROTATION MATRIX "R" PASSED INTO THE ROUTINE AND STORES THE ROTATED
; VERTICES TO "rotv".
;=====
rotate_vertices_loop:
vxorps   xmm8, xmm8, xmm8
vmovlps  xmm4, xmm8, QWORD PTR [rdx+8*r8] ; XMM4=[, ,v1,v0]
vmovlps  xmm6, xmm8, QWORD PTR [rdx+8*r8+8] ; XMM6=[, ,v3,v2]
vunpcklps xmm4, xmm4, xmm4 ; XMM4=[v1,v1,v0,v0]
vunpcklps xmm6, xmm6, xmm6 ; XMM6=[v3,v3,v2,v2]
vmovhlps xmm5, xmm4, xmm4 ; XMM5=[, ,v1,v1]
vmovhlps xmm7, xmm6, xmm6 ; XMM7=[, ,v3,v3]
vmovlhps xmm4, xmm4, xmm4 ; XMM4=[v0,v0,v0,v0]
vmulps   xmm4, xmm4, xmm0 ; XMM4=[R30*v0,R20*v0,R10*v0,R00*v0]
vmovlhps xmm5, xmm5, xmm5 ; XMM5=[v1,v1,v1,v1]
vfmaddps xmm4, xmm5, xmm1, xmm4 ; XMM4=[R30*v0+R31*v1,R20*v0+R21*v1,
; R10*v0+R11*v1,R00*v0+R01*v1]

vmovlhps xmm6, xmm6, xmm6 ; XMM6=[v2,v2,v2,v2]
vmulps   xmm6, xmm6, xmm2 ; XMM6=[R32*v2,R22*v2,R12*v2,R02*v2]
vmovlhps xmm7, xmm7, xmm7 ; XMM7=[v3,v3,v3,v3]
vfmaddps xmm6, xmm7, xmm3, xmm6 ; XMM6=[R32*v2+R33*v3,R22*v2+R23*v3,
; R12*v2+R13*v3,R02*v2+R03*v3]

vaddps   xmm4, xmm4, xmm6 ; XMM4=New rotated vertex
movntps  XMMWORD PTR[r9+8*r8], xmm4 ; Store rotated vertex to rotv.
add r8, 2 ; Decrement the # of QWORDS to rotate by 2.
jnz rotate_vertices_loop
sfence ; Finish all memory writes.
;=====
ret 0
_matrix_x_vector_simd ENDP

_TEXT ENDS
END
END

```

To greatly enhance performance, the previous function can perform the matrix multiplication not only on one four-column vector, but on many. Creating a separate function to transform a single vertex and repeatedly calling the function is prohibitively expensive because of the overhead in pushing and popping registers from the stack. This applies to routines that negate a single vector, nullify a single vector, and add two vectors.

10.13 Floating-Point-to-Integer Conversion

Optimization

Floating-point-to-integer conversion in C and C++ requires the use of truncation. Use one of the instructions from VCVTTSS2SI, VCVTTSD2SI to convert a floating-point number to integer when

truncation is required. See the *AMD64 Architecture Programmer's, Volume 4: 128-Bit Media Instructions*, order# 26568, for details.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

These instructions provide the fastest means by which to convert floating-point types to integers in AMD Family 15h processors.

10.14 Reuse of Dead Registers

Optimization

On AMD Family 15h processors, when it is necessary to save the contents of a register that is a single-precision floating-point scalar to another unused (or *dead*) register, use `VMOVAPS dest, src` instead of `VMOVSS dest, src`.

When saving a register that is a double-precision floating-point scalar to another register, where the contents are unknown, then use `VMOVAPD dest, src` instead of `VMOVSD dest, src`.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

On the AMD Family 15h processors, the `VMOVSS dest, src` instruction takes additional time to execute if any of the upper three fields of `dest` is a denormal. Additionally, the `VMOVSS dest, src` instruction has a dependency on previous instructions that change `dest`, either partially or in full, and the `VMOVAPS dest, src` instruction breaks such dependency chains by changing `dest` as a whole.

The `VMOVSD dest, src` instruction also takes additional time to execute, if the previous value in `xmm1` is a denormal. Moreover, the `VMOVSD dest, src` instruction has a dependency on previous instructions that change `dest` either partially or in full. On the other hand, the `VMOVAPD dest, src` instruction breaks such dependency chains by writing to all of `dest`.

10.15 Floating-Point Scalar Conversions

Optimization

Use the recommended instruction sequences given in Table 5 and Table 6 to convert integer data to floating-point data.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

On AMD Family 15h processors, some SIMD conversion instructions are VectorPath and/or add a false dependency on previous instructions that change the same destination register. In the cases for which there are alternatives in Tables 5 and 6, these instruction sequences use DirectPath instructions and provide better performance. (All recommendations apply to both 32-bit and 64-bit software, unless stated otherwise.)

Several instructions may be required to perform some conversions from unsigned integer to floating-point, due to the lack of a suitable conversion instruction, therefore signed integers should be favored when converting to floating-point.

Table 5. Single-Precision Floating-Point Scalar Conversion

Conversion	From a Register	From Memory
32-Bit Signed Integer to Single-Precision	<code>vmovd xmm, reg32</code> <code>vcvtdq2ps xmm, xmm</code>	<code>vmovd xmm, mem32</code> <code>vcvtdq2ps xmm, xmm</code>
32-Bit Unsigned Integer to Single-Precision	64-bit software: <code>vxorps xmm, xmm, xmm²</code> <code>mov mem64, reg64</code> <code>vcvtsi2ss xmm, xmm, mem64</code>	64-bit software: <code>vxorps xmm, xmm, xmm²</code> <code>vcvtsi2ss xmm, xmm, mem64</code>
64-Bit Signed Integer to Single-Precision	64-bit software: <code>vxorps xmm, xmm, xmm²</code> <code>mov mem64, reg64</code> <code>vcvtsi2ss xmm, xmm, mem64</code>	64-bit software: <code>vxorps xmm, xmm, xmm²</code> <code>vcvtsi2ss xmm, xmm, mem64</code>
Double-Precision to Single-Precision	<code>vunpcklpd xmm2, xmm2, xmm2¹</code> <code>vcvtpd2ps xmm1, xmm2</code>	<code>vmovsd xmm, mem64</code> <code>vcvtpd2ps xmm, xmm</code>
Notes:		
<ol style="list-style-type: none"> 1. If the contents of [127:64] of <code>xmm2</code> is known to be a normal number, this instruction can be omitted. 2. This avoids a merge dependency for contents of [127:32] of <code>xmm</code> as a result of a previous long latency instruction that has written to contents of [127:0] of <code>xmm</code>. 		

Table 6. Double-Precision Floating-Point Scalar Conversion

Conversion	From a Register	From Memory
32-Bit Signed Integer to Double-Precision	<code>vmovd xmm, reg32</code> <code>vcvtdq2pd xmm, xmm</code>	<code>vmovd xmm, mem32</code> <code>vcvtdq2pd xmm, xmm</code>
32-Bit Unsigned Integer to Double-Precision	64-bit software: <code>vxorpd xmm, xmm, xmm²</code> <code>mov mem64, reg64</code> <code>vcvtsi2sd xmm, xmm, mem64</code>	64-bit software: <code>vxorpd xmm, xmm, xmm²</code> <code>vcvtsi2sd xmm, xmm, mem64</code>
64-Bit Signed Integer to Double-Precision	64-bit software: <code>vxorpd xmm, xmm, xmm²</code> <code>mov mem64, reg64</code> <code>vcvtsi2sd xmm, xmm, mem64</code>	64-bit software: <code>vxorpd xmm, xmm, xmm²</code> <code>vcvtsi2sd xmm, xmm, mem64</code>
Single-Precision to Double-Precision	<code>unpcklps xmm2, xmm2¹</code> <code>cvtps2pd xmm1, xmm2</code>	<code>vmovss xmm, mem32</code> <code>cvtps2pd xmm, xmm</code>
Notes:		
<ol style="list-style-type: none"> 1. If the contents of [63:32] of <code>xmm2</code> is known to be a normal number, this instruction can be omitted. 2. This avoids a merge dependency for contents of [127:32] of <code>xmm</code> as a result of a previous long latency instruction that has written to contents of [127:0] of <code>xmm</code>. 		

In loops which involve the conversion of a scalar operand, it is helpful to clear the targeted destination register to zeros prior to the conversion. When performing scalar conversions require merging the floating point unit first checks whether the upper 64 or 96 bits of the destination register are set. If any of these bits are set to 1, then no merge occurs and merge dependencies are eliminated.

10.16 Move/Compute Optimization

Optimization

The latency of certain XMM move instructions that provide an input operand to a subsequent compute instruction can be hidden in many cases. The hardware will only recognize such optimization opportunities relative to the most recent move instruction, so instructions must be ordered with this in mind. Use these specific move instructions, appropriately ordered as described below, wherever possible.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Carefully ordering move instructions with respect to subsequent computations having a source XMM that is identical to the destination XMM of the move increases throughput by effectively canceling the latency cost of move operations associated with FPU computations.

This hardware optimization is designed to work with the VMOVAPD, VMOVAPS, VMOVDQA, VMOVDQU, VMOVUPD, and VMOVUPS instructions. Other SIMD move instructions cause a two-cycle delay in executing the dependent compute instruction. If at all possible, every effort should be made to use move instructions that the processor hardware can optimize.

Move-Compute Execution Rules

Move instructions themselves cannot short-circuit the latency of a prior move instruction on which they are dependent; this can only be done by compute instructions.

Only one move instruction's output-to-input mapping can be recognized at a time and eliminated for dependent compute instructions. So a given move-compute optimization opportunity exists only until one of the following occurs:

- The destination or source XMM of the currently-recognized move instruction is overwritten by a younger instruction.
- Another eligible move instruction is encountered.
- A redirect occurs (such as from a mispredicted branch).

Example 1

Optimal sequence—move instructions are recognized as move-compute eligible:

```
vmovapd xmm2, xmm1; hardware recognizes xmm2 = xmm1
vmulsd xmm2, xmm2, xmm2; executed as xmm2 = xmm1 * xmm1
vmovapd xmm4, xmm3; hardware recognizes xmm4 = xmm3
vmulsd xmm4, xmm4, xmm4; executed as xmm4 = xmm3 * xmm3
```

Less optimal sequence due to poor instruction ordering:

```
vmovapd xmm2, xmm1; hardware recognizes xmm2 = xmm1
vmovapd xmm4, xmm3; hardware loses sight of xmm2 = xmm1
vmulsd xmm2, xmm2, xmm2; not optimized by hardware -- stalls on movapd result
vmulsd xmm4, xmm4, xmm4; executed as xmm4 = xmm3 * xmm3 -- no stall
```

Non-interacting instructions, including ineligible move instructions, may appear anywhere in these sequences without affecting the hardware optimization opportunities, although the biggest gain occurs when a dependent compute instruction can issue at the same time as the move.

Example 2

In this example, an intervening VSUBPD instruction does not affect the destination register of the previous VMOVAPD instruction, which the optimization is able to cut out of the critical path for the subsequent VADDPD instruction.

```
vaddpd xmm2, xmm2, xmm1
vmovapd xmm3, xmm2 ; Hardware recognizes xmm3 = xmm2.
vsubpd xmm4, xmm4, xmm0 ; Optimization is not applicable, since neither
                        ; source operand matches the previous destination
                        ; of the previous vmovapd.
vaddpd xmm4, xmm4, xmm3 ; Executes as xmm4 = xmm4 + xmm2 to cut VMOVAPD
                        ; out of critical path.
vaddpd xmm3, xmm3, xmm2 ; Executes as xmm3 = xmm2 + xmm2, then cancels further
                        ; use of this optimization because xmm3 is overwritten
```

Example 3

Here, the interfering instruction blocks the hardware optimization opportunity.

```
vaddpd xmm2, xmm2, xmm1
vmovapd xmm3, xmm2 ; Hardware recognizes that xmm3 = xmm2.
vsubpd xmm2, xmm2, xmm0 ; xmm3 no longer equal to xmm2, hence move-compute
                        ; optimization not possible
vaddpd xmm4, xmm4, xmm3 ; Optimization not enabled due to
                        ; intervening VSUBPD xmm2, xmm0
```

10.17 Using SIMD Instructions for Rounding

Optimization

Use AVX instructions (VROUNDSS, VROUNDSD, VROUNDPS, and VROUNDPD) to round floating-point values to integers.

Application

This optimization applies to

- 32 bit software
- 64 bit software

Rationale

The AVX instruction set provides the VROUNDSS and VROUNDSD instructions for rounding scalar single- and double-precision floating-point values and the packed variants of these instructions (VROUNDPS and VROUNDPD) to round four single-precision or two double-precision values at a

time. These instructions allow you to perform rounding without setting the Rounding Control field of the MXCSR Control and Status Register.

For details of how these instructions are used, see the *AMD64 Architecture Programmer's Manual Volume 6: 128-Bit and 256-Bit XOP, FMA4 and CVT16 Instructions*, order# 43479.

10.18 Using SIMD Instructions for Floating-Point Comparisons

Optimization

Whenever possible, use the three-operand AVX instructions (VCMPSD, VCMPPS and CMPPD) to compare floating point values.

Application

This optimization applies to

- 32 bit software
- 64 bit software

Rationale

The AVX instruction set provides three-operand instructions (VCMPSD, VCMPPS and VCMPPD) for comparing single- and double-precision scalar and vector floating point values. These “non-destructive” instructions write their resulting values into a separate destination register, rather than overwriting one of the source registers. This improves performance by increasing the work per instruction and reduces the the need to save and reload register operands.

Example

The VCMPSD/VCMPPS instructions are especially useful when used in conjunction with the VPCMOV instruction to compute the minimum, maximum and median elements in an array of floating point numbers.

The following pseudocode shows how to compute the minimum of an array of double precision floating-point values, using these instructions.

Original loop:

```
double min = a[0]

for(int i = 0; i < a.length; i++)
{
    if(a[i] < min)
        min = a[i];
}
```

```
}
```

Assembly code for the loop body using AVX instructions:

```
        vmovsd xmm1,QWORD PTR [rcx] ;xmm1 = a[0]
loopmin:
        vmovsd xmm2,QWORD PTR [rcx+r11*8] ;xmm2=a[i]
        vcmpsd xmm0,xmm1,xmm2,1
        vpcmov xmm1,xmm1,xmm2,xmm0 ;xmm1=(xmm1<xmm2)?xmm1:xmm2
        add r11,1
        cmp r11,r8
        jl loopmin
```

Chapter 11 Multiprocessor Considerations

This chapter covers the following topics:

Topic	Page
ccNUMA Optimizations	191
Writing Instruction Bytes to Memory on Multiprocessor Systems	200
Multithreading	202
Data Organization	203
Data Caching	203
False Data Sharing	204
Data-Parallel Threading	205
Stream Processing	206
Multithreaded Libraries	207
Locked Instructions as Memory Barriers	208
Store/Store Barriers in WB Memory	210
Optimizing Inter-Core Data Transfer	210

11.1 ccNUMA Optimizations

AMD multiprocessor systems use cache coherent non-uniform memory access (ccNUMA). For details on optimizing applications for ccNUMA systems, see *Performance Guidelines for AMD ccNUMA Multiprocessor Systems*, order# 40555.

11.1.1 Scheduling Single and Multithreaded Applications on Multiprocessor Systems

Optimization

On AMD family 15h multi-core multiprocessor systems, schedule threads in such a way as to maintain a balanced system load. In most cases, it is advisable to rely on the ccNUMA-aware operating system to make the correct scheduling decisions for single and multi-threaded applications.

Be sure the operating system is properly configured to support ccNUMA. Most versions of Microsoft® Windows™-based operating systems support ccNUMA. For 64-bit Linux™, there may be separate kernels supporting ccNUMA that should be selected. The 2.6.x Linux kernels feature NUMA awareness in the scheduler. Most SuSE and Red Hat 64-bit Linux distributions have the ccNUMA-aware kernel. All versions of Solaris™ for AMD64 support ccNUMA without change.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Some multiple processor systems available today employ a symmetric multiprocessing (SMP) architecture. Processors on an SMP platform generally share a common or centralized memory bus having identical memory access latencies regardless of the processor position. Because the processors use the same bus and memory, system performance may be negatively affected when bottlenecks occur due to increased demands on the single memory bus. Figure 3 shows a simplified diagram of a two processor(2P) SMP system.




Figure 7. Simple SMP Block Diagram

AMD family 15h and later multiprocessor systems implement cache coherent non-uniform memory access (ccNUMA) architecture to connect two or more processors. In a ccNUMA design, each processor has its own memory system. In AMD family 15h and later multiprocessor systems, each processor has its own memory controller and its own local memory. When a processor accesses its local memory, the latency is relatively low, especially when compared to that of a similar SMP system. If a processor accesses remote memory—that is, memory located on a different processor—then the access latency is higher. The phrase 'non-uniform memory access' refers to this potential difference in latency. Figure 4 on page 197 shows a simplified diagram of a two processor (2P) AMD processor system in a ccNUMA configuration.




Figure 8. AMD 2P System

AMD family 15h dual processor systems can have up to four compute units on each processor chip that share the on-chip integrated memory controller and memory. Figure 5 shows a simplified diagram of a two processor (2P) AMD family 15h system in a ccNUMA configuration. (Recall that a “compute unit” consists of two independent sets of integer units and a floating point unit that all share a level 2 cache, as well as instruction fetch, decode, and dispatch units.)

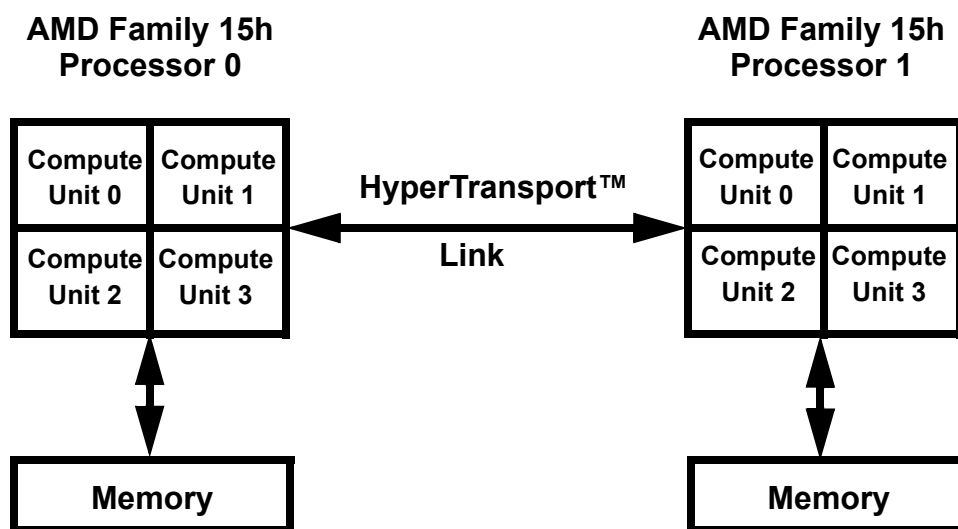


Figure 9. Dual AMD Family 15h Processor Configuration

An operating system running on an AMD family 15h platform transparently coordinates and manages the memory configuration. Thus, it is not necessary for applications to be aware of memory configuration details. Thanks to the OS, the platform simply appears to have one contiguous block of memory, regardless of how many processors are in the platform. The architecture simultaneously ensures that the entire shared memory space gives consistent values despite potentially parallel accesses from different processors. The phrase “cache coherence” in a ccNUMA system refers to this guaranteed memory consistency.

In an AMD 2P multiprocessor system, each processor is directly connected to the other processor. In addition to the 2P configuration, AMD offers 4P and larger configurations.

Figure 6 shows an example of a four processor AMD family 15h system in a ccNUMA configuration. The processors, also called nodes, are numbered Node 0, Node1, Node2 and Node 3 clockwise from the top left. Each node has four compute units that are labeled CU0, CU1, CU2 and CU3, respectively.

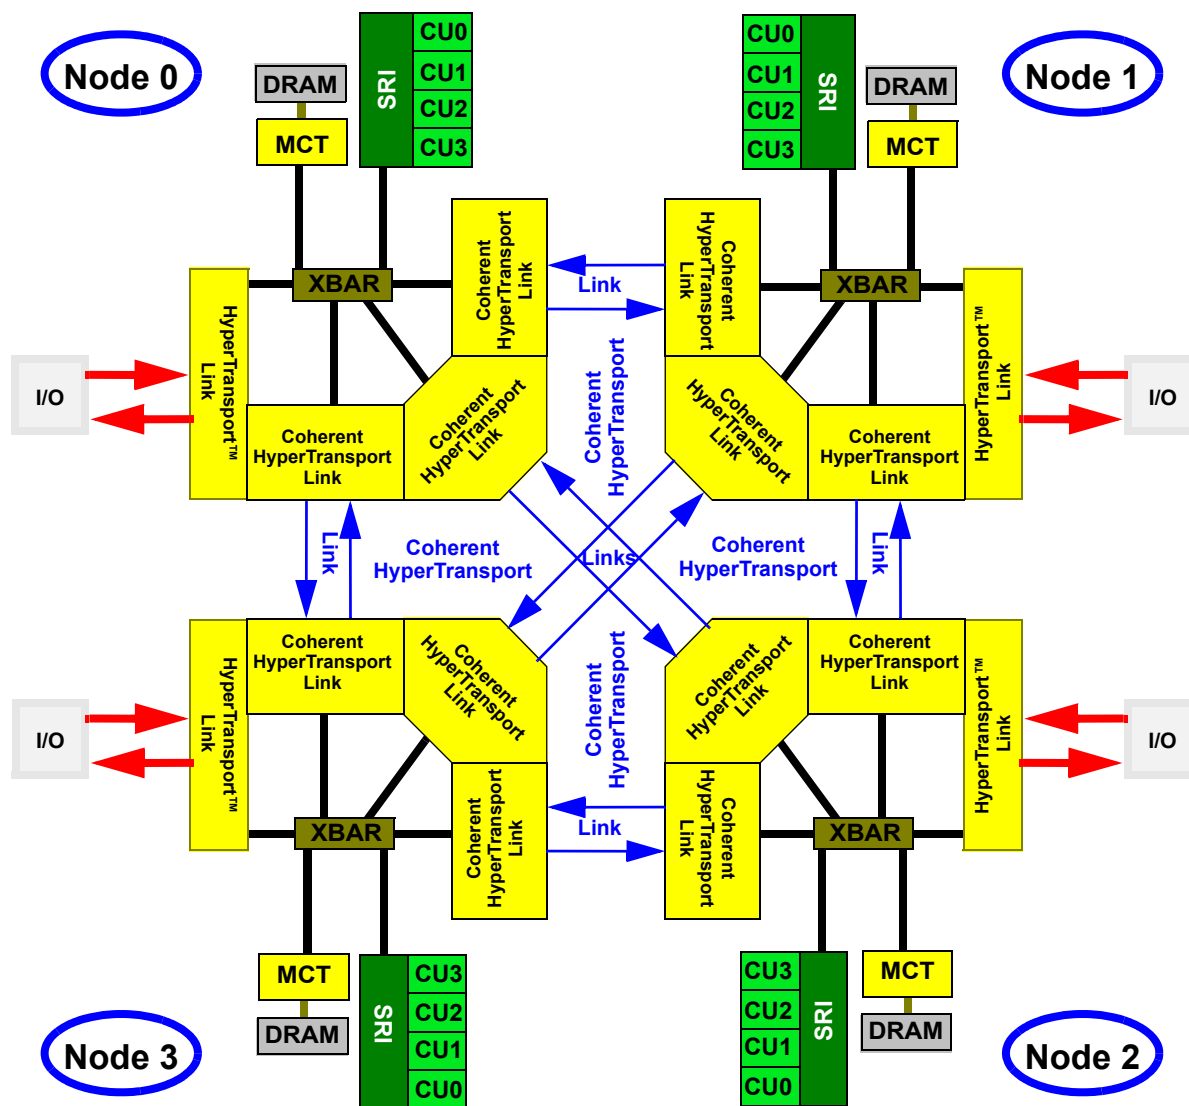


Figure 10. Block Diagram of a ccNUMA AMD Family 15h Quad-Core Multiprocessor System

The four processors are connected by coherent HyperTransport™ links. In a typical configuration, each processor has one bidirectional non-coherent link that is dedicated to I/O and three bidirectional coherent HyperTransport links that each connect to each quad-compute unit processor in the configuration. In a 4-way configuration, this assures a direct connection for any given quad-compute unit processor to all the other quad-compute unit processors in the system but one. Each node is connected to its own memory.

The term *hop* is commonly used to describe access distances on NUMA systems. If a thread accesses memory on the same node as that on which the thread is running, the memory access is considered a zero-hop access or *local* access. If a thread is running on one node but accessing memory that is resident on a different node, the access is considered a *remote* access. If the node on which the thread is running and the node on which the memory is resident are directly connected to each other, it is a one-hop access. If they are indirectly connected to each other (no direct coherent HyperTransport link), it is considered a multi-hop access.

Family 15 AMD processors provide four Hypertransport links per processor, enabling four processor systems to be connected with remote memory that is never more than one hop away from any requestor. This is an improvement upon earlier implementations of AMD64 processors that only provided three Hypertransport links, requiring some two-hop remote memory references for four-processor systems.

Large configurations of more than four processors may contain processors that are not directly connected. However, Hypertransport links may now be divided in half, enabling more direct connections between processors but at a reduced bandwidth. System OEMs may choose to use either divided Hypertransport links or configurations with multi-hop remote memory references. Depending on the number of processors and configuration, two- and even three-hop memory accesses are possible. For large-scale systems, check with the manufacturer's documentation for specifics on the Hypertransport interconnection scheme.

The four compute units on each node of the AMD family 15h processor share the Northbridge, memory and HyperTransport technology resources available on that node. Scheduling should be carried out in such a way as to avoid overloading the resources on a single node, while leaving the resources on the rest of the system unused—in other words, loads should be balanced.

Scheduling multiple threads across nodes and cores of a system is complicated by a number of factors:

- Whether multiple threads access independent data.
- Whether multiple threads access shared data.
- Whether the system is idle.

Multiple Threads-Independent Data

When scheduling multiple threads that access independent data on an idle system, it is preferable, first, to schedule the threads to an idle core of each node until all nodes are exhausted and, then, to schedule the other idle core of each node. In other words, schedule using node major order first,

followed by core major order. Since each compute unit has a single L2 cache, this is the preferred scheduling approach for independent data for AMD family 15 multi processors.

For example, when scheduling threads that access independent data on a four-way quad-core AMD family 15h system, scheduling the threads in the following order is recommended (see Figure 10 on page 194):

- compute unit 0 on node 0, node 1, node 2 and node 3 in any order
- compute unit 1 on node 0, node 1, node 2 and node 3 in any order
- compute unit 2 on node 0, node 1, node 2 and node 3 in any order
- compute unit 3 on node 0, node 1, node 2 and node 3 in any order

Multiple Threads-Shared Data

When scheduling multiple threads that share data on an idle system, it is preferable to schedule the threads on the compute units of an idle node first, then on compute units of the the next idle node, and so on. In other words, schedule using core major order first followed by node major order. However, for AMD family 15 processors, there are some variations to core major ordering due to the configuration of each compute unit. For applications with a lot of integer processing, scheduling two threads per compute unit is beneficial, before scheduling the processor's next compute unit. This enables these two threads to share data via the L2 cache, which enables higher performance than sharing via the L3 cache. On the other hand, for floating point intensive applications, it may be more desirable to schedule one thread per compute unit to avoid scheduling conflicts in the single floating point unit per compute unit. Because the floating point unit has a high capacity, and since floating point programs also contain many integer operations (address arithmetic, for example), we assume in general that scheduling two threads per compute unit is preferred for data sharing applications.

For example, when scheduling threads that share data on a four-way quad-core AMD family 15h system, AMD recommends using the following order:

- compute unit 0, compute unit 1, compute unit 1, compute unit 2, compute unit 3, compute unit 3 on node 0 in any order
- compute unit 0, 1, 2, or 3 on node 1 in any order
- Core 0, 1, 2, or 3 on node 2 in any order
- Core 0, 1, 2, or 3 on node 3 in any order

Scheduling on a Non-Idle System

It is a more difficult task to schedule multiple threads optimally for an application on a non-idle system. It requires that the application make global holistic decisions about machine resources, coordinate itself with other applications already running, and balance decisions between them. In such cases, it is better to rely on the OS to do the appropriate load balancing. In general, most

developers will achieve good performance by relying on the ccNUMA-aware OS to make the right scheduling decisions on idle and non-idle systems.

In addition to the scheduler, several NUMA-aware operating systems provide tools and APIs to allow the developer to explicitly bind a thread (set thread affinity) to a certain core or node. Using these tools or APIs overrides the scheduler and hands over control for thread placement to the program. For additional details on the thread/process affinity tools and APIs supported in various OSs, refer to Appendix C, “Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems”.

11.1.2 Data Locality Considerations on Multiprocessor Systems

Optimization

Keep data accessed by a thread local to the node on which the thread runs. In a multithreaded application in which each thread operates on largely independent data, each thread should allocate and initialize the data it accesses and allow the ccNUMA-aware operating system to make the right data locality decisions.

In multithreaded applications, performance may benefit from taking advantage of API functions or tools for thread and memory placement (thread and memory affinity) offered by the operating system.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

It is best to keep data local to the node from which it is being accessed. Accessing data remotely is slower than accessing data locally. The further the hop distance to the data, the greater the cost of accessing remote memory. For most memory latency-sensitive applications, keeping data local is the single most important recommendation to consider. While the ratio of latencies of remote accesses to local accesses may be close to 1.5 to 1 for small-scale systems, this ratio can grow to 2 or 3 or more to 1 for larger-scale topologies.

Almost all ccNUMA-aware operating systems by default rely on the first-touch policy: the physical memory for data is only committed on the node on which the thread or process writing to it for the first time runs. In general, commitment implies mapping of virtual pages to zeroed out physical pages. This is done by the OS when it detects a first-touch and takes a page fault. Thus, data is kept local on the node where the thread or process that writes to it for the first time is run.

The OS keeps data local on the node where first-touch occurs as long as there is enough physical memory available on that node. If enough physical memory is not available on the node, then different OSs use various advanced techniques to determine where to bind the data.

Memory once bound to a node by the first-touch policy normally resides on that node for its lifetime. However, the OS scheduler could migrate the thread or process that first touched the memory from one core to another core even on a different node. This can be done by the OS for the purpose of load balancing.

This can move the process/thread farther from its allocated memory. Most schedulers will try to bring the thread or the process back to the core on which the thread was previously running and on which its memory was local, but this is not guaranteed. Furthermore, the thread or process can dynamically allocate and first-touch more memory on the node to which it was moved before it is moved back. This is a difficult problem for the OS to resolve, since it has no prior information as to how long the thread or process is going to run and, hence, whether migrating it back is optimal or not.

If an application demonstrates that threads are being moved away from their associated memory by the scheduler, it is typically useful to explicitly set thread placement. By explicitly pinning a thread to a node, the application can tell the OS to keep the thread on that node and, thus, keep data accessed by the thread local to it by the virtue of the first-touch policy.

The performance improvement obtained by explicit thread placement may vary depending on whether the application is multithreaded, whether it needs more memory than available on a node, whether threads are being moved away from their data, etc.

In cases in which threads are scheduled from the outset on a core that is remote from their data, it might be useful to explicitly control data placement. This is discussed in detail in the “Scheduling on a Non-Idle System” on page 196. Advanced software developers can refer to Appendix C, “Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems,” on page 313 for additional details on support for these tools and APIs in various OSs.

11.1.3 Techniques to Minimize and Alleviate Data Sharing

Optimization

Avoid accessing data in memory that was first touched by a thread running on a different node.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

When data is shared between threads running on different nodes, the default policy of local allocation by first-touch used by the OS can become non-optimal.

For example, a multithreaded application may have a startup thread that sets up the environment, allocates and initializes a data structure and forks off worker threads. As per the default local

allocation policy, the data structure is placed in the physical memory of the node where the start up thread performed the first-touch. Forked worker threads are then spread around by the scheduler to be balanced across all nodes and their cores. A worker thread starts accessing the data structure remotely from the memory on the node where the first-touch occurred. This scenario could lead to significant memory and HyperTransport traffic in the system, with the node where the data resides becoming a potential bottleneck. This situation is especially bad for performance, firstly, if the startup thread only performs the initialization and afterwards no longer needs the data structure and, secondly, if only one of the worker threads needs the data structure. In other words, the data structure is not truly shared between the worker threads.

It is best in this case to use a data initialization scheme that avoids incorrect data placement due to first-touch. This is done by allowing each worker thread to first-touch its own data or by explicitly pinning the data associated with each worker thread on the node where the worker thread runs.

Certain OSs provide memory placement tools and APIs that also permit data migration. A worker thread can use these to migrate the data from the node where the start up thread performed the first-touch to the node where the worker thread needs it. There is a cost associated with the migration and it would be less efficient than using the correct data initialization scheme in the first place.

If it is not possible to modify the application to use a correct data initialization scheme or if data is truly being shared by the various worker threads—as in a database application—then a technique called node interleaving can be used to improve performance. Node interleaving allows for memory to be interleaved across any subset of nodes in the multiprocessor system. When the node interleaving policy is used, it overrides the default local allocation policy used by the OS on first-touch.

Let us assume that the data structure shared between the worker threads in this case is of size 16 KB. If the default policy of local allocation is used, then the entire 16 KB data structure resides on the node where the startup thread does first-touch. However, using the policy of node interleaving, the 16-KB data structure can be interleaved on first-touch such that the first 4 KB ends up on node 0, the next 4 KB ends up on node 1, and the next 4 KB ends up on node 2 and so on. This assumes that there is enough physical memory available on each node. Thus, instead of having all memory resident on a single node and making that the bottleneck, memory is now spread out across all nodes.

The tools and APIs that support explicit thread and memory placement mentioned in the previous sections can also be used by an application to use the node interleaving policy for its memory. (See Appendix C, “Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems”.)

By default, the granularity of interleaving offered by the tools/APIs is usually set to the size of the virtual page supported by the hardware, which is 4 K (when system is configured for normal pages, which is the default) and 2 M (when system is configured for large pages). Therefore any benefit from node interleaving will only be obtained if the data being accessed is significantly larger than a virtual page size.

If data is being accessed by three or more cores, then it is better to interleave data across the nodes that access the data than to leave it resident on a single node. We anticipate that using this rule of thumb could give a significant performance improvement. However, developers are advised to experiment with their applications to measure any performance change.

11.1.4 Keep Locks Cacheable and Aligned to a Cache Line Boundary

Optimization

In general, it is good practice for user-level and kernel-level code to keep locks aligned to their natural boundaries. In some hardware implementations, locks that are not naturally aligned are handled with the mechanisms used for legacy memory mapped I/O and should absolutely be avoided if possible.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

If a lock is aligned properly, it is treated as a faster cache lock. The significantly slower alternative to a cache lock is a bus lock, which should be avoided at all costs. Bus locks are very slow and force serialization of many operations unrelated to the lock within the processor. Furthermore, bus locks prevent the entire HyperTransport fabric from making forward progress until the bus lock completes. Cache locks on the other hand are guaranteed atomicity by using the underlying cache coherence of the ccNUMA system and are much faster.

11.2 Writing Instruction Bytes to Memory on Multiprocessor Systems

A common situation in dynamically optimized applications is that in which a thread on one processor in a multiprocessor system (which we will call the writer) is required to replace an original code segment with some new code segment, while there are one or more other threads (executors) on other processors that could possibly execute the original code. This can occur, for example, when a function is recompiled or reoptimized at run time.

For simplicity, this section discusses the case in which the original code consists of a single instruction. If the original code consists of multiple instructions, the writer must always ensure in some way that an executor is not in the middle of the original code.

Rule 1

If the part of the original code that needs to be patched fits within an aligned 8-byte boundary, then no special considerations are necessary. The writer may simply store the new code into memory.

In the following example, the instruction itself crosses an aligned 8-byte boundary but since the first byte is not changing, the part to be changed does not cross an aligned 8-byte boundary and so can be changed with a single store.

```
Original Code xxxxxF: E8 78 56 34 12 Call $+12345678
New Code xxxxxF:      E8 44 33 22 11 Call $+11223344
```

When a modification does cross an aligned 8-byte boundary, then care must be taken that the executor not see an invalid combination of the original code and the new code. There is no architectural store instruction, including instructions that use the lock prefix, to ensure that an executor will not see a combination of the original code and the new code. Instead, one of the following methods can be used:

- Software semaphores can be used between the writer and the executors to prevent executors from entering the original code

or

- The original code can be modified in stages by first writing a branch at the beginning of the original code to catch an executor and then modifying the remaining code. In this case a system-dependent delay must be used after writing the branch. This delay is necessary to ensure that any executor that had already fetched the first bytes of the original code (before the branch was written) has finished fetching the rest of the original code.

To modify in stages, the writer uses the following steps:

1. Modify the beginning of the original code with a branch that will catch any executor that enters the original code. The easiest branch to use is the two-byte short JMP to self (bytes EB, FE). This requires that the first two bytes of the original code not cross an 8-byte boundary. When the original code is generated and the compiler knows that it is a candidate for patching, it is recommended that a NOP be inserted. If the first two bytes of the original code do cross an 8-byte boundary, a one-byte BPT instruction and a special BPT trap handler that returns to the BPT instruction must be used.
2. Wait for a system-dependent delay. This delay ensures that any Executor who had fetched the beginning of the Original Code before the branch was written has finished fetching the rest of the Original Code.
 - The maximum amount of delay can be lessened by avoiding patches across a 4K byte page boundary and lessened further by avoiding patches which cross a 64 byte cache-line boundary.
3. Leaving the self-branch in place, modify the rest of the original code with the new code.
4. Replace the self-branch with the corresponding bytes of the new code.

11.3 Multithreading

The subject of creating multithreaded software is quite broad, and many resources exist that address its various aspects. Here we briefly discuss those aspects of multithreading that are most relevant from a hardware perspective.

To fully utilize the CPU power of multicore processors, applications must implement *scalable* threading. In other words, the application must be able to partition the work load into a variable number of threads, to match the available resources on the particular machine.

Many of these problems can be resolved by the implementation of various programming practices, including *task decomposition*, *careful data organization*, and *data caching and sharing*. Two of the most important and straightforward ways to implement scalable threading are by means of *data-parallel threading* and *stream processing*. These methods are described in detail in the following sections.

11.3.1 Task Decomposition

Optimization

For each task, use multiple threads in parallel to process equal workloads involving different data items.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Historically, multithreading has been used to implement separate functions. For example, one thread might perform I/O operations while another thread handles user input. This approach, called functional threading or task-parallel threading, can sometimes simplify the structure of a program, especially when the program is performing several asynchronous tasks.

However, functional threading has limitations. Only a fixed, limited number of threads are used. Also, the workloads in different threads are not balanced. For these reasons, functional threading is not a good match for present and future multicore processors. It doesn't scale up to utilize the hardware.

A much better approach is data-parallel threading. In data-parallel threading, each CPU-intensive task is handled in sequence. For each task, multiple threads are used in parallel to process equal workloads involving different data items. Ideally, the application can use N threads, to match an N -core processor or system. Not every processing task can be implemented using data-parallel threading. For example, data decompression and decryption are often inherently sequential tasks.

11.3.2 Data Organization

Optimization

Divide data cleanly into many largely independent sets.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Clearly, data parallel threading requires a certain class of algorithms. For example, if data is organized as a single linked list, the operation of accessing the list is not well suited to multithreading. On the other hand, an array of uniformly sized structures can usually be accessed in parallel as N chunks.

Double buffering can be used to good effect. Creating one set of data which is 100% "read only" can be valuable, even if this involves total replication of data sets in memory. Processing can read one copy of the data, while writing to the other copy. This can greatly reduce or eliminate cache thrashing and race conditions between threads. Copying all the data might not be a performance win if you are just running two threads, but it can pay off as the number of threads grows.

11.3.3 Data Caching

Optimization

Make good use of the data caches.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Single-threaded applications are influenced by cache effects. The processor stores recently used data in a local cache memory, making subsequent operations on that data run faster. All of the traditional cache factors apply to multithreaded code: limited cache size, data replacement policy, set associativity of the cache, L1 vs. L2 cache, and related criteria. (For additional information on cache

architecture and optimizations, see section 7.5, “Memory Caches” in the *AMD64 Architecture Programmer’s Manual: Volume 2 System Instructions* (order# 24593), and Chapter 6 “Cache and Memory Optimizations” on page 95. Two additional factors enter the picture when multiple threads are running on multiple cores: data sharing between caches, and false sharing between caches.

11.3.4 Data Sharing between Caches

Optimization

Design threads so that each thread operates on separate data.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

All threads in a process have a coherent view of memory. If certain data is used by multiple threads, then every time that data is modified, it must be copied into more than one cache. This data copying is avoided if threads are designed so that each thread operates on separate data. Of course, if threads are only reading the data and not modifying it, they can all share the same data, without the additional communication of updated values.

11.3.5 False Data Sharing

Optimization

To avoid false data sharing, keep each thread's data carefully separate by enforcing, for example, 64-byte alignment during allocation.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

False sharing is a subtle variation on data sharing. The data cache is managed on a cache-line basis, where each naturally aligned 64-byte cache line is treated as a unit. If any byte is modified, the entire cache line is tagged as modified. So if multiple threads access different parts of the same cache line,

and at least one thread is modifying the data, that cache line must be copied into the other caches to maintain coherence. The threads are functionally independent, but they incur a performance penalty as if they were actually sharing data. False sharing can be avoided by keeping each thread's data carefully separate, for example by enforcing 64-byte alignment during allocation.

Clean data separation at the algorithm level will minimize the occurrence of real or false data sharing.

11.3.6 Data-Parallel Threading

Optimization

Break up a workload into multiple data sets and use threads to perform the same operations on different data in parallel.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

Data-parallel threading involves breaking up a workload into multiple data sets and using threads to perform the same operations on different data in parallel. The data organization and algorithms used by the application must be carefully designed to efficiently support this form of parallelism to avoid race conditions or expensive synchronization mechanisms.

Data-parallel threading can usually be made to achieve very good load balancing between threads, efficiently utilizing all available CPU resources. Furthermore, if an application is designed for data-parallel threading, the threads do not alter the overall logical order of operations in the application, so they do not introduce as many potential deadlocks or race conditions that can complicate other threading strategies.

A trivial example of data-parallel threading involves the addition of two arrays of numbers. The basic operation might be expressed in C++ as:

```
for (int i=0; i < 10000; i++) {  
    c[i] = a[i] + b[i];  
}
```

A data-parallel multithreaded implementation would process the arrays in chunks, for example if four threads are used, the values of array index *i* would be assigned to each of the four threads as follows:

thread #0	i = 0 through 2499
thread #1	i = 2500 through 4999

```
thread #2      i = 5000 through 7499
thread #3      i = 7500 through 9999
```

Threading can be implemented explicitly, for example, using Win32 thread APIs on Windows[®] or Pthreads on Linux[®]. The application must choose how many threads to run, create and/or start the threads, and detect their completion. The application is also typically responsible for detecting the number of processors available at run time.

Most modern compilers also support OpenMP, which greatly simplifies the syntax for data-parallel threading in loops. For the following loop, using OpenMP requires just one extra line of code (a pragma), which partitions the workload across an appropriate number of threads:

```
#pragma omp parallel for
for (int i=0; i < 10000; i++) {
    c[i] = a[i] + b[i];
}
```

OpenMP has a simple API, and it supports many options for controlling how the data-parallel threading is executed. For details see <http://www.openmp.org>.

On platforms that support multiple CPU nodes (as opposed to simply supporting one multicore CPU), additional performance gains can be achieved because of greater system memory bandwidth available. Memory buffers for storing thread-specific data should be allocated locally to the NUMA node, by calling the allocation function from within the thread. In some heavily threaded applications, it also makes sense to set thread affinity at the time of thread creation to distribute them across multiple NUMA nodes. Care must be taken when manually setting affinity.

11.3.7 Stream Processing

Use stream processing to operate on large arrays of related data.

Application

This optimization applies to:

- 32-bit software
- 64-bit software

Rationale

With the advent of truly programmable graphics processing units (GPUs), the programming paradigm of stream processing has become much more relevant. Strictly speaking, stream processing is not a form of multithreading, but it shares many of the same constraints on data organization and algorithm choice as data-parallel threading does.

In stream processing, a set of kernels (i.e., functions) operate on *streams*—large arrays of related data. Typically, kernels implement math operations that can be vectorized, for instance, by using vector

SIMD instructions for best performance. For maximum efficiency, kernels consume streams that are generated as output by other kernels; these streams persist locally in the low-latency processor data cache, instead of making a trip through system memory.

At the appropriate time, streams are explicitly moved between processor cache and system memory, as a logically separate process from the kernel operations, so data movement is only loosely coupled to processing. In principle, this decoupling can enable more efficient gather/scatter operations on blocks of data that comprise the streams. For maximum efficiency, stream data should be organized contiguously in memory. In many cases, for best performance the stream data can be read from memory using software prefetch instructions and, finally, written back to memory using the streaming store instructions, which avoid disturbing the L2 cache.

If the application's algorithms and data structures are mappable onto the stream/kernel model, then a stream processing approach can be profitably implemented. This can result in increased processor performance because data cache locality and memory bandwidth are well utilized and also because data-parallel threading can usually be employed in conjunction with the stream processing. Perhaps even more importantly, organizing an application to fit the stream processing model can pave the way for off-loading the heavy computational workloads to a highly parallel GPU chip or other specialized processor.

11.3.8 Multithreaded Libraries

Programmers can use multithreaded code libraries, such as the Framewave and AMD Core Math Library (ACML), to great advantage in writing applications that incorporate the multithreading paradigm.

11.4 Memory Barrier Operations

Memory barriers of type A/B, where A and B represent either a load or store memory operation and A is ordered prior to B in program order, allow the programmer to specify that *older* memory operations of type A (load or store) cannot appear to be passed by any *younger* memory operations of type B (load or store). Here, B *passing* A means that although A precedes B in program order, the results of instructions A and B may be returned in any order.

There are four types of memory barriers:

- Load/Load—older loads are not passed by younger loads.
- Store/Store—older stores are not passed by younger stores.
- Load/Store—older loads are not passed by younger stores.
- Store/Load—older stores are not passed by younger loads.

Memory Barriers in WB Memory

On the AMD64 architecture, when using writeback (WB) type memory without streaming stores, the only type of barrier that requires an explicit barrier instruction is Store/Load. When streaming stores are used, the Store/Store barrier also requires an explicit barrier instruction. In WB memory, all other barriers are implicit in the AMD64 architecture. For additional information on memory and memory barrier instructions, see “Forcing Memory Order” in the *AMD64 Architecture Programmer’s Manual Volume 1: Application Programming* and Chapter 7, “Memory System” in the *AMD64 Architecture Programmer’s Manual Volume 2: System Programming*.

Memory barriers in WB memory are unnecessary in systems consisting of a single processor core.

Store/Load Barriers in WB Memory

On the AMD64 architecture there are three ways to achieve the Store/Load barrier in WB memory (see section 3.9 in APM volume 1):

- MFENCE instruction
- A locked instruction that reads and writes memory—any instruction of the form `LOCK op mem, reg` or `LOCK op mem, imm`. (The specific instruction `XCHG mem, reg` is treated as locked whether or not a LOCK prefix is used.)
- Architecturally serializing instructions such as CPUID.

11.4.1 Locked Instructions as Memory Barriers

Optimization

Use locked instructions to implement Store/Load barriers.

Application

Applies to programs running on multicore processors or on multiple single core processors.

Rationale

On AMD family 15h processors, the MFENCE instruction is a serializing instruction. This stalls the pipeline and the processor core cannot begin processing any further instructions until all *previous* instructions are completed and any outstanding memory operations (such as prefetches and stores) have completed. (This stall applies only to the individual integer unit of the compute unit where the MFENCE instruction is executed.) Architecturally serializing instructions such as CPUID have the same pipeline stall behavior as MFENCE. The LOCKed instructions do not stall the pipeline and, thus, allow more instruction-level parallelism.

LOCKed instructions that access shared memory (memory shared between processor cores) incur a delay while the cache line is changed to modified state and data is (potentially) transferred between

caches in the system. LOCKed instructions that are not naturally aligned incur the very high overhead of a bus lock.

When possible, make the LOCKed instruction perform a useful store (an XCHG *mem,reg* instruction can be used for this purpose, assuming the *reg* can be overwritten). The memory location that is the target of the store (e.g., XCHG) instruction should be in the exclusive state in the processor core's local L1 cache. Avoid using a memory location that is shared with other processor cores (even if it is only written by the local processor core). It is also very important to avoid using a memory location that is not naturally aligned.

Thus, for a pattern such as:

```
mov    localmem2, rax        ;; store to local memory
mov    sharedmem1, rbx      ;; store to shared memory
StoreLoad_Barrier
mov    rcx, sharedmem3      ;; load from shared memory
```

Preferred:

```
mov    sharedmem1, rbx
xchg   localmem2, rax      ;; performs local store and StoreLoad barrier
                               ;; in one instruction (note: modifies rax)
mov    rcx, sharedmem3
```

Avoid:

```
mov    localmem2, rax
xchg   sharedmem1, rbx    ;; avoid using shared mem for locked operation
mov    rcx, sharedmem3
```

Avoid:

```
mov    sharedmem1, rbx
mov    localmem2, rax
mfence                               ;; avoid MFENCE which is serializing
mov    rcx, sharedmem3
```

When the locked instruction cannot be made to do a useful store, there are several variations of LOCK *op mem, imm* that do not modify the memory contents or any registers other than the FLAGS register, for example:

```
LOCK OR DWORD PTR localmem, 0
```

To repeat, the memory location that is the target of the locked instruction should be in the exclusive state in the processor's local L1 cache. A location on the stack such as [RSP] in 64-bit mode or [ESP] in 32-bit mode generally meets this criteria. To avoid Store-To-Load forwarding issues, the location should be addressed using the same data width with which it is otherwise accessed.

11.4.2 Store/Store Barriers in WB Memory

When performing a Store/Store in WB memory, a Store/Store barrier is only required when using streaming stores and then only in systems with more than one processor core. Store/Store barriers can be achieved in one of the following ways (see section 3.9, “Memory Optimization” in *AMD64 Architecture Programmer’s Manual Volume 1: Application Programming*):

- Using the SFENCE instruction
- Using any of the methods covered under the topic of Store/Load barriers.

Optimization

Use the SFENCE instruction to implement a Store/Store barrier.

Application

Applies to programs running on multicore processors or on multiple single core processors.

Rationale

The SFENCE instruction is not a serializing instruction, so it achieves the desired effect of waiting for the write-combining buffers to drain, while allowing parallelism with other non-store instructions.

11.5 Optimizing Inter-Core Data Transfer

AMD family 15h processors incorporate four distinct cores on a single die and have a cache that all the cores share. Shared caches provide an efficient way to handle a group of computational problems belonging to a producer/consumer model: a program thread running on one core produces data that is intended for consumption by a thread that is running on another core. In such cases, round trips to and from main memory can be avoided by arranging for pairs of cores to communicate through the shared cache.

A naïve implementation of a producer/consumer program will produce bandwidth results that appear to be throttled by main memory speeds. Yet, with some knowledge of the cache architecture, it is possible to boost throughput significantly.

The producer/consumer program is handled by setting up a system in which the producer and consumer threads chase each other around a ring buffer through which they communicate and share data. When a thread reaches the end of the buffer, it wraps back around to the beginning and keeps reading or writing. There are three scenarios that affect the achievable bandwidth of communication between a producer and a consumer thread based on the amount of modified data and how it relates to cache sizes.

- If the producer and consumer are executing together on a single compute unit, there is no minimum time lag required between writes and reads of the ring buffer. This is true because the

L1 cache for each execution unit in the compute unit is write-through. Thus data produced on one execution unit is available for consumption in the shared L2 cache.

- If the producer and consumer threads are executing on separate compute units, then ideally communication is through the shared L3 cache. In this case, the consumer thread must lag the producer thread by an amount determined by the size of the L2 cache. The ring buffer cannot be smaller than this minimum distance or the producer thread will not be able to spill its data into the L3 cache when the consumer wants to read it, causing a longer latency read that misses the L3 cache and snoops the producer thread's L2 cache.
- If the ring buffer is of significant size, as explained below, the producer thread must not lead the consumer thread by too much or the data it places in the L3 cache will more than likely be evicted to main memory before the producer thread wraps back around again.

These caveats are discussed in the following sections.

Cache sizes and thread distances

For data, the AMD Family 15h processor has three levels of cache, a 16-Kbyte L1 data cache, a 1–2MB unified L2, and a 16MB or greater L3. There are two L1 data caches, one private to each integer execution unit; the L2 cache is shared between the execution units for each compute unit. The highest level L3 cache is shared by all four compute units. AMD family 15h cache design is a combination of write through and “mostly exclusive” “victim” cache.

“Mostly” exclusive cache hierarchies. In an *exclusive* cache hierarchy, only one copy of the data exists anywhere in the entire cache hierarchy. In such a system, when a thread hits in the L3 cache, it moves the cache line to a lower level of cache (which is private to the particular compute unit running the thread) without leaving a copy in the L3 cache. However, in AMD Family 15h processors, for purposes of optimizing for multiple readers, multiple cores may generate local copies when they access the shared L3 cache line. For this reason, the cache is termed “mostly” exclusive.

Victims. Bringing a chunk of new data into a cache requires allocation of cache space by displacing an older chunk, the so-called “victim,” from that cache. This victim cache line is then pushed into the next higher level cache in the hierarchy; of course this, in turn, triggers the displacement of a victim from the higher level cache. For this reason, the L3 (shared) cache is only filled with victims from the L2 cache. For AMD family 15 processors, victims from the L1 cache are not pushed to the L2 cache. This is due to the write-through nature of the L1 cache as described above. There is no need to push a victim from the L1 cache to the L2 cache, as the line is already present in the L2 cache.

Thread sizes and distances. Because the L1 cache is write-through, when the producer thread produces a cache-line of data, it exists both in the L1 cache and the L2 of the producer core; it does not get written-through to the L3 caches. When the L2 cache is full, data is evicted to the L3 cache. For a program to maximize the bandwidth between two cores through the processor L3 cache, the lag distance between the consumer threads and the producer threads must be greater than the size of the L2 cache (1MB or 2MB), so that the consumer compute unit can read victim blocks from the producer compute unit.

If the producer and consumer are both executing on the same compute unit, the communication takes place through the L2 cache, which is updated immediately with every write. In this case, no lag distance is required.

When communicating through the L3 cache, the producer thread must get too far ahead of the consumer thread, or it will flood the L3 cache with victims and start spilling unread cache-lines to memory. On a four compute-unit processor, four producer/consumer pairs can run simultaneously, so the most optimal allocation would be to divide the L3 cache evenly between pairs; for example, on a microprocessor with a 16MB L3 cache and a 2MB L2 cache, each pair should be allocated 16MB/4 of L3 cache. Therefore, the producer thread should be at least 2MB ahead of the consumer, but no more than $4\text{MB} + 2\text{MB} = 6\text{MB}$ or it will evict its own data from the L3 cache. If only one producer/consumer pair is running on the processor, the entire L3 cache can be dedicated to the pair. On the other hand, if the distance between the producer and consumer can be kept less than the size of the L2 cache (2MB in our example), then the producer and consumer threads should be scheduled to the same compute unit.

MOESI protocol issues

In the case of communication via the L3 cache, even if the above mentioned distance constraints are followed, the measured bandwidth will still be limited by the performance of the DRAM controller. This results from the subtle interaction of the cache lines as they are touched by the producer and consumer threads and the MOESI protocol that AMD processors implement to maintain cache coherency. (For a complete description of the MOESI protocol, see the *AMD64 Architecture Programmer's Manual Vol 2: System Programming*, order# 24593.)

The MOESI cache coherency protocol is defined by the state of data in a cache line in relation to other copies of the data (in memory, another processor cache, etc.). These states are summarized as follows:

- **Modified (M)**—The cache line holds the most recent correct copy of the data and the copy in memory is stale. (No other copies exist.)
- **Owned (O)**—A cache line in the owned state holds the most recent, correct copy of the data, which may be shared by other processors. This copy is responsible for updating main memory, when evicted. (The copy in memory may be stale and other processors may hold a copy in the S state.)
- **Exclusive (E)**—A cache line holds the most recent, correct copy of the data, which is identical to the copy in memory. (No other processor holds a copy of the data.)
- **Shared (S)**—A cache line in the shared state holds the most recent, correct copy of the data, which may be shared by other processors. (The copy in memory may be stale.)
- **Invalid (I)**—A cache line does not hold a valid copy of the data. (valid copies are in main memory or another processor cache.)

The producer thread allocates cache-lines in the modified (M) state, as an automatic consequence of writing a new entry. Eventually, these M-marked cache lines will start to fill the L3 cache, thanks to

the adherence to the abovementioned rules defining the allowable distance between threads. When the the consumer reads the cache line, the MOESI protocol changes the state of the cache line to owned (O) in the L3 cache and pulls down a shared (S) copy for its own use. Now, the producer thread circles the ring buffer to arrive back to the same cache line it had previously written. However, when the producer attempts to write new data to the owned (marked 'O') cache line, it finds that it cannot, since a cache line marked 'O' by the previous consumer read does not have sufficient permission for a write request (in the MOESI parlance). To maintain coherence, the memory controller must initiate probes in the other caches (to handle any other S copies that may exist)—and this is slow.

Thus, it is preferable to keep the cache line in the 'M' state in the L3 cache. Then, when the producer comes back around the ring buffer, it finds the previously written cache line still marked 'M', to which it is safe to write without coherency concerns. This is exactly what happens when the producer and consumer are communicating through a shared L2 cache within the same compute unit.

The PREFETCHW instruction provides the means to control this cache allocation restriction. The PREFETCHW instruction provides a hint to the processor that the program intends to modify the cache line, so the processor keeps the cache line in the 'M' state. To clarify how this works, we will step through a scenario in which the consumer thread uses PREFETCHW to proactively fetch cache lines.

As previously mentioned, the producer thread first spills a cache line marked 'M' into the L3 cache. At some later time, the consumer thread executes the PREFETCHW instruction to load the cache line into the L1 and L2 data caches; this time, the processor keeps the cache line marked 'M', removes the cache line from L3 and pulls the line down the cache hierarchy, into the consuming core's L1 and L2 caches. No 'S' copy is handed to the consumer core, and no 'O' copy remains in the L3 cache. As far as the producer compute unit and its associated L2 cache are concerned, the cache line is gone.

If the producing thread were to wrap around the buffer and attempt to write to the cache line, it would register as a cache miss and a request would be sent to the memory controller. To avoid this, the consumer needs to evict the cache line back out to the processor shared L3. As discussed previously, the only way to achieve this is to have the cache line trickle back up the cache hierarchy, until it eventually becomes a victim block into the L3 cache again; in other words, the consumer reads enough memory to equal its L2 cache size (1MB or 2MB), forcing the cache line to evict before the producer thread needs to write to that memory. When that happens, it's as if the producer never knew that the cache line was gone. When it writes new data to the cache line and it finds it in the L3 cache; the producer and consumer are communicating through the L3 cache, fully utilizing the inherent speed and bandwidth.

To assure that the consumer has enough time to spill its contents into the L3 cache, the producer thread must lag the consumer thread by a distance that is at least equal to the size of the L2 cache. This has further implications for the size of the ring buffer. If the consumer must lag the producer by X bytes, and the producer must lag the consumer by X bytes, then the buffer must be at least $2 \times X$ in size to achieve maximum performance. In practice, it is not a bad idea to pad the numbers by at least two units of granularity to allow for some extra space, since cache eviction is not controllable or precise. For instance, assuming that the producer and consumer threads write data in 16K chunks

before checking their positions relative to each other, it would be safe to have the producer and consumer enforce a distance of $((L2) + 2 \times \textit{granularity})$ or $(1\text{MB}) + 2 \times 16\text{K} = 1056\text{K}$ apart (assuming a 1MB L2 cache size). This in turn implies that the ring buffer size should be at least 2112K.

Alternatively, a producer and consumer executing within the same compute unit can communicate directly through the L2 cache. In this case, the data written in the L1 (and thus the L2) is in the modified state, and remains so, as long as no other processors access the data and the L2 is not forced to purge the data to the L3. In this scenario, higher communication bandwidth can be achieved between the two threads, but the ring buffer must fit in the L2 cache and remain resident there. In this scenario, a size between one-half and three-quarters the size of the L2 cache is recommended.

Summary

A producer/consumer program can achieve harmony by successfully bouncing an ‘M’ marked cache line back and forth between the consuming and producing threads either through the fast L3 cache or through the even faster L2 cache. For a producer and consumer executing on the same compute unit, the ring buffer should be sized at roughly one-half the size of the L2 cache. The consumer need not lag the producer by more than a single cache line.

If the producer and consumer are executing on different compute units, keep in mind the following constraints:

- The consumer thread needs to ‘lag’ the producer thread by at least the L2 cache size.
- The producer thread needs to lag the consumer thread by at least the L2 cache size.
- The ring buffer should be at least $2 \times (L2)$.
- The producer thread should not get so far ahead of the consumer to flood the L3, if larger ring buffers are used.
- Use `PREFETCHW` on the consumer side, even if the consumer will not modify the data.
- Add a small extra factor to the calculated sizes to give the threads additional space when communicating through the caches.

In general, the AMD cache is optimized for widely shared data, i.e. one core produces data that may be of interest to several other compute units. The ‘S’ and ‘O’ states provide coherence for multiple readers of the same data. One compute unit is responsible for the data in the ‘O’ state, but that data can be safely shared with many other compute units through the ‘S’ state. In the producer/consumer program however, it is known ahead of time that the data the producer creates is only interesting to the matching consumer thread, and not to any other thread. Following the constraints listed above, it is possible to achieve a large increase in throughput for two producer/consumer pairs on AMD processors.

Chapter 12 Optimizing Secure Virtual Machines

The goal of this chapter is to enable virtual machine monitor (VMM) software engineers to minimize the performance overhead imposed by the virtualization of a guest. A significant consumer of processor cycles on microprocessors enabled for AMD Virtualization™ (AMD-V™) is the world switch, which refers to the process of running either a VMRUN instruction to enter a guest context or running the #VMEXIT mechanism to leave a guest context. World switch can also broadly apply to the requisite software effort surrounding VMRUN and #VMEXIT; software effort for some intercepts may be significantly longer than the VMRUN/#VMEXIT portion of the world switch. Several of the optimizations proposed in this chapter attempt to reduce the frequency of world switches. Other optimizations provide techniques to reduce software or processor effort required for performing other virtualization tasks.

For additional information on virtualization and related topics, see Chapter 15, “Secure Virtual Machine,” in the *AMD64 Architecture Programmer’s Manual Volume 2: System Programming* (order# 24593).

This chapter covers the following topics:

Topic	Page
Use Nested Paging	216
VMCB.G_PAT Configuration	217
State Swapping	217
Economizing Interceptions	218
Nested Page Size	219
Shadow Page Size	220
Setting VMCB.TLB_Control	220
TLB Flushes in Shadow Paging	221
Use of Virtual Interrupt VMCB Field	222
Avoid Instruction Fetch for Intercepted Instructions	223
Share IOIO and MSR Protection Maps	224
Obey CPUID Results	224
Using Time Sources	225
Paravirtualized Resources	226

12.1 Use Nested Paging

Optimization

- ❖ Use nested paging instead of shadow paging.

Application

This optimization applies to:

- VMMs

Rationale

To virtualize guests fully, virtual machine monitor (VMM or hypervisor) software must virtualize guests' physical memory mappings without the guests' knowledge. On processors that do not implement nested paging, a method called shadow paging is commonly used for this purpose. But that method is complex to implement efficiently, it is significantly slower than native virtual-to-physical address translation, and performance tuning often requires significant memory to store cached shadow page tables for each guest page table. (There is typically one page table per guest user process.) Shadow paging requires both significant time for the VMM to manage shadow page tables and frequent VMM intervention during guest page faults, guest CR0, CR3, and CR4 accesses, guest INVLPG execution, and guest modifications to page table contents.

In contrast to shadow paging, nested paging requires minimal VMM attention. The CR_x, INVLPG, and page fault intercepts are unnecessary, and the VMMs need only set up an initial nested page table that maps guest physical addresses to system physical addresses. Each guest requires its own nested page table. A VMM that uses nested paging is significantly less complex and, thus, is easier to validate and verify than a VMM using shadow paging.

A TLB miss under nested paging incurs potentially more memory accesses than a TLB miss under shadow paging, but AMD-V microprocessors that support nested paging employ intelligent caching to minimize the latency of a nested paging TLB miss.

TLB Miss Latency under Nested Paging

TLB entries cache translations from the virtual address to the system physical address for non-virtualized programs and from the guest virtual address to the system physical address for shadow and nested paging. A TLB hit under nested paging performs the same as a TLB hit under shadow paging or in a non-virtualized environment. A TLB miss under nested paging is potentially more expensive than a non-nested TLB miss; nested paging page table walks are accelerated by the CPU's caching of page table information.

12.2 VMCB.G_PAT Configuration

Optimization

❖ Properly configure the guest page attribute table (G_PAT) in the virtual machine control block (VMCB).G_PAT field.

Application

This optimization applies to:

- VMMs using nested paging

Rationale

When nested paging is enabled, the VMCB.G_PAT field is used to virtualize the guest's PAT register. For a description of how the final memory type of a guest page is determined, see section 15.24.8 “Combining Memory Types, MTRRs” in the *AMD64 Architecture Programmer's Manual Volume 2: System Programming* (order# 24593). For details on the organization and layout of the VMCB, see Appendix B “Layout of VMCB” in the same volume.

Operating systems typically leave the PAT at its default reset value of 0x00070406_00070406, although they are free to change the PAT register's contents. The VMM software should start up guest virtual machines with the same default value. A VMM that leaves the G_PAT value equal to 0x0 will experience significant performance degradation in the guest because all guest memory accesses will be forced to the effective PAT type of uncacheable (UC).

12.3 State Swapping

Optimization

Avoid unnecessary VMSAVE, VMLOAD, STGI, CLGI, and guest GPR and FPR state swapping.

Application

This optimization applies to:

- VMMs for guests in all modes

Rationale

Avoiding unnecessary instructions that would occur on every world switch can reduce the cost of a world switch.

For example, a VMM may need only a small subset of the state swapped by VMSAVE and VMLOAD, so the VMM that expects to return to the same guest can skip VMSAVEing the guest's

state and, instead, leave that guest state active in the CPU. If the VMM needs to use any of the VMSAVE values, such as the task register (TR), the VMM can use the LTR instruction to install the VMM's TR value, while leaving the other guest values intact. Upon returning to the guest, the VMM can VMLOAD the guest or execute the LTR instruction to restore guest state from an artificial TR entry in VMM context. To ensure that the guest VMCB contains the correct TR values, the VMM must intercept the LTR instruction in the guest.

Similar work can be done on other pieces of state represented in VMLOAD and VMSAVE.

VMRUN sets the global interrupt flag (GIF) to 1—equivalent to executing an implicit STGI instruction. Similarly, #VMEXIT clears GIF with an implicit CLGI instruction. A VMM that performs only a minimal amount of work between a #VMEXIT and the next VMRUN may wish to skip executing explicit STGI and CLGI instructions.

VMMs can use methods similar to callee-save to avoid saving and restoring all guest general-purpose registers and floating-point registers if the VMM intends to return to the same guest. This approach is probably most useful for performing lazy floating-point state saves and saving debug registers DR0-DR3.

12.4 Economizing Interceptions

Optimization

Intercept as few MSRs, events, and instructions as possible.

Application

This optimization applies to:

- VMMs

Rationale

To minimize virtualization overhead, VMMs should try to minimize the number of #VMEXITs due to MSR and instruction intercepts.

The VMM should intercept only those MSRs that are critical for system function or security, and which, therefore, must be protected from guest access. The VMM can avoid intercepting MSRs that are frequently used and changed by operating systems, such as GSBASE and KernelGSBase, and all other MSRs that are loaded by VMLOAD, since these MSRs have no system-level side-effects and can be efficiently context switched. VMM writers may evaluate the frequency of reads to specific MSRs that must be intercepted to determine if the following optimization is worthwhile: if the read value is equal to the value the guest expects, then the MSR write may be intercepted while leaving the MSR read unintercepted.

The state that is context switched by AMD-V instructions often does not require intercepts. For example, the IDTR, GDTR, LDTR, and TR read and write intercepts, and PUSHF and POPF intercepts often do not need to be set because VMRUN/#VMEXIT and VMLOAD/VMSAVE appropriately virtualize the related state.

Under nested paging, the paging-related control registers (CR0, CR2, CR3, CR4) and PAT MSR are context switched by VMRUN and #VMEXIT and, thus, often do not need to be intercepted. Similarly, the INVLPG intercept is not necessary under nested paging. In comparison, most shadow paging implementations need to intercept CR0, CR3, and CR4 read and write accesses and the INVLPG instruction, although they can avoid intercepting CR2 accesses.

To avoid the overhead of context switching floating-point state, VMMs can use lazy floating point context switching methods by controlling guest CR0.TS. When the VMM forces CR0.TS to a value other than the value the guest had written, the VMM should intercept CR0 reads and writes in order to properly virtualize CR0.TS.

12.5 Nested Page Size

Optimization

Where possible, use large pages in nested page tables.

Application

This optimization applies to:

- VMMs using nested paging

Rationale

VMMs can realize several performance advantages by using large (2 MB or 1 GB) pages in nested page tables, when it is possible for the VMM to allocate naturally-aligned large pages for portions of guest physical memory images.

The first performance increase comes from reducing multiplicative factors in the cost of TLB misses under nested paging.

Secondly, a common use of large pages is to reduce TLB pressure. For best performance, nested page table entries should be larger than or equal to the size of the corresponding guest page size.

Large pages allow the reduction of the memory footprint used by nested page tables. For each 2-MB large page in a nested page table, an entire 4-KB bottom-level page table becomes unnecessary. For each 1-GB large page, a 4-KB page-directory table becomes unnecessary, as do up to 512 bottom level page tables (each of which occupy 4 KB).

12.6 Shadow Page Size

Optimization

Use large pages where possible in shadow paging.

Application

This optimization applies to:

- VMMs using shadow paging

Rationale

For reasons similar to those enumerated in section 12.5, “Nested Page Size” above, VMMs should attempt to use large pages in shadow page tables for address ranges that the guest maps using large pages. This avoids increasing TLB pressure by the fracturing of large pages into smaller TLB entries and reduces software complexity and memory usage.

When a VMM encounters a 2-MB (or 4-MB or 1-GB) guest page and decides to map it using 4-KB shadow page table entries, the VMM must use memory to store an additional 512 derived entries, or 4 KB, for a shadow page table that does not correspond to any page table in the guest. Additionally, if the guest performs an INVLPG instruction to the guest's 2-MB page, the VMM must clear all 512 of the derived 4-KB entries and must invalidate each 4 KB derived page (in which case it is likely to be more efficient to flush the entire TLB, than to execute 512 INVLPG instructions).

12.7 Setting VMCB.TLB_Control

Optimization

When possible, avoid setting VMCB.TLB_Control to 1.

Application

This optimization applies to:

- VMMs

Rationale

Setting `VMCB.TLB_Control` to 1 and then VMRUNning that VMCB flushes the entire TLB of all entries, local and global, for all ASID values. Flushing the entire TLB can have minor but noticeable adverse effects on performance by unnecessarily flushing TLB entries from ASIDs other than the current ASID. If capacity misses would not have evicted the other ASIDs' TLB entries, then those TLB entries would be available and useful for avoiding page table walks when the VMM or other guests are executed.

12.8 TLB Flushes in Shadow Paging

Optimization

Use specific `TLB_CONTROL` encodings to flush a specific ASID's TLB entries when emulating guest TLB flushing events.

Application

This optimization applies to:

- VMMs using shadow paging

Rationale

When a VMM is using shadow paging, it must intercept every event in the guest that is defined to cause a TLB flush or TLB line invalidation. The most commonly, TLB flushes are triggered by the `MOV CR3` instruction, while the `INVLPG` instruction invalidates specified TLB entries. When these instructions are intercepted by an AMD-V processor, the TLB flush or invalidation is suppressed; the VMM assumes the responsibility for carrying out the appropriate invalidations after performing the appropriate shadow page table manipulations. A simplistic solution is to set `VMCB.TLB_CONTROL` to 1 to cause a complete TLB flush on the next VMRUN. This simplistic solution may have a negative performance impact due to the complete flushing of global entries and TLB entries for all other ASIDs. Notably, when the `TLB_CONTROL` field is set to 1, the VMM TLB entries that might otherwise be useful after `#VMEXIT` are lost, as well as potentially useful TLB entries of other guests. Earlier generations of AMD-V processors did not support Flush By ASID VMCB commands, which led to an optimization that involved managing ASIDs. This optimization is functionally correct but is superseded on processors that support Flush By ASID. Earlier generations of AMD-V processors did not support Flush By ASID VMCB commands, which led to an optimization that involved managing ASIDs. This optimization is functionally correct but is superseded on processors that support Flush By ASID.

When the hypervisor emulates a local TLB flush (for example, a guest `CR3` write), it should write 111b to the `TLB_CONTROL` field. The next VMRUN will flush the appropriate local entries for that VMCB's current ASID, without flushing other ASIDs' TLB entries. Similarly, when the hypervisor emulates a global TLB flush, it should write 011b to the `TLB_CONTROL` field. The hypervisor

should clear the TLB_CONTROL field before any subsequent VMRUN to avoid redundant redundant TLB flushes.

These optimizations are functionally compatible with AMD-V processors that do not support Flush By ASID because each of the enhanced encodings of the TLB_COMMAND sets bit 1; however, the ASID changing optimization is likely to result in better performance on those processors.

Note that an intercepted INVLPG instruction can be turned into a shadow page table operation followed by an INVLPGA instruction and does not necessarily require a TLB flush.

12.9 Use of Virtual Interrupt VMCB Field

Optimization

Use the Virtual Interrupt VMCB field instead of event injection when there is only one interrupt pending for the guest.

Application

This optimization applies to:

- VMMs, when VMCB.V_INTR_MASKING == 1 for the guest.

Rationale

VMMs commonly do not allow guests direct access to physical interrupts, choosing instead to virtualize the interrupts using the V_INTR_MASKING and virtual interrupt mechanisms.

AMD-V processors automatically deliver a pending virtual interrupt to the guest when the guest is not masking interrupts due to any of the following:

- Guest EFLAGS.IF == 0
- Guest TPR > priority of pending virtual interrupt
- Guest is in an interrupt shadow

VMMs can avoid the overhead and complexity in software of determining if a guest is ready to take the interrupt by appropriately filling the virtual interrupt fields in the guest VMCB, and can avoid one or more unnecessary world switches. An AMD-V processor automatically clears the V_IRQ valid bit when the interrupt is taken.

By taking these steps, VMMs can provide correct interrupt behavior to the guest while using the smallest possible number of world switches.

12.10 Avoid Instruction Fetch for Intercepted Instructions

Optimization

Avoid all guest instruction fetches.

Application

This optimization applies to:

- VMMs

Rationale

On all instruction intercepts, Family 15h processors provide sufficient information in the VMCB to allow the hypervisor to avoid fetching guest instruction bytes. This saves the hypervisor the overhead of traversing guest page tables and the complexity and overhead of enforcing synchronization with other potentially active guest virtual CPUs. The information delivered by the Decode Assists corresponds to the instruction as originally fetched from the instruction stream and can be used directly by the hypervisor without additional memory synchronization. (For a full description of instruction intercepts, see section 15.8, “Instruction Intercepts” in the AMD64 Architecture Programmer’s Manual Volume 2: System Programming.)

Additionally, all instructions that directly cause an intercepted data page fault or a nested data page fault will deliver up to 15 instruction bytes to allow the hypervisor to decode the instruction without traversing the guest page tables. This is useful, for example, when a hypervisor is emulating an MMIO access to a device. It should be noted that the instruction bytes may be fetched by the processor using data paths. If a hypervisor uses the instruction bytes delivered in the VMCB, it must ensure consistency between the instruction and data TLBs by issuing an INVLPGA instruction to invalidate the address containing the guest’s RIP (the hypervisor should issue a pair of INVLPGA instructions to the beginning and end bytes of the instruction if the instruction crosses a 4KB aligned boundary). This INVLPGA sequence is required to be executed once before returning to the guest vCPU after the use of the instruction bytes.

12.11 Share IOIO and MSR Protection Maps

Optimization

Share IOIO and MSR protection maps, if possible, to save memory.

Application

This optimization applies to:

- VMMs

Rationale

A VMM running multiple guests typically enforces the same I/O port and MSR restrictions on most or all guests in the system. While a VMM must allocate one VMCB per guest virtual CPU, the VMM can conserve memory by sharing common IOIO and MSR protection maps. These structures can be shared because they are read, but never written, by the CPU. The VMM should be careful about using proper mutual exclusion to handle modifications done to protection maps that are in use on other CPUs.

12.12 Obey CPUID Results

Optimization

- ❖ Guests should obey CPUID results.

Application

This optimization applies to:

- All programs, operating systems, and libraries

Rationale

Any existing or future operating system, program, or library may be executed in a virtualized environment. A VMM may control the results of CPUID to hide certain capabilities from the guest for various reasons. The VMM may wish to enable migration of a guest from one processor to a processor of a different generation with different features enabled. The VMM provides a set of CPUID results to its guests that represents a common subset of features. That subset may not represent any existing physical processor.

To ensure that programs, libraries, and operating systems work properly in the face of virtualization, all software should obey the results returned by CPUID. The most straightforward way to obey CPUID is to execute CPUID once per program or library initialization and then record the result in an

internal data structure. For example, a program may detect the RDTSCP indicator in CPUID and then configure code paths to reflect the presence or absence of RDTSCP. The VMM's control over the RDTSCP CPUID bit will cause the program to exhibit the correct behavior based on whether the VMM wishes to advertise the fact that the current CPU implements the RDTSCP instruction.

This restriction is eased for existing programs and existing methods to detect processor features that already existed at the time AMD-V microprocessors were introduced. For example, before using legacy SSE1 instructions, user programs are required to do a try-catch sequence to determine if the operating system has enabled the XMM registers. This try-catch sequence is still required for SSE1 instructions, but software must adhere to the results of CPUID instruction without a try-catch sequence for detecting new instructions like the SSE3 instruction set.

Future CPU versions may add new instruction encodings to replace formerly undefined encodings. Software should never depend on #UD exceptions from instructions that are currently undefined on any given processor. The UD2 opcode should be used if software wishes to create #UD exceptions.

12.13 Using Time Sources

Optimization

Guests should be careful about using time sources.

Application

This optimization applies to:

- All programs, operating systems, and libraries

Rationale

Programs and operating systems that are not virtualization-aware might assume that the RDTSC instruction, high precision event timers (HPETs), programmable interrupt timers (PITs), and other time sources are monotonically increasing by constant amounts and are usable as a measure of both elapsed time and wall-clock time. When a VMM is present, it necessarily intercepts guest operation for variable lengths of time, and must make adjustments to the time values read by the guest. These adjustments may break one or more assumptions about time sources. A VMM may choose to adjust the time sources to synchronize them with a wall-clock time so that the guest's time of day measurements are correct, in which case a guest that is continuously monitoring the time will see occasional jumps in the apparent wall-clock time; this may cause fairness problems with the guest's process scheduling. A VMM may choose to adjust the time sources so the guest correctly measures elapsed guest time, which would cause the guest's TSC-based measurement of wall-clock time to be incorrect and may affect time-critical applications such as media playback.

It is unlikely that a guest that is unaware of virtualization will be able to use time sources for all common purposes at the same time. Users should be aware of these pitfalls and understand their

implications. As operating systems and programs are written to be aware of virtualization, they should take advantage of any available paravirtualized access to time resources. For their part, VMMs should strive to provide a sufficiently rich and standardized set of paravirtualized timer resources.

12.14 Paravirtualized Resources

Optimization

Guests should detect VMM presence and use paravirtualized resources

Application

This optimization applies to:

- All guests that are aware of virtualization

Rationale

An OS does not implicitly know whether it is a guest or if the OS is running without a VMM present. Some VMMs may support paravirtualization as a means to improve performance or create features. When this is the case, guests should use industry-standard methods to detect VMMs and enumerate the available paravirtualized functions. System resources, such as paging controls in non-nested paging environments, time references, network and video drivers, storage and other device drivers can benefit from paravirtualization.

Appendix A Implementation of Write-Combining

This appendix describes the memory write-combining feature implemented in AMD Family 15h processors. Write-combining is the merging of multiple memory write cycles that target locations within the address range of a write buffer.

AMD Family 15h processors support the memory type range register (MTRR) and the page attribute table (PAT) extensions, which allow software to define ranges of memory as either writeback (WB), write-protected (WP), writethrough (WT), uncacheable (UC), or write-combining (WC).

Defining the memory type for a range of memory as WC or WT allows the processor to conditionally combine data from multiple write cycles that are addressed within this range into a merge buffer. Merging multiple write cycles into a single write cycle reduces processor bus utilization and processor stalls. Write combining buffers are also used for streaming store instructions such as MOVNTQ and MOVNTI. See “Prefetch and Streaming Instructions” on page 105.

This appendix covers the following topics:

Topic	Page
Write-Combining Definitions and Abbreviations	227
Programming Details	228
Write-Combining Operations	228
Sending Write-Buffer Data to the System	229
Write Combining to MMIO Devices that Support Write Chaining	229

A.1 Write-Combining Definitions and Abbreviations

This appendix uses the following definitions and abbreviations:

- MTRR—Memory type range register
- PAT—Page attribute table
- UC—Uncacheable memory type
- WC—Write-combining memory type
- WT—Writethrough memory type
- WP—Write-protected memory type
- WB—Writeback memory type

A.2 Programming Details

Write-combining regions are controlled by the MTRRs and PAT extensions. Write-combining should be enabled for the appropriate memory ranges. (For more information on the MTRRs and the PAT extensions, see the *AMD64 Architecture Programmer's Manual, Volume 2*, order# 24593, and the *BIOS and Kernel Developer's Guide for AMD Family 15h Processors*, order# 31116.)

A.3 Write-Combining Operations

To improve system performance, AMD Family 15h processors aggressively combine multiple memory-write cycles of any data size that address locations within a 64-byte write buffer that is aligned to a cache-line boundary. The processor continues to combine writes to this buffer without writing the data to the system, as long as certain rules apply (see Table 7 for more information). The data sizes can be bytes, words, doublewords, or quadwords.

- WC memory type writes can be combined in any order up to a full 64-byte write buffer.
- All other memory types for stores that go through the write buffer (UC, WP, WT and WB) cannot be combined except when the WB memory type is over-ridden for streaming store instructions such as the MOVNTQ and MOVNTI instructions, etc. These instructions use the write buffers and will be write-combined in the same way as address spaces mapped by the MTTR registers and PAT extensions. When WC is used for streaming store instructions, then the buffers are subject to the same flushing events as write-combined address spaces.

Combining continues until interrupted by one of the conditions listed in Table 7. When combining is interrupted, one or more bus commands are issued to the system for that write buffer and all older write buffers, even if they are not full, as described in “Sending Write-Buffer Data to the System” on page 229.

Table 7. Write-Combining Completion Events

Event	Comment
I/O Read or Write	Any IN/INS or OUT/OUTS instruction closes combining. The implied memory type for all IN/OUT instructions is UC, which cannot be combined.
Serializing instructions	Any serializing instruction closes combining. These instructions include: MOVCRx, MOVDRx, WRMSR, INVD, INVLPG, WBINVD, LGDT, LLDT, LIDT, LTR, CPUID, IRET, RSM, INIT, and HALT.
Flushing instructions	Any flush instruction causes the WC to complete.
Locks	Any instruction or processor operation that requires a cache or bus lock closes write-combining before starting the lock. Writes within a lock can be combined.
Uncacheable Read	A UC read closes write-combining. A WC read closes combining only if a cache block address match occurs between the WC read and a write in the write buffer.

Table 7. Write-Combining Completion Events (Continued)

Event	Comment
Different memory type	When a store hits on a write buffer that has been written to earlier with a different memory type than that store, the buffer is closed and flushed.
Buffer full	Write-combining is closed if all 64 bytes of the write buffer are valid.
TLB AD bit set	Write-combining is closed whenever a TLB reload sets the accessed [A] or dirty [D] bits of a PDE or PTE.
Executing an SFENCE (Store Fence) instruction.	The SFENCE instruction forces the completion of pending stores, including those within the WC memory type, making these globally visible and emptying the store buffer and all write-combining buffers.
An interrupt or exception occurs.	Interrupts and exceptions are serializing events that force the processor to write all results to memory before fetching the first instruction from the interrupt or exception service routine

A.4 Sending Write-Buffer Data to the System

Maximum throughput is achieved by write combining when all quadwords or doublewords are valid and the AMD Family 15h processors can use one efficient 64-byte memory write instead of multiple 8-byte memory writes.

A.5 Write Combining to MMI/O Devices that Support Write Chaining

AMD Family 15 processors implement write combining through the use of a 4KB, 4-way set associative write coalescing cache (WCC) and a 64 byte, 4 entry write combining buffer (WCB). The data in the WCC is inclusive with the L2 data cache, implying that write data in this cache is globally visible. It will only contain writes for the writeback memory type. (Recall that streaming store instructions enable write combining for the writeback memory type. These writes, as well as writes to other memory types, including the write combining (WC) memory type, occur in the WCB, and are not globally visible until flushed from the WCB.) Because writes to writeback memory are globally visible in the WCC, they may require data to be flushed from the WCB to ensure the enforcement of ordering rules between writes to different memory types. The 4K WCC thus provides 64 write-combining buffers for the WB memory type, each 64 bytes wide. For the WC memory type, writes can be combined up to the capacity of the WCB, allowing for 64 byte writes to the system.

Having multiple write-combining buffers that can combine independent WC streams has implications on data throughput rates (bandwidth), especially when data is written by the CPU to WC memory mapped I/O devices, residing on the AGP, PCI, PCI-X[®] and PCI Express[®] buses including:

- Memory Mapped I/O registers—command FIFO, etc.
- Memory Mapped I/O apertures—windows to which the CPU uses programmed I/O to send data to a hardware device

- Sequential block of 2D/3D graphic engine registers written using programmed I/O
- Video memory residing on the graphics accelerator—frame buffer, render buffers, textures, etc.

HyperTransport™ Tunnels and Write Chaining

HyperTransport™ tunnels are HyperTransport-to-bus bridges. Many HyperTransport tunnels use a hardware optimization feature called write-chaining. In write-chaining, the tunnel device buffers and combines separate HyperTransport packets of data sent by the CPU, creating one large burst on the underlying bus when the data is received by the tunnel in sequential address order. Using larger bursts results in better throughput since bus efficiency is increased. This is because bus arbitration overhead is lower: only one address/attribute phase is issued per burst in the PCI-X case, and one address/command phase is issued for the AGP Fast Writes case.

For reasons cited in the preceding paragraph, to utilize hardware write chaining efficiently, software should flush the CPU write-combining buffer in sequential linear address order, any time a target GB/s hardware device is capable of receiving large bursts of CPU write data.

Software should be aware that on AMD64 processors that have multiple write-combining buffers, events that flush the write-combining buffers (see Table 7 on page 228) do so in the order that the streams were opened. For example, if the CPU writes to the WC space in the 64-byte buffer at the highest address first (say, address 40h), followed by a write to a lower 64-byte buffer (for example, address 00h), the CPU first sends the highest addressed 64-byte buffer by HyperTransport to the tunnel, followed by the second (lower address) 64-byte buffer. Since the addressing is not sequential the tunnel device will not chain both 64-byte WC buffers and must issue two separate transactions on the target bus.

If the buffers in this example were targeted for AGP fast writes, issuing two fast write transactions (rather than issuing one fast write transaction) would reduce the bandwidth (data throughput) by one-third.

Optimizations

Adhere to the following guidelines to ensure that AMD Family 15h processors issue WC buffers in sequential address order:

- When practical, shadow the data structure in memory (rather than writing the actual WC buffer in MMI/O space), prior to copying the structure to WC MMI/O space. This will also ensure that the write-combining buffers are not emptied prematurely by external events (such as a UC read—perhaps issued by another device driver thread or a hardware interrupt, etc.). Shadowing also ensures that writes that occur to different cache lines in the structure do not send out the WC buffers, since the number of WC buffers that can be open at one time is CPU implementation dependent.
- When ready to update the actual WC MMI/O address space, copy the shadowed structure from memory to MMI/O, from the lowest address 64-byte block upward. To do the copy, use discrete loads and stores for up to 64 bytes of data. Use a loop of discrete loads and stores for up to 4KB of

data. Use REP MOVS instructions for up to 32KB of data. To do discrete loads use assembly language, or, if available, compiler intrinsic functions available (`__movsb()`, `__movsw()`, `__movsd()`), etc. (For more information, see section 6.9 “Memory and String Routines” on page 117.)

- In general, using these methods to do the copy will exhibit less overhead in a data movement function than calling a `memcpy()` LIBC function, which is usually optimized for copying larger blocks of memory.

Appendix B Instruction Latencies

This appendix provides a listing of AMD64 instructions, decode types, and execution latencies. For more information on these instructions, see the *AMD64 Architecture Programmer's Manual, Volumes 3, 4, and 5 (order# 24594, 26568, and 26569)*.

The instruction entries in this appendix are grouped into categories as follows and are presented within each category in alphabetical order by mnemonic:

Topic	Page
Understanding Instruction Entries	233
General Purpose and Integer Instruction Latencies	237
System Instruction Latencies	250
FPU Instruction Latencies	253
Amended Latency for Selected FMA Instructions	309

B.1 Understanding Instruction Entries

To use the information in this appendix effectively, you need to understand how the entry for an instruction is organized and how to interpret certain items.

Example: Instruction Entry

The entry for an instruction begins with its syntax. Subsequent columns provide additional information about the instruction.

Syntax	Decode Type	FPU Pipe(s)	Latencies	Comments
ADDPD_mem	FMA(P0 P1)	FastPath Single	10	4

Components of the Instruction Entry

Columns in the latency tables are defined as follows. Not all categories are relevant to all instruction sets. Thus, only the pipe and latency are relevant to system instructions, general purpose instruction latency tables use all five categories.

Category	Description
Syntax	Shows the syntax for the instruction—the permitted arrangement of its parts. Items in italics are placeholders for operands that you must provide. For information on how to interpret the placeholders, see “Interpreting Placeholders” on page 235
Instruction	This category is used by the FPU instruction latency table to specify the format of an FPU instruction (e.g., reg–mem or reg–reg). For details on how to interpret the instruction format, see “Interpreting Instruction Format” on page.
Pipes	Lists the possible floating-point unit (FPU) pipeline available for use by any particular FastPath single, FastPath Double or microcoded operation.
Decode type	Shows the method that the processor uses to decode the instruction—FastPath Single, FastPath Double, or microcode.
Latency	Shows the static execution latency for the instruction. For details on how to interpret the latency information, see “Interpreting Latencies” on page 235.
Comments	Specifies clarifying information

Note: Each instruction mnemonic that is commented as “Repeat After 2 Cycles” can issue the instruction at the indicated number of interval cycles.

Decode Type

The decode type hierarchy, from simplest to most complex is:

FastPath Single ↔ FastPath Double ↔ Microcode.

Pipes

EX0 and EX1 represent the ALU execution units.

AG0 and AG1 represent the address generation units.

Intepreting FPU Pipe Assignments

The following table shows hows each of the four pipes are mapped to floating-point units in the AMD family 15h architecture.

Table 8. Mapping of Pipes to Floating-Point Units

Pipe 0	Pipe 1	Pipe 2	Pipe 3
FPFMA fmul, fadd, fmac	FPFMA fmul, fadd, fmac	FPMAL avx, simd, mmx, ALU	FPMAL avx, simd, mmx, ALU
FPCVT fconverts	FPXBR shuffles, packs, permutes		FPSTO fpstore
FPFMA avx, simd, mmx, multiplier			

As indicated above, some floating point units (FPFMA, FPMAL) map to multiple pipes. This means, for example, that the VFMADDPS instruction can map to either pipe 0 or pipe 1. In the latency table for the FPU instructions, each instruction mapping is displayed in the Pipes column.

Interpreting Placeholders

The Syntax column for an instruction entry shows the mnemonic for the instruction followed by any operands. Items in italics are placeholders for operands that you must provide. A placeholder indicates the size and type of operand that is allowed.

This operand	Is a placeholder for
reg	A general-purpose register
mmrx	An MMX™ register
xmm	A 128-bit SIMD register
ST(<i>i</i>)	X87 stack register
mem, mem(32/64)	A memory location
imm	An immediate value
disp	A memory displacement or offset
pm32	A relative offset
x/y	Operand type x or y
x (mem)	Operand type x or mem (used only for media instructions)
fn0x2	CPUID-specific information
near/far	Semantics for CALL instructions
<p>Note: Operands with numbers indicate operand sizes, for example <i>mem32/64</i> indicates that this operand can either be a 32-bit or a 64-bit memory location. When sizes are not indicated, the information in the entry is identical for any legal operand size. Please consult the AMD64 Architecture Programmer's Manual Volumes 3–6 to determine the legal operand sizes for a given instruction type.</p>	

Interpreting Instruction Format

An entry in the 'Instruction' category takes the form *mnemonic_optype*, combining an instruction name with the type of its operand. This signature uniquely identifies the particular form of the instruction. For example, *ADDPD_mem* represents the memory form of the instruction (*ADDPD xmm1, mem128*) while *ADDPD_reg* represents the register form (*ADDPD xmm1, xmm2*). In some cases, a form such as *mem32* is used for sake of brevity to identify special cases of memory size. For more information on these instructions, see the AMD64 Architecture Programmer's Manual, Volumes 3, 4, 5 and 6.

Interpreting Latencies

The Latency column for an instruction entry shows the static execution latency for the instruction. The static execution latency is the number of clock cycles it takes to execute the serially dependent sequence of micro-ops that comprise the instruction.

The latencies in this appendix are estimates and are subject to change. They assume that:

- The instruction is an L1-cache hit that has already been fetched and decoded, with the operations loaded into the scheduler.
- Memory operands are in the L1 data cache.
- There is no contention for execution resources or load-store unit resources.

Each latency in the table denotes the typical execution time of the instruction when run in isolation on a processor with any referenced memory locations already in the L1 cache. For real programs executed on this highly aggressive superscalar family of processors, multiple instructions can execute simultaneously; therefore, the effective latency for any given instruction's execution may be overlapped with the latency of other instructions executing in parallel. An example of this effect can be seen for an SIMD load-compute instruction like `ADDPD reg, mem`, which effectively adds 4 cycles of latency (10 cycles total) over `ADDPD reg, reg`, which uses 6 cycles when run in isolation. In a real program, however, the load portion of the instruction often occurs in parallel with earlier work, effectively hiding the extra 4 cycles from the critical execution path. There are also other cases of additional latencies that may be incurred in a real program that are not described in the latency table, such as delays caused by L1 cache misses or contention for execution or load-store unit resources.

The following formats are used to indicate the static execution latency:

Table 9. Latency Formats

Latency format	Description	Example
x	The latency is the indicated value.	3
x/y	The latency is additive according to the pipe used for executing the instruction. A latency entry of this format occurs when the pipe entry is of the form a/b ; latency x corresponds to pipe a and latency y corresponds to pipe b , meaning total latency is effectively $x + y$.	6/10

B.2 General Purpose and Integer Instruction Latencies

The latency table for general purpose and integer instructions gives the decode type and latency corresponding to each instruction mnemonic. For more detailed information on the operation of a particular general purpose integer instruction, as well as encoding information, see the *AMD64 Architecture Programmer's Manual Volume 3: General-Purpose and System Instructions*, order# 24594.

Table 10. General Purpose and Integer Instruction Latencies

Syntax	Pipes	Decode Type	Latencies	Comments
AAA	microcode	microcode	NA	
AAD	microcode	microcode	NA	
AAM	microcode	microcode	NA	
AAS	microcode	microcode	NA	
ADC <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
ADC <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
ADC <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
ADC <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
ADC <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
ADD <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
ADD <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
ADD <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
ADD <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
ADD <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
AND <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
AND <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
AND <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
AND <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
AND <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
ANDN <i>reg, reg, reg</i>	EX0 EX1	FastPath Double	NA	
ANDN <i>reg, reg, mem</i>	EX0 EX1	FastPath Double	NA	
BEXTR <i>reg, reg, reg</i>	EX0 EX1	FastPath Double	NA	
BEXTR <i>reg, mem, reg</i>	EX0 EX1	FastPath Double	NA	
BEXTR <i>reg, reg, imm</i>	EX0 EX1	FastPath Double	NA	
BEXTR <i>reg, mem, imm</i>	EX0 EX1	FastPath Double	NA	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
BLCFILL <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLCFILL <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLCI <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLCI <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLCIC <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLCIC <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLCMSK <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLCMSK <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLCS <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLCS <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLSFILL <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLSFILL <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLSI <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLSI <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLSIC <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLSIC <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLSMSK <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLSMSK <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BLSR <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
BLSR <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
BOUND <i>reg32, mem64</i>	microcode	microcode	NA	
BSF <i>reg, reg</i>	microcode	microcode	NA	
BSF <i>reg, mem</i>	microcode	microcode	NA	
BSR <i>reg, reg</i>	microcode	microcode	NA	
BSR <i>reg, mem</i>	microcode	microcode	NA	
BSWAP <i>reg</i>	EX0 EX1	FastPath Single	1	
BT <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
BT <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
BT <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
BT <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
BTC <i>reg, reg</i>	EX0 EX1	FastPath Double	2	
BTC <i>reg, imm</i>	EX0 EX1	FastPath Double	2	
BTC <i>mem, imm</i>	EX0 EX1	FastPath Double	6	
BTC <i>mem, reg</i>	EX0 EX1	FastPath Double	6	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
BTR <i>reg, reg</i>	EX0 EX1	FastPath Double	2	
BTR <i>reg, imm</i>	EX0 EX1	FastPath Double	2	
BTR <i>mem, imm</i>	EX0 EX1	FastPath Double	6	
BTR <i>mem, reg</i>	EX0 EX1	FastPath Double	6	
BTS <i>reg, reg</i>	EX0 EX1	FastPath Double	2	
BTS <i>reg, imm</i>	EX0 EX1	FastPath Double	2	
BTS <i>mem, imm</i>	EX0 EX1	FastPath Double	6	
BTS <i>mem, reg</i>	EX0 EX1	FastPath Double	6	
CALL <i>disp (near)</i>	AG0 AG1 EX0 EX1	FastPath Double	2	First op to AG0 AG1, Second to EX0 EX1
CALL <i>reg (near)</i>	AG0 AG1 EX0 EX1	FastPath Double	2	First op to AG0 AG1, Second to EX0 EX1
CALL <i>mem (near)</i>	AG0 AG1 EX0 EX1	FastPath Double	6	First op to AG0 AG1, Second to EX0 EX1
CBW	EX0 EX1	FastPath Single	1	
CWDE	EX0 EX1	FastPath Single	1	
CDQE	EX0 EX1	FastPath Single	1	
CWD	EX0 EX1	FastPath Double	1	
CDQ	EX0 EX1	FastPath Single	1	
CQO	EX0 EX1	FastPath Single	1	
CLC	EX0 EX1	FastPath Single	1	
CLD	microcode	microcode	NA	
CMC	EX0 EX1	FastPath Single	1	
CMOVcc <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
CMOVcc <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
CMP <i>reg, reg</i>	EX0 EX1	FastPath Single	1	If branch fused, op to EX1, else op to EX0 EX1
CMP <i>reg, imm</i>	EX0 EX1	FastPath Single	1	If branch fused, op to EX1, else op to EX0 EX1
CMP <i>mem, reg</i>	EX0 EX1	FastPath Single	5	If branch fused, op to EX1, else op to EX0 EX1
CMP <i>mem, imm</i>	EX0 EX1	FastPath Single	5	If branch fused, op to EX1, else op to EX0 EX1
CMP <i>reg, mem</i>	EX0 EX1	FastPath Single	5	If branch fused, op to EX1, else op to EX0 EX1
CMPS	microcode	microcode	NA	
CMPSB	microcode	microcode	NA	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
CMPSW	microcode	microcode	NA	
CMPD	microcode	microcode	NA	
CMPSQ	microcode	microcode	NA	
CMPXCHG <i>reg, reg</i>	microcode	microcode	3	
CMPXCHG <i>mem8, reg8</i>	microcode	microcode	7	
CMPXCHG <i>mem16, reg16</i>	microcode	microcode	6	
CMPXCHG <i>mem32, reg32</i>	microcode	microcode	6	
CMPXCHG <i>mem 64, reg64</i>	microcode	microcode	6	
CMPXCHG8B <i>mem64</i>	microcode	microcode	7	
CMPXCHG16B <i>mem128</i>	microcode	microcode	6	
CPUID fn0x0	microcode	microcode	114	
CPUID fn0x1	microcode	microcode	53	
CPUID fn0x2	microcode	microcode	132	
CPUID fn0x3	microcode	microcode	132	
CPUID fn0x4	microcode	microcode	132	
CPUID fn0x5	microcode	microcode	118	
CPUID fn0x6	microcode	microcode	124	
CPUID fn0x7	microcode	microcode	132	
CPUID fn0x8	microcode	microcode	132	
CPUID fn0x9	microcode	microcode	132	
CPUID fn0xA	microcode	microcode	132	
CPUID fn0xB	microcode	microcode	85	
CPUID fn0xC	microcode	microcode	132	
CPUID fn0xD	microcode	microcode	162	
DAA	microcode	microcode	NA	
DAS	microcode	microcode	NA	
DEC <i>reg</i>	EX0 EX1	FastPath Single	1	
DEC <i>mem</i>	EX0 EX1	FastPath Single	5	
DIV <i>reg</i>	microcode	FastPath Single		See "Optimizing Integer Division" on page 165.
DIV <i>mem</i>	microcode	FastPath Single		See "Optimizing Integer Division" on page 165.
ENTER <i>imm32, 0</i>	microcode	FastPath Single	NA	
ENTER <i>imm32, 1</i>	microcode	FastPath Single	NA	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
ENTER <i>imm32</i> , 2	microcode	FastPath Single	NA	
IDIV <i>reg</i>	microcode	FastPath Single		See “Optimizing Integer Division” on page 165.
IDIV <i>mem</i>	microcode	FastPath Single		See “Optimizing Integer Division” on page 165.
IMUL <i>reg8</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
IMUL <i>reg16</i>	EX1	FastPath Double	4	Repeat after 2 cycles.
IMUL <i>reg16</i> , <i>imm16</i>	EX1	FastPath Double	4	Repeat after 2 cycles.
IMUL <i>reg16</i> , <i>mem16</i>	EX1	FastPath Double	8	Repeat after 2 cycles.
IMUL <i>reg16</i> , <i>mem16</i> , <i>imm</i>	EX1	FastPath Double	8	Repeat after 2 cycles.
IMUL <i>reg16</i> , <i>reg16</i>	EX1	FastPath Double	4	Repeat after 2 cycles.
IMUL <i>reg16</i> , <i>reg16</i> , <i>imm</i>	EX1	FastPath Double	4	Repeat after 2 cycles.
IMUL <i>reg32</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
IMUL <i>reg32</i> , <i>imm32</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
IMUL <i>reg32</i> , <i>mem32</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
IMUL <i>reg32</i> , <i>mem32</i> , <i>imm</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
IMUL <i>reg32</i> , <i>reg32</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
IMUL <i>reg32</i> , <i>reg32</i> , <i>imm</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
IMUL <i>reg64</i>	EX1	FastPath Single	6	Repeat after 4 cycles.
IMUL <i>reg64</i> , <i>imm32</i>	EX1	FastPath Single	6	Repeat after 4 cycles.
IMUL <i>reg64</i> , <i>mem64</i>	EX1	FastPath Single	10	Repeat after 4 cycles.
IMUL <i>reg64</i> , <i>mem64</i> , <i>imm</i>	EX1	FastPath Single	10	Repeat after 4 cycles.
IMUL <i>reg64</i> , <i>reg64</i>	EX1	FastPath Single	6	Repeat after 4 cycles.
IMUL <i>reg64</i> , <i>reg64</i> , <i>imm32</i>	EX1	FastPath Single	6	Repeat after 4 cycles.
IMUL <i>mem8</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
IMUL <i>mem16</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
IMUL <i>mem32</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
IMUL <i>mem64</i>	EX1	FastPath Single	8	Repeat after 4 cycles.
INC <i>reg</i>	EX0 EX1	FastPath Single	1	
INC <i>mem</i>	EX0 EX1	FastPath Single	5	
Jcc <i>disp</i>	EX0 EX1	FastPath Single	1	
JCXZ <i>disp</i>	EX0 EX1	FastPath Single	1	
JECXZ <i>disp</i>	EX0 EX1	FastPath Single	1	
JRCXZ <i>disp</i>	EX0 EX1	FastPath Single	1	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
JMP <i>reg</i> (near)	EX0 EX1	FastPath Single	1	
JMP <i>disp</i> (near)	EX0 EX1	FastPath Single	1	
JMP <i>mem</i> (near)	EX0 EX1	FastPath Single	5	
JMP <i>disp</i> (far, no call gate)	microcode	microcode	NA	
JMP <i>mem</i> (far, no call gate)	microcode	microcode	NA	
LAHF	EX0 EX1	FastPath Single	4	
LEA <i>reg16, mem</i> (2 operands)	EX0 EX1	FastPath Single	1	
LEA <i>reg16, mem</i> (3 operands)	EX0 EX1 AG0 AG1	FastPath Double	2	First op to AG0 AG1, Second to EX0 EX1.
LEA <i>reg32, mem</i> (2 operands)	EX0 EX1	FastPath Single	1	
LEA <i>reg32, mem</i> (3 operands)	EX0 EX1 AG0 AG1	FastPath Double	2	First op to AG0 AG1, Second to EX0 EX1.
LEA <i>reg64, mem</i> (2 operands)	EX0 EX1	FastPath Single	1	
LEA <i>reg64, mem</i> (3 operands)	EX0 EX1 AG0 AG1	FastPath Double	2	First op to AG0 AG1, Second to EX0 EX1.
LEAVE	microcode	microcode	NA	
LLWPCB <i>reg</i>	microcode	microcode	NA	
LODS	microcode	microcode	NA	
LODSB	microcode	microcode	NA	
LODSW	microcode	microcode	NA	
LODSD	microcode	microcode	NA	
LOOP/LOOPcc <i>pm32</i>	EX0 EX1	FastPath Single	1	
LOOPcc <i>pm32</i>	EX0 EX1	FastPath Single	1	
LOOP	EX0 EX1	FastPath Single	1	
LOOPcc <i>pm64</i>	EX0 EX1	FastPath Single	1	
LWPVAL <i>reg, reg32, imm32</i>	microcode	microcode	NA	
LWPVAL <i>reg, mem32, imm32</i>	microcode	microcode	NA	
LWPINS <i>reg, reg, imm</i>	microcode	microcode	NA	
LWPINS <i>reg, reg, imm</i>	microcode	microcode	NA	
LZCNT <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
LZCNT <i>reg, mem</i>	EX0 EX1	FastPath Double	6	
LZCNT <i>reg, mem16</i>	EX0 EX1	FastPath Double	NA	
LZCNT <i>reg, mem32</i>	EX0 EX1	FastPath Double	NA	
LZCNT <i>reg, mem64</i>	EX0 EX1	FastPath Double	NA	
MOV <i>reg, reg</i>	EX0 EX1	FastPath Single	1	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
MOV <i>reg, mem8</i>	EX0 EX1	FastPath Single	5	
MOV <i>reg, mem16</i>	EX0 EX1	FastPath Single	5	
MOV <i>reg, mem32</i>	EX0 EX1	FastPath Single	4	
MOV <i>reg, mem63</i>	EX0 EX1	FastPath Single	4	
MOV <i>mem, reg</i>	EX0 EX1	FastPath Single	4	
MOV <i>mem, imm</i>	EX0 EX1	FastPath Single	4	
MOV <i>mem16, FS</i>	microcode	microcode	NA	
MOV <i>mem32, SS</i>	microcode	microcode	NA	
MOV <i>mem32, DS</i>	microcode	microcode	NA	
MOV <i>reg32, SS</i>	microcode	microcode	NA	
MOV <i>reg32, DS</i>	microcode	microcode	NA	
MOV <i>reg32, FS</i>	microcode	microcode	NA	
MOV <i>reg64, FS</i>	microcode	microcode	NA	
MOV <i>SS, mem32</i>	microcode	microcode	NA	
MOV <i>SS, reg32</i>	microcode	microcode	NA	
MOV <i>DS, mem32</i>	microcode	microcode	NA	
MOV <i>DS, reg32</i>	microcode	microcode	NA	
MOV <i>FS, mem16</i>	microcode	microcode	NA	
MOV <i>FS, reg32</i>	microcode	microcode	NA	
MOV <i>FS, reg64</i>	microcode	microcode	NA	
MOVNTI <i>mem, reg</i>	EX0 EX1	FastPath Single	1	
MOVS	microcode	microcode	NA	
MOVSB	microcode	microcode	NA	
MOVSW/	microcode	microcode	NA	
MOVSD	microcode	microcode	NA	
MOVSQ	microcode	microcode	NA	
MOVSX <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
MOVSX <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
MOVXSD <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
MOVXSD <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
MOVZX <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
MOVZX <i>reg, mem</i>	EX0 EX1	FastPath Single	4	
MUL <i>reg8</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
MUL <i>reg16</i>	EX1	FastPath Single	4	Repeat after 2 cycles.

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
MUL <i>reg32</i>	EX1	FastPath Single	4	Repeat after 2 cycles.
MUL <i>reg64</i>	EX1	FastPath Single	6	Repeat after 4 cycles.
MUL <i>mem8</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
MUL <i>mem16</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
MUL <i>mem32</i>	EX1	FastPath Single	8	Repeat after 2 cycles.
MUL <i>mem64</i>	EX1	FastPath Single	10	Repeat after 4 cycles.
NEG <i>reg</i>	EX0 EX1	FastPath Single	1	
NEG <i>mem</i>	EX0 EX1	FastPath Single	5	
NOP	EX0 EX1	FastPath Single	0	No resources mapped.
NOT <i>reg</i>	EX0 EX1	FastPath Single	1	
NOT <i>mem</i>	EX0 EX1	FastPath Single	5	
OR <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
OR <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
OR <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
OR <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
OR <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
POP <i>reg16</i>	EX0 EX1	FastPath Double	1	
POP <i>reg32</i>	EX0 EX1	FastPath Single	1	
POP <i>reg64</i>	EX0 EX1	FastPath Single	1	
POP <i>mem</i>	EX0 EX1	FastPath Double	1	
POP DS	microcode	microcode	NA	
POP ES	microcode	microcode	NA	
POP FS	microcode	microcode	NA	
POP GS	microcode	microcode	NA	
POP SS	microcode	microcode	NA	
POPA	microcode	microcode	NA	
POPAD	microcode	microcode	NA	
POPCNT <i>reg16, reg16</i>	EX0 EX1	FastPath Single	4	
POPCNT <i>reg32, reg32</i>	EX0 EX1	FastPath Single	4	
POPCNT <i>reg64, reg64</i>	EX0 EX1	FastPath Single	4	
POPCNT <i>reg16, mem16</i>	EX0 EX1	FastPath Single	4	
POPCNT <i>reg32 mem32</i>	EX0 EX1	FastPath Single	4	
POPCNT <i>reg64 mem64</i>	EX0 EX1	FastPath Single	10	
POPF	microcode	microcode	NA	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
POPFD	microcode	microcode	NA	
POPFQ	microcode	microcode	NA	
PUSH <i>reg</i>	EX0 EX1	FastPath Single	1	
PUSH <i>imm</i>	EX0 EX1	FastPath Single	1	
PUSH <i>mem</i>	EX0 EX1	FastPath Double	1	
PUSH CS	EX0 EX1	FastPath Double	1	
PUSH DS	EX0 EX1	FastPath Double	1	
PUSH ES	EX0 EX1	FastPath Double	1	
PUSH FS	EX0 EX1	FastPath Double	1	
PUSH GS	EX0 EX1	FastPath Double	1	
PUSH SS	EX0 EX1	FastPath Double	1	
PUSHA	microcode	microcode	NA	
PUSHAD	microcode	microcode	NA	
PUSHF	microcode	microcode	NA	
PUSHFD	microcode	microcode	NA	
PUSHFQ	microcode	microcode	NA	
RCL <i>reg</i> , 1	EX0 EX1	FastPath Single	1	
RCL <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
RCL <i>reg</i> , CL	EX0 EX1	FastPath Single	1	
RCL <i>mem</i> , 1	EX0 EX1	FastPath Single	5	
RCL <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
RCL <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
RCR <i>reg</i> , 1	EX0 EX1	FastPath Single	1	
RCR <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
RCR <i>reg</i> , CL	EX0 EX1	FastPath Single	1	
RCR <i>mem</i> , 1	EX0 EX1	FastPath Single	5	
RCR <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
RCR <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
RET	EX0 EX1	FastPath Single	1	
RET <i>imm16</i>	EX0 EX1	microcode	NA	
ROL <i>reg</i> , 1	EX0 EX1	FastPath Single	1	
ROL <i>reg</i> , CL	EX0 EX1	FastPath Single	1	
ROL <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
ROL <i>mem</i> , 1	EX0 EX1	FastPath Single	5	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
ROL <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
ROL <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
ROR <i>reg</i> , 1	EX0 EX1	FastPath Single	1	
ROR <i>reg</i> , CL	EX0 EX1	FastPath Single	1	
ROR <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
ROR <i>mem</i> , 1	EX0 EX1	FastPath Single	5	
ROR <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
ROR <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
SAHF	EX0 EX1	FastPath Double	2	
SAL <i>reg</i>	EX0 EX1	FastPath Single	1	
SAL <i>reg</i> , 1	EX0 EX1	FastPath Single	1	
SAL <i>reg</i> , CL	EX0 EX1	FastPath Single	1	
SAL <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
SAL <i>mem</i> , 1	EX0 EX1	FastPath Single	5	
SAL <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
SAL <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
SHL <i>reg</i> , 1	EX0 EX1	FastPath Single	5	
SHL <i>reg</i> , CL	EX0 EX1	FastPath Single	5	
SHL <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
SHL <i>mem</i> , 1	EX0 EX1	FastPath Single	5	
SHL <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
SHL <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
SAR <i>reg</i> , 1	EX0 EX1	FastPath Single	1	
SAR <i>reg</i> , CL	EX0 EX1	FastPath Single	1	
SAR <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
SAR <i>mem</i> , 1	EX0 EX1	FastPath Single	5	
SAR <i>mem</i> , CL	EX0 EX1	FastPath Single	5	
SAR <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
SBB <i>reg</i> , <i>reg</i>	EX0 EX1	FastPath Single	1	
SBB <i>reg</i> , <i>imm</i>	EX0 EX1	FastPath Single	1	
SBB <i>mem</i> , <i>reg</i>	EX0 EX1	FastPath Single	5	
SBB <i>mem</i> , <i>imm</i>	EX0 EX1	FastPath Single	5	
SBB <i>reg</i> , <i>mem</i>	EX0 EX1	FastPath Single	5	
SCAS	microcode	microcode	NA	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
SCASB	microcode	microcode	NA	
SCASW	microcode	microcode	NA	
SCASD	microcode	microcode	NA	
SCASQ	microcode	microcode	NA	
SETcc <i>reg</i>	EX0 EX1	FastPath Single	1	
SETcc <i>mem</i>	EX0 EX1	FastPath Single	5	
SHLD <i>reg, reg, CL</i>	EX0 EX1	microcode	NA	
SHLD <i>reg, reg, imm</i>	EX0 EX1	microcode	NA	
SHLD <i>mem, reg, CL</i>	EX0 EX1	microcode	NA	
SHLD <i>mem, reg, imm</i>	EX0 EX1	microcode	NA	
SHR <i>reg, 1</i>	EX0 EX1	FastPath Single	1	
SHR <i>reg, CL</i>	EX0 EX1	FastPath Single	1	
SHR <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
SHR <i>mem, 1</i>	EX0 EX1	FastPath Single	5	
SHR <i>mem, CL</i>	EX0 EX1	FastPath Single	5	
SHR <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
SHRD <i>reg, reg, CL imm</i>	EX0 EX1	microcode	NA	
SHRD <i>reg, reg, imm</i>	EX0 EX1	microcode	NA	
SHRD <i>mem, reg, CL</i>	EX0 EX1	microcode	NA	
SHRD <i>mem, reg, imm</i>	EX0 EX1	microcode	NA	
SLWPCB <i>reg</i>	microcode	microcode	NA	
STC	EX0 EX1	FastPath Single	1	
STD	microcode	microcode	1	
STOS	microcode	microcode	NA	
STOSB	microcode	microcode	NA	
STOSW	microcode	microcode	NA	
STOSD	microcode	microcode	NA	
STOSQ	microcode	microcode	NA	
SUB <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
SUB <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
SUB <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
SUB <i>mem, imm</i>	EX0 EX1	FastPath Single	5	
SUB <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
T1MSKC <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
T1MSKC <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
TEST <i>reg, reg</i>	EX0 EX1	FastPath Single	1	If branch fused, op to EX1, else op to EX0 EX1.
TEST <i>reg, imm</i>	EX0 EX1	FastPath Single	1	If branch fused, op to EX1, else op to EX0 EX1.
TEST <i>mem, reg</i>	EX0 EX1	FastPath Single	5	If branch fused, op to EX1, else op to EX0 EX1.
TEST <i>mem, imm</i>	EX0 EX1	FastPath Single	5	If branch fused, op to EX1, else op to EX0 EX1.
TZCNT <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
TZCNT <i>reg, mem16</i>	EX0 EX1	FastPath Double	NA	
TZCNT <i>reg, mem32</i>	EX0 EX1	FastPath Double	NA	
TZCNT <i>reg, mem64</i>	EX0 EX1	FastPath Double	NA	
TZMSK <i>reg, reg</i>	EX0 EX1	FastPath Double	NA	
TZMSK <i>reg, mem</i>	EX0 EX1	FastPath Double	NA	
XADD <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
XADD <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
XCHG <i>reg8, reg8</i>	EX0 EX1	FastPath Double	1	
XCHG <i>reg16, reg16</i>	EX0 EX1	FastPath Double	1	
XCHG <i>reg32, reg32</i>	EX0 EX1	FastPath Double	1	
XCHG <i>reg64, reg64</i>	EX0 EX1	FastPath Double	1	
XCHG <i>reg8, mem8</i>	EX0 EX1	FastPath Double	5	
XCHG <i>reg16, mem16</i>	EX0 EX1	FastPath Double	5	
XCHG <i>reg32, mem32</i>	EX0 EX1	FastPath Double	5	
XCHG <i>reg64, mem64</i>	EX0 EX1	FastPath Double	5	
XCHG <i>mem8, reg8</i>	EX0 EX1	FastPath Double	5	
XCHG <i>mem16, reg16</i>	EX0 EX1	FastPath Double	5	
XCHG <i>mem32, reg32</i>	EX0 EX1	FastPath Double	5	
XCHG <i>mem64, reg64</i>	EX0 EX1	FastPath Double	5	
XLAT	microcode	microcode	NA	
XLATB	microcode	microcode	NA	
XOR <i>reg, reg</i>	EX0 EX1	FastPath Single	1	
XOR <i>reg, imm</i>	EX0 EX1	FastPath Single	1	
XOR <i>mem, reg</i>	EX0 EX1	FastPath Single	5	
XOR <i>reg, imm</i>	EX0 EX1	FastPath Single	5	

Table 10. General Purpose and Integer Instruction Latencies (Continued)

Syntax	Pipes	Decode Type	Latencies	Comments
XOR <i>reg, mem</i>	EX0 EX1	FastPath Single	5	
XRSTOR <i>mem</i>	microcode	microcode	NA	
XSAVE <i>mem</i>	microcode	microcode	NA	

B.3 System Instruction Latencies

The latency table for system instructions gives the decode type and latency corresponding to each instruction mnemonic. For more detailed information on the operation of a particular system instruction, as well as encoding information, see the *AMD64 Architecture Programmer's Manual Volume 3: General-Purpose and System Instructions*, order# 24594 .

Table 11. System Instruction Latencies

Syntax	Decode Type ¹	Latency
ARPL <i>reg16, reg16</i>	microcode	NA
ARPL <i>mem16, reg16</i>	microcode	NA
CLI	microcode	NA
CLTS	microcode	NA
INVLPG <i>mem8</i>	microcode	NA
IRETQ	microcode	NA
LAR <i>reg, reg</i>	microcode	NA
LAR <i>reg, mem</i>	microcode	NA
LGDT <i>mem32</i>	microcode	NA
LIDT <i>mem32</i>	microcode	NA
LLDT <i>reg32</i>	microcode	NA
LLDT <i>mem32</i>	microcode	NA
LMSW <i>reg</i>	microcode	NA
LMSW <i>mem</i>	microcode	NA
LSL <i>reg, reg16</i>	microcode	NA
LSL <i>reg, reg32</i>	microcode	NA
LSL <i>reg, reg64</i>	microcode	NA
LSL <i>reg, mem16</i>	microcode	NA
LSL <i>reg, mem32</i>	microcode	NA
LSL <i>reg, mem64</i>	microcode	NA
MONITOR	microcode	Variable
MOV CR0, <i>reg32</i>	microcode	NA
MOV CR0, <i>reg64</i>	microcode	NA
MOV CR2, <i>reg32</i>	microcode	NA
MOV CR4, <i>reg32</i>	microcode	NA
MOV CR4, <i>reg64</i>	microcode	NA
MOV CR8, <i>reg32</i>	microcode	NA

Table 11. System Instruction Latencies (Continued)

Syntax	Decode Type ¹	Latency
MOV DR0–3, <i>reg32</i>	microcode	NA
MOV DR0, <i>reg64</i>	microcode	NA
MOV DR6–7, <i>reg32</i>	microcode	NA
MOV DR6, <i>reg64</i>	microcode	NA
MOV <i>reg32</i> , CR0	microcode	NA
MOV <i>reg32</i> , CR2	microcode	NA
MOV <i>reg32</i> , CR3	microcode	NA
MOV <i>reg32</i> , CR4	microcode	NA
MOV <i>reg32</i> , CR8	microcode	NA
MOV <i>reg32</i> , DR0–3	microcode	NA
MOV <i>reg32</i> , DR6–7	microcode	NA
MOV <i>reg64</i> , CR0	microcode	NA
MOV <i>reg64</i> , CR3	microcode	NA
MOV <i>reg64</i> , CR4	microcode	NA
MOV <i>reg64</i> , DR0	microcode	NA
MOV <i>reg64</i> , DR6	microcode	NA
MWAIT	microcode	Variable
RDMSR APIC base	microcode	NA
RDMSR FS base	microcode	NA
RDMSR GS base	microcode	NA
RDMSR	microcode	Machine Dependent
RDPMC	microcode	NA
RDTSC	microcode	43
RDTSCP	microcode	NA
SGDT <i>mem</i>	microcode	NA
SIDT <i>mem</i>	microcode	NA
SLDT <i>reg</i>	microcode	NA
SLDT <i>mem</i>	microcode	NA
SMSW <i>reg</i>	microcode	NA
SMSW <i>mem</i>	microcode	NA
STI	microcode	NA
STR <i>reg</i>	microcode	NA
STR <i>mem</i>	microcode	NA

Table 11. System Instruction Latencies (Continued)

Syntax	Decode Type ¹	Latency
SWAPGS	microcode	NA
VERR <i>reg16</i>	microcode	NA
VERW <i>reg16</i>	microcode	NA
VERR <i>mem16</i>	microcode	NA
VERW <i>mem16</i>	microcode	NA
WRMSR APIC base	microcode	NA
WRMSR FS base	microcode	NA
WRMSR GS base	microcode	NA
WRMSR	microcode	NA
XSETBV	microcode	NA

B.4 FPU Instruction Latencies

The table that follow provides the name, pipes, decode type and latency for the 128-bit and 256-bit media instructions, comprised of the x87 amd SIMD instruction sets. For detailed information on the operation of these instructions, as well as opcodes, see the *AMD64 Architecture Programmer's Manual, Volume 4: 128-Bit Media Instructions*, order# 26568, *AMD64 Architecture Programmer's Manual Volume 5: 64-bit Media and x87 Floating-Point Instructions*, order# 26569, and *AMD64 Architecture Programmer's Manual Volume 6: 128-Bit XOP, FMA4 and CVT16 Instructions*, order# 43479.

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
ADDPD_mem	FMA[P0 P1]	FastPath Single	10
ADDPD_reg	FMA[P0 P1]	FastPath Single	5
ADDPS_mem	FMA[P0 P1]	FastPath Single	10
ADDPS_reg	FMA[P0 P1]	FastPath Single	5
ADDSD_mem	FMA[P0 P1]	FastPath Single	10
ADDSD_reg	FMA[P0 P1]	FastPath Single	5
ADDSS_mem	FMA[P0 P1]	FastPath Single	10
ADDSS_reg	FMA[P0 P1]	FastPath Single	5
ADDSUBPD_mem	FMA[P0 P1]	FastPath Single	10
ADDSUBPD_reg	FMA[P0 P1]	FastPath Single	5
ADDSUBPS_mem	FMA[P0 P1]	FastPath Single	10
ADDSUBPS_reg	FMA[P0 P1]	FastPath Single	5
AESDEC_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
AESDEC_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
AESDECLAST_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
AESDECLAST_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
AESENC_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
AESENC_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
AESENCLAST_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
AESENCLAST_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
AESIMC_mem	FMA[P0 P1]	FastPath Single	9
AESIMC_reg	FMA[P0 P1]	FastPath Single	5
AESKEYGENASSIST_mem	FMA[P0 P1]	FastPath Single	9
AESKEYGENASSIST_reg	FMA[P0 P1]	FastPath Single	5
ANDNPD_mem	MAL[P2 P3]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
ANDNPD_reg	MAL[P2 P3]	FastPath Single	2
ANDNPS_mem	MAL[P2 P3]	FastPath Single	6
ANDNPS_reg	MAL[P2 P3]	FastPath Single	2
ANDPD_mem	MAL[P2 P3]	FastPath Single	6
ANDPD_reg	MAL[P2 P3]	FastPath Single	2
ANDPS_mem	MAL[P2 P3]	FastPath Single	6
ANDPS_reg	MAL[P2 P3]	FastPath Single	2
BLENDPD_mem	MAL[P2 P3]	FastPath Single	6
BLENDPD_reg	MAL[P2 P3]	FastPath Single	2
BLENDPS_mem	MAL[P2 P3]	FastPath Single	6
BLENDPS_reg	MAL[P2 P3]	FastPath Single	2
BLENDVPD_mem	XBR[P1]	FastPath Single	6
BLENDVPD_reg	XBR[P1]	FastPath Single	2
BLENDVPS_mem	XBR[P1]	FastPath Single	6
BLENDVPS_reg	XBR[P1]	FastPath Single	2
CMPPD_mem	FMA[P0 P1]	FastPath Single	5
CMPPD_reg	FMA[P0 P1]	FastPath Single	2
CMPPS_mem	FMA[P0 P1]	FastPath Single	5
CMPPS_reg	FMA[P0 P1]	FastPath Single	2
CMPSD_mem	FMA[P0 P1]	FastPath Single	5
CMPSD_reg	FMA[P0 P1]	FastPath Single	2
CMPSS_mem	FMA[P0 P1]	FastPath Single	5
CMPSS_reg	FMA[P0 P1]	FastPath Single	2
COMISD_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
COMISD_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
COMISS_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
COMISS_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
CRC32_reg16_mem16	microcode	microcode	6
CRC32_reg16_reg16	microcode	microcode	2
CRC32_reg32_mem32	microcode	microcode	10
CRC32_reg32_reg32	microcode	microcode	6
CRC32_reg8_mem8	microcode	microcode	6
CRC32_reg8_reg8	microcode	microcode	2
CVTDQ2PD_mem	XBR[P1]/CVT[P0]	FastPath Double	8/2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
CVTDQ2PD_reg	XBR[P1]/CVT[P0]	FastPath Double	4/2
CVTDQ2PS_mem	CVT[P0]	FastPath Single	8
CVTDQ2PS_reg	CVT[P0]	FastPath Single	4
CVTPD2DQ_mem	CVT[P0]/XBR[P1]	FastPath Double	6/2
CVTPD2DQ_reg	CVT[P0]/XBR[P1]	FastPath Double	2/2
CVTPD2PI_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
CVTPD2PI_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
CVTPD2PS_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
CVTPD2PS_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
CVTPI2PD_mem	XBR[P1]/CVT[P0]	FastPath Double	6/4
CVTPI2PD_reg	XBR[P1]/CVT[P0]	FastPath Double	2/4
CVTPI2PS_mem	CVT[P0]	FastPath Single	8
CVTPI2PS_reg	CVT[P0]	FastPath Single	4
CVTPS2DQ_mem	CVT[P0]	FastPath Single	8
CVTPS2DQ_reg	CVT[P0]	FastPath Single	4
CVTPS2PD_mem	XBR[P1]/CVT[P0]	FastPath Double	6/4
CVTPS2PD_reg	XBR[P1]/CVT[P0]	FastPath Double	2/4
CVTPS2PI_mem	CVT[P0]	FastPath Single	8
CVTPS2PI_reg	CVT[P0]	FastPath Single	4
CVTSD2SI_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTSD2SI_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTSD2SI_reg32	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTSD2SI_reg64	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTSD2SS_mem	CVT[P0]	FastPath Single	8
CVTSD2SS_reg	CVT[P0]	FastPath Single	4
CVTSI2SD_mem32	CVT[P0]	FastPath Single	8
CVTSI2SD_mem64	CVT[P0]	FastPath Single	8
CVTSI2SD_reg32	CVT[P0]	FastPath Double	4
CVTSI2SD_reg64	CVT[P0]	FastPath Double	4
CVTSI2SS_mem32	CVT[P0]	FastPath Single	8
CVTSI2SS_mem64	CVT[P0]	FastPath Single	8
CVTSI2SS_reg32	CVT[P0]	FastPath Double	4
CVTSI2SS_reg64	CVT[P0]	FastPath Double	4
CVTSS2SD_mem	CVT[P0]	FastPath Single	8

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
CVTSS2SD_reg	CVT[P0]	FastPath Single	4
CVTSS2SI_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTSS2SI_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTSS2SI_reg32	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTSS2SI_reg64	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTTPD2DQ_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
CVTTPD2DQ_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
CVTTPD2PI_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
CVTTPD2PI_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
CVTTPS2DQ_mem	CVT[P0]	FastPath Single	8
CVTTPS2DQ_reg	CVT[P0]	FastPath Single	4
CVTTPS2PI_mem	CVT[P0]	FastPath Single	8
CVTTPS2PI_reg	CVT[P0]	FastPath Single	4
CVTTSD2SI_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTTSD2SI_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTTSD2SI_reg32	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTTSD2SI_reg64	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTTSS2SI_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTTSS2SI_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
CVTTSS2SI_reg32	CVT[P0]/STO[P3]	FastPath Double	4/2
CVTTSS2SI_reg64	CVT[P0]/STO[P3]	FastPath Double	4/2
DIVPD_mem	FMA[P0 P1]	FastPath Single	31
DIVPD_reg	FMA[P0 P1]	FastPath Single	27
DIVPS_mem	FMA[P0 P1]	FastPath Single	28
DIVPS_reg	FMA[P0 P1]	FastPath Single	24
DIVSD_mem	FMA[P0 P1]	FastPath Single	31
DIVSD_reg	FMA[P0 P1]	FastPath Single	27
DIVSS_mem	FMA[P0 P1]	FastPath Single	28
DIVSS_reg	FMA[P0 P1]	FastPath Single	24
DPPD_mem	microcode	microcode	19
DPPD_reg	microcode	microcode	15
DPPS_mem	microcode	microcode	29
DPPS_reg	microcode	microcode	25
EMMS	microcode	FastPath Single	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
EXTRACTPS_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
EXTRACTPS_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
EXTRQ_reg	XBR[P1]	FastPath Single	3
F2XM1_reg	microcode	microcode	189
FABS_reg	FMA[P0 P1]	FastPath Single	2
FADD_mem	FMA[P0 P1]	FastPath Single	10
FADD_reg	FMA[P0 P1]	FastPath Single	5
FADDP_reg	FMA[P0 P1]	FastPath Single	5
FBLD_mem	microcode	microcode	NA
FBSTP_mem	microcode	microcode	NA
FCHS_reg	FMA[P0 P1]	FastPath Single	2
FCLEX_reg	microcode	microcode	NA
FCMOVB_reg	microcode	microcode	NA
FCMOVBE_reg	microcode	microcode	NA
FCMOVE_reg	microcode	microcode	NA
FCMOVNB_reg	microcode	microcode	NA
FCMOVNBE_reg	microcode	microcode	NA
FCMOVNE_reg	microcode	microcode	NA
FCMOVNU_reg	microcode	microcode	NA
FCMOVU_reg	microcode	microcode	NA
FCOM_mem	FMA[P0 P1]	FastPath Single	5
FCOM_reg	FMA[P0 P1]	FastPath Single	2
FCOM2_reg	FMA[P0 P1]	FastPath Single	2
FCOMI_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
FCOMIP_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
FCOMP_mem	FMA[P0 P1]	FastPath Single	5
FCOMP_reg	FMA[P0 P1]	FastPath Single	2
FCOMP3_reg	FMA[P0 P1]	FastPath Single	2
FCOMP5_reg	FMA[P0 P1]	FastPath Single	2
FCOMPP_reg	FMA[P0 P1]	FastPath Single	2
FCOS_reg	microcode	microcode	151
FDECSTP_reg	None	FastPath Single	0
FDISI_reg	microcode	microcode	NA
FDIV_mem	FMA[P0 P1]	FastPath Single	46

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
FDIV_reg	FMA[P0 P1]	FastPath Single	42
FDIVP_reg	FMA[P0 P1]	FastPath Single	42
FDIVR_mem	FMA[P0 P1]	FastPath Single	46
FDIVR_reg	FMA[P0 P1]	FastPath Single	42
FDIVRP_reg	FMA[P0 P1]	FastPath Single	42
FENI_reg	microcode	microcode	NA
FFREE_reg	None	FastPath Single	0
FFREEP_reg	None	FastPath Single	0
FIADD_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
FICOM_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/2
FICOMP_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/2
FIDIV_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	42
FIDIVR_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	42
FILD_mem	CVT[P0]	FastPath Single	4
FIMUL_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
FINCSTP_reg	None	FastPath Single	0
FINIT_reg	microcode	microcode	NA
FIST_mem	CVT[P0]/STO[P3]	FastPath Double	4/2
FISTP_mem	CVT[P0]/STO[P3]	FastPath Double	4/2
FISTTP_mem	CVT[P0]/STO[P3]	FastPath Double	4/2
FISUB_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
FISUBR_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
FLD_mem	FMA[P0 P1]	FastPath Single	5
FLD_reg	FMA[P0 P1]	FastPath Single	2
FLD1_reg	CVT[P0]	FastPath Single	4
FLDCW_mem	microcode	microcode	NA
FLDENV_mem	microcode	microcode	NA
FLDL2E_reg	CVT[P0]	FastPath Single	4
FLDL2T_reg	CVT[P0]	FastPath Single	4
FLDLG2_reg	CVT[P0]	FastPath Single	4
FLDLN2_reg	CVT[P0]	FastPath Single	4
FLDPI_reg	CVT[P0]	FastPath Single	4
FLDZ_reg	CVT[P0]	FastPath Single	4
FMADDPD_mem	FMA[P0 P1]	FastPath Single	10

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
FMADDPD_reg	FMA[P0 P1]	FastPath Single	5
FMADDPD_mem	FMA[P0 P1]	FastPath Single	10
FMADDPS_reg	FMA[P0 P1]	FastPath Single	5
FMADDPS_mem	FMA[P0 P1]	FastPath Single	10
FMADDSD_reg	FMA[P0 P1]	FastPath Single	5
FMADDSD_mem	FMA[P0 P1]	FastPath Single	10
FMADDSS_reg	FMA[P0 P1]	FastPath Single	5
FMADDSS_mem	FMA[P0 P1]	FastPath Single	10
FMSUBPD_reg	FMA[P0 P1]	FastPath Single	5
FMSUBPD_mem	FMA[P0 P1]	FastPath Single	10
FMSUBPS_reg	FMA[P0 P1]	FastPath Single	5
FMSUBPS_mem	FMA[P0 P1]	FastPath Single	10
FMSUBSD_reg	FMA[P0 P1]	FastPath Single	5
FMSUBSD_mem	FMA[P0 P1]	FastPath Single	10
FMSUBSS_reg	FMA[P0 P1]	FastPath Single	5
FMSUBSS_mem	FMA[P0 P1]	FastPath Single	10
FMUL_reg	FMA[P0 P1]	FastPath Single	5
FMUL_mem	FMA[P0 P1]	FastPath Double	10
FMULP_reg	FMA[P0 P1]	FastPath Single	5
FNCLEX_reg	microcode	microcode	NA
FNINIT_reg	microcode	microcode	NA
FNMADDPD_mem	FMA[P0 P1]	FastPath Single	10
FNMADDPD_reg	FMA[P0 P1]	FastPath Single	5
FNMADDPS_mem	FMA[P0 P1]	FastPath Single	10
FNMADDPS_reg	FMA[P0 P1]	FastPath Single	5
FNMADDSD_mem	FMA[P0 P1]	FastPath Single	10
FNMADDSD_reg	FMA[P0 P1]	FastPath Single	5
FNMADDSS_mem	FMA[P0 P1]	FastPath Single	10
FNMADDSS_reg	FMA[P0 P1]	FastPath Single	5
FNMSUBPD_mem	FMA[P0 P1]	FastPath Single	10
FNMSUBPD_reg	FMA[P0 P1]	FastPath Single	5
FNMSUBPS_mem	FMA[P0 P1]	FastPath Single	10
FNMSUBPS_reg	FMA[P0 P1]	FastPath Single	5
FNMSUBSD_mem	FMA[P0 P1]	FastPath Single	10
FNMSUBSD_reg	FMA[P0 P1]	FastPath Single	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
FNMSUBSS_mem	FMA[P0 P1]	FastPath Single	10
FNMSUBSS_reg	FMA[P0 P1]	FastPath Single	5
FNOP_reg	None	FastPath Single	0
FNSAVE_mem	microcode	microcode	NA
FNSTCW_mem	microcode	microcode	NA
FNSTENV_mem	microcode	microcode	NA
FNSTSW_mem	microcode	microcode	NA
FNSTSW_reg	microcode	microcode	NA
FPATAN_reg	microcode	microcode	433
FPREM_reg	FMA[P0 P1]	FastPath Single	10
FPREM1_reg	FMA[P0 P1]	FastPath Single	10
FPTAN_reg	microcode	FastPath Single	240
FRNDINT_reg	CVT[P0]	FastPath Single	4
FRSTOR_mem	microcode	microcode	NA
FSAVE_mem	microcode	microcode	NA
FSCALE_reg	microcode	microcode	NA
FSETPM_reg	microcode	microcode	NA
FSIN_reg	microcode	microcode	148
FSINCOS_reg	microcode	microcode	143
FSQRT_reg	FMA[P0 P1]	FastPath Single	52
FST_mem	CVT[P0]/STO[P3]	FastPath Double	4/2
FST_reg	FMA[P0 P1]	FastPath Single	2
FSTCW_mem	microcode	microcode	NA
FSTENV_mem	microcode	microcode	NA
FSTP_mem	CVT[P0]/STO[P3]	FastPath Double	4/2
FSTP_reg	FMA[P0 P1]	FastPath Single	2
FSTP1_reg	FMA[P0 P1]	FastPath Single	2
FSTP8_reg	FMA[P0 P1]	FastPath Single	2
FSTP9_reg	FMA[P0 P1]	FastPath Single	2
FSTSW_mem	microcode	microcode	NA
FSTSW_reg	microcode	microcode	NA
FSUB_mem	FMA[P0 P1]	FastPath Single	10
FSUB_reg	FMA[P0 P1]	FastPath Single	5
FSUBP_reg	FMA[P0 P1]	FastPath Single	10

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
FSUBR_mem	FMA[P0 P1]	FastPath Single	5
FSUBR_reg	FMA[P0 P1]	FastPath Single	5
FSUBRP_reg	FMA[P0 P1]	FastPath Single	5
FTST_reg	FMA[P0 P1]	FastPath Single	2
FUCOM_reg	FMA[P0 P1]	FastPath Single	2
FUCOMI_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
FUCOMIP_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
FUCOMP_reg	FMA[P0 P1]	FastPath Single	2
FUCOMPP_reg	FMA[P0 P1]	FastPath Single	2
FWAIT(WAIT)_reg	None	FastPath Single	0
FXAM_reg	FMA[P0 P1]	FastPath Single	2
FXCH_reg	FMA[P0 P1]	FastPath Single	2
FXCH4_reg	FMA[P0 P1]	FastPath Single	2
FXCH7_reg	FMA[P0 P1]	FastPath Single	2
FXRSTOR_mem	microcode	microcode	NA
FXSAVE_mem	microcode	microcode	NA
FXTRACT_reg	microcode	microcode	NA
FYL2X_reg	microcode	microcode	NA
FYL2XP1_reg	microcode	microcode	241
HADDPD_mem	microcode	microcode	15
HADDPD_reg	microcode	microcode	11
HADDPS_mem	microcode	microcode	15
HADDPS_reg	microcode	microcode	11
HSUBPD_mem	microcode	microcode	15
HSUBPD_reg	microcode	microcode	11
HSUBPS_mem	microcode	microcode	15
HSUBPS_reg	microcode	microcode	11
INSERTPS_mem	XBR[P1]	FastPath Single	6
INSERTPS_reg	XBR[P1]	FastPath Single	2
INSERTQ_reg	XBR[P1]	FastPath Single	3
LDDQU_mem	None	FastPath Single	4
LDMXCSR_mem	microcode	microcode	NA
MASKMOVDQU_reg	microcode	microcode	NA
MASKMOVQ_mmx_reg	microcode	microcode	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
MAXPD_mem	FMA[P0 P1]	FastPath Single	5
MAXPD_reg	FMA[P0 P1]	FastPath Single	2
MAXPS_mem	FMA[P0 P1]	FastPath Single	5
MAXPS_reg	FMA[P0 P1]	FastPath Single	2
MAXSD_mem	FMA[P0 P1]	FastPath Single	5
MAXSD_reg	FMA[P0 P1]	FastPath Single	2
MAXSS_mem	FMA[P0 P1]	FastPath Single	5
MAXSS_reg	FMA[P0 P1]	FastPath Single	2
MFENCE_reg	microcode	microcode	NA
MINPD_mem	FMA[P0 P1]	FastPath Single	5
MINPD_reg	FMA[P0 P1]	FastPath Single	2
MINPS_mem	FMA[P0 P1]	FastPath Single	5
MINPS_reg	FMA[P0 P1]	FastPath Single	2
MINSD_mem	FMA[P0 P1]	FastPath Single	5
MINSD_reg	FMA[P0 P1]	FastPath Single	2
MINSS_mem	FMA[P0 P1]	FastPath Single	5
MINSS_reg	FMA[P0 P1]	FastPath Single	2
MOVAPD_mem	None	FastPath Single	4
MOVAPD_reg	MAL[P2 P3]	FastPath Single	0
MOVAPS_mem	None	FastPath Single	4
MOVAPS_reg	MAL[P2 P3]	FastPath Single	0
MOVD_mem32_reg32	None	FastPath Single	4
MOVD_mem64_reg64	None	FastPath Single	4
MOVD_mmx_mem32	None	FastPath Single	4
MOVD_mmx_mem64	STO[P3]	FastPath Single	4
MOVD_reg32_xmm	STO[P3]	FastPath Single	2
MOVD_reg64_xmm	STO[P3]	FastPath Single	2
MOVD_xmm_reg32	None	FastPath Double	8
MOVD_xmm_reg64	None	FastPath Double	8
MOVDDUP_mem	None	FastPath Single	4
MOVDDUP_reg	XBR[P1]	FastPath Single	2
MOVDQ2Q_mem	XBR[P1]	FastPath Single	4
MOVDQ2Q_reg	MAL[P2 P3]	FastPath Single	2
MOVDQA_mem	None	FastPath Single	4

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
MOVDQA_reg	MAL[P2 P3]	FastPath Single	2
MOVDQU_mem	None	FastPath Single	4
MOVDQU_reg	MAL[P2 P3]	FastPath Single	2
MOVHLPD_reg	XBR[P1]	FastPath Single	2
MOVHPD_xmm_mem	XBR[P1]	FastPath Single	4
MOVHPD_mem_xmm	XBR[P1]	FastPath Double	4
MOVHPS_xmm_mem	XBR[P1]	FastPath Single	4
MOVHPS_mem_xmm	XBR[P1]	FastPath Double	4
MOVLHPS_reg	XBR[P1]	FastPath Single	2
MOVLPD_mem	FMA[P0 P1]	FastPath Single	4
MOVLPS_mem	FMA[P0 P1]	FastPath Single	4
MOVMSKPD_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
MOVMSKPS_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
MOVNTDQ_mem	STO[P3]	FastPath Single	4
MOVNTDQA_mem	STO[P3]	FastPath Single	4
MOVNTPD_mem	STO[P3]	FastPath Single	4
MOVNTPS_mem	STO[P3]	FastPath Single	4
MOVNTQ_mmx_mem	STO[P3]	FastPath Single	4
MOVNTSD_mem	STO[P3]	FastPath Single	4
MOVNTSS_mem	STO[P3]	FastPath Single	4
MOVQ_mem	MAL[P2 P3]	FastPath Single	4
MOVQ_mmx_mem	MAL[P2 P3]	FastPath Single	4
MOVQ_mmx_reg	MAL[P2 P3]	FastPath Single	2
MOVQ_reg	MAL[P2 P3]	FastPath Single	2
MOVQ2DQ_reg	MAL[P2 P3]	FastPath Single	2
MOVSD_mem	FMA[P0 P1]	FastPath Single	4
MOVSD_reg	FMA[P0 P1]	FastPath Single	2
MOVSHDUP_mem	XBR[P1]	FastPath Single	4
MOVSHDUP_reg	XBR[P1]	FastPath Single	2
MOVSLDUP_mem	XBR[P1]	FastPath Single	4
MOVSLDUP_reg	XBR[P1]	FastPath Single	2
MOVSS_mem	FMA[P0 P1]	FastPath Single	4
MOVSS_reg	FMA[P0 P1]	FastPath Single	2
MOVUPD_mem	MAL[P2 P3]	FastPath Single	4

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
MOVUPD_reg	MAL[P2 P3]	FastPath Single	0
MOVUPS_mem	MAL[P2 P3]	FastPath Single	4
MOVUPS_reg	MAL[P2 P3]	FastPath Single	0
MPSADBW_mem	microcode	microcode	12
MPSADBW_reg	microcode	microcode	8
MULPD_mem	FMA[P0 P1]	FastPath Single	10
MULPD_reg	FMA[P0 P1]	FastPath Single	5
MULPS_mem	FMA[P0 P1]	FastPath Single	10
MULPS_reg	FMA[P0 P1]	FastPath Single	5
MULSD_mem	FMA[P0 P1]	FastPath Single	10
MULSD_reg	FMA[P0 P1]	FastPath Single	5
MULSS_mem	FMA[P0 P1]	FastPath Single	10
MULSS_reg	FMA[P0 P1]	FastPath Single	5
ORPD_mem	MAL[P2 P3]	FastPath Single	6
ORPD_reg	MAL[P2 P3]	FastPath Single	2
ORPS_mem	MAL[P2 P3]	FastPath Single	6
ORPS_reg	MAL[P2 P3]	FastPath Single	2
PABSB_mem	MAL[P2 P3]	FastPath Single	6
PABSB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PABSB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PABSB_reg	MAL[P2 P3]	FastPath Single	2
PABSD_mem	MAL[P2 P3]	FastPath Single	6
PABSD_mmx_mem	MAL[P2 P3]	FastPath Single	6
PABSD_mmx_reg	MAL[P2 P3]	FastPath Single	2
PABSD_reg	MAL[P2 P3]	FastPath Single	2
PABSW_mem	MAL[P2 P3]	FastPath Single	6
PABSW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PABSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PABSW_reg	MAL[P2 P3]	FastPath Single	2
PACKSSDW_mem	XBR[P1]	FastPath Single	6
PACKSSDW_mmx_mem	XBR[P1]	FastPath Single	6
PACKSSDW_mmx_reg	XBR[P1]	FastPath Single	2
PACKSSDW_reg	XBR[P1]	FastPath Single	2
PACKSSWB_mem	XBR[P1]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PACKSSWB_mmx_mem	XBR[P1]	FastPath Single	6
PACKSSWB_mmx_reg	XBR[P1]	FastPath Single	2
PACKSSWB_reg	XBR[P1]	FastPath Single	2
PACKUSDW_mem	XBR[P1]	FastPath Single	6
PACKUSDW_reg	XBR[P1]	FastPath Single	2
PACKUSWB_mem	XBR[P1]	FastPath Single	6
PACKUSWB_mmx_mem	XBR[P1]	FastPath Single	6
PACKUSWB_mmx_reg	XBR[P1]	FastPath Single	2
PACKUSWB_reg	XBR[P1]	FastPath Single	2
PADDB_mem	MAL[P2 P3]	FastPath Single	6
PADDB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDB_reg	MAL[P2 P3]	FastPath Single	2
PADDD_mem	MAL[P2 P3]	FastPath Single	6
PADDD_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDD_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDD_reg	MAL[P2 P3]	FastPath Single	2
PADDQ_mem	MAL[P2 P3]	FastPath Single	6
PADDQ_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDQ_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDQ_reg	MAL[P2 P3]	FastPath Single	2
PADDSB_mem	MAL[P2 P3]	FastPath Single	6
PADDSB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDSB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDSB_reg	MAL[P2 P3]	FastPath Single	2
PADDSW_mem	MAL[P2 P3]	FastPath Single	6
PADDSW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDSW_reg	MAL[P2 P3]	FastPath Single	2
PADDUSB_mem	MAL[P2 P3]	FastPath Single	6
PADDUSB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDUSB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDUSB_reg	MAL[P2 P3]	FastPath Single	2
PADDUSW_mem	MAL[P2 P3]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PADDUSW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDUSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDUSW_reg	MAL[P2 P3]	FastPath Single	2
PADDW_mem	MAL[P2 P3]	FastPath Single	6
PADDW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PADDW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PADDW_reg	MAL[P2 P3]	FastPath Single	2
PALIGNR_mem	XBR[P1]	FastPath Single	6
PALIGNR_mmx_mem	XBR[P1]	FastPath Single	6
PALIGNR_mmx_reg	XBR[P1]	FastPath Single	2
PALIGNR_reg	XBR[P1]	FastPath Single	2
PAND_mem	MAL[P2 P3]	FastPath Single	6
PAND_mmx_mem	MAL[P2 P3]	FastPath Single	6
PAND_mmx_reg	MAL[P2 P3]	FastPath Single	2
PAND_reg	MAL[P2 P3]	FastPath Single	2
PANDN_mem	MAL[P2 P3]	FastPath Single	6
PANDN_mmx_mem	MAL[P2 P3]	FastPath Single	6
PANDN_mmx_reg	MAL[P2 P3]	FastPath Single	2
PANDN_reg	MAL[P2 P3]	FastPath Single	2
PAVGB_mem	MAL[P2 P3]	FastPath Single	6
PAVGB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PAVGB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PAVGB_reg	MAL[P2 P3]	FastPath Single	2
PAVGW_mem	MAL[P2 P3]	FastPath Single	6
PAVGW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PAVGW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PAVGW_reg	MAL[P2 P3]	FastPath Single	2
PBLENDVB_mem	XBR[P1]	FastPath Single	6
PBLENDVB_reg	XBR[P1]	FastPath Single	2
PBLENDW_mem	MAL[P2 P3]	FastPath Single	6
PBLENDW_reg	MAL[P2 P3]	FastPath Single	2
PCLMULQDQ_mem	microcode	microcode	16
PCLMULQDQ_reg	microcode	microcode	12
PCMPEQB_mem	MAL[P2 P3]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PCMPEQB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PCMPEQB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PCMPEQB_reg	MAL[P2 P3]	FastPath Single	2
PCMPEQD_mem	MAL[P2 P3]	FastPath Single	6
PCMPEQD_mmx_mem	MAL[P2 P3]	FastPath Single	6
PCMPEQD_mmx_reg	MAL[P2 P3]	FastPath Single	2
PCMPEQD_reg	MAL[P2 P3]	FastPath Single	2
PCMPEQQ_mem	MAL[P2 P3]	FastPath Single	6
PCMPEQQ_reg	MAL[P2 P3]	FastPath Single	2
PCMPEQW_mem	MAL[P2 P3]	FastPath Single	6
PCMPEQW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PCMPEQW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PCMPEQW_reg	MAL[P2 P3]	FastPath Single	2
PCMPESTRM_mem	microcode	microcode	34
PCMPESTRM_reg	microcode	microcode	30
PCMPESTRM_mem	microcode	microcode	NA
PCMPESTRM_reg	microcode	microcode	NA
PCMPGTB_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PCMPGTB_reg	MAL[P2 P3]	FastPath Single	2
PCMPGTD_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTD_mmx_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTD_mmx_reg	MAL[P2 P3]	FastPath Single	2
PCMPGTD_reg	MAL[P2 P3]	FastPath Single	2
PCMPGTQ_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTQ_reg	MAL[P2 P3]	FastPath Single	2
PCMPGTW_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PCMPGTW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PCMPGTW_reg	MAL[P2 P3]	FastPath Single	2
PCMPISTRM_mr	microcode	microcode	10
PCMPISTRM_mr	microcode	microcode	NA
PEXTRB_mem	XBR[P1]/STO[P3]	FastPath Double	6/2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PEXTRB_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
PEXTRD_mem32_xmm	XBR[P1]/STO[P3]	FastPath Double	6/2
PEXTRD_reg32_xmm	XBR[P1]/STO[P3]	FastPath Double	2/2
PEXTRQ_mem64_xmm	XBR[P1]/STO[P3]	FastPath Double	6/2
PEXTRQ_reg64_xmm	XBR[P1]/STO[P3]	FastPath Double	2/2
PEXTRW_mem	XBR[P1]/STO[P3]	FastPath Double	6/6
PEXTRW_mmx_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
PEXTRW_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
PHADDD_mem	microcode	microcode	9
PHADDD_mmx_mem	microcode	microcode	9
PHADDD_mmx_reg	microcode	microcode	5
PHADDD_reg	microcode	microcode	5
PHADDSW_mem	microcode	microcode	9
PHADDSW_mmx_mem	microcode	microcode	9
PHADDSW_mmx_reg	microcode	microcode	5
PHADDSW_reg	microcode	microcode	5
PHADDW_mem	microcode	microcode	9
PHADDW_mmx_mem	microcode	microcode	NA
PHADDW_mmx_reg	microcode	microcode	NA
PHADDW_reg	microcode	microcode	5
PHMINPOSUW_mem	MAL[P2 P3] XBR[P1]	FastPath Single	6/2
PHMINPOSUW_reg	MAL[P2 P3] XBR[P1]	FastPath Single	2/2
PHSUBD_mem	microcode	microcode	9
PHSUBD_mmx_mem	microcode	microcode	9
PHSUBD_mmx_reg	microcode	microcode	5
PHSUBD_reg	microcode	microcode	5
PHSUBSW_mem	microcode	microcode	9
PHSUBSW_mmx_mem	microcode	microcode	9
PHSUBSW_mmx_reg	microcode	microcode	5
PHSUBSW_reg	microcode	microcode	5
PHSUBW_mem	microcode	microcode	9
PHSUBW_mmx_mem	microcode	microcode	9
PHSUBW_mmx_reg	microcode	microcode	5
PHSUBW_reg	microcode	microcode	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PINSRB_mem	XBR[P1]	FastPath Single	6
PINSRB_reg	XBR[P1]	FastPath Single	2
PINSRD_OP32_mem	XBR[P1]	FastPath Single	6
PINSRD_OP32_reg	XBR[P1]	FastPath Single	2
PINSRQ_OP64_mem	XBR[P1]	FastPath Single	6
PINSRQ_OP64_reg	XBR[P1]	FastPath Single	2
PINSRW_mem	XBR[P1]	FastPath Single	6
PINSRW_mmx_mem	XBR[P1]	FastPath Single	6
PINSRW_mmx_reg	XBR[P1]	FastPath Single	2
PINSRW_reg	XBR[P1]	FastPath Double	2
PMADDUBSW_mem	MMA[P0]	FastPath Single	8
PMADDUBSW_mmx_mem	MMA[P0]	FastPath Single	8
PMADDUBSW_mmx_reg	MMA[P0]	FastPath Single	4
PMADDUBSW_reg	MMA[P0]	FastPath Single	4
PMADDWD_mem	MMA[P0]	FastPath Single	8
PMADDWD_mmx_mem	MMA[P0]	FastPath Single	8
PMADDWD_mmx_reg	MMA[P0]	FastPath Single	4
PMADDWD_reg	MMA[P0]	FastPath Single	4
PMAXSB_mem	MAL[P2 P3]	FastPath Single	6
PMAXSB_reg	MAL[P2 P3]	FastPath Single	2
PMAXSD_mem	MAL[P2 P3]	FastPath Single	6
PMAXSD_reg	MAL[P2 P3]	FastPath Single	2
PMAXSW_mem	MAL[P2 P3]	FastPath Single	6
PMAXSW_mmx__mem	MAL[P2 P3]	FastPath Single	6
PMAXSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PMAXSW_reg	MAL[P2 P3]	FastPath Single	2
PMAXUB_mem	MAL[P2 P3]	FastPath Single	6
PMAXUB_mmx__mem	MAL[P2 P3]	FastPath Single	6
PMAXUB_mmx__reg	MAL[P2 P3]	FastPath Single	2
PMAXUB_reg	MAL[P2 P3]	FastPath Single	2
PMAXUD_mem	MAL[P2 P3]	FastPath Single	6
PMAXUD_reg	MAL[P2 P3]	FastPath Single	2
PMAXUW_mem	MAL[P2 P3]	FastPath Single	6
PMAXUW_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PMINSB_mem	MAL[P2 P3]	FastPath Single	6
PMINSB_reg	MAL[P2 P3]	FastPath Single	2
PMINSD_mem	MAL[P2 P3]	FastPath Single	6
PMINSD_reg	MAL[P2 P3]	FastPath Single	2
PMINSW_mem	MAL[P2 P3]	FastPath Single	6
PMINSW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PMINSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PMINSW_reg	MAL[P2 P3]	FastPath Single	2
PMINUB_mem	MAL[P2 P3]	FastPath Single	6
PMINUB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PMINUB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PMINUB_reg	MAL[P2 P3]	FastPath Single	2
PMINUD_mem	MAL[P2 P3]	FastPath Single	6
PMINUD_reg	MAL[P2 P3]	FastPath Single	2
PMINUW_mem	MAL[P2 P3]	FastPath Single	6
PMINUW_reg	MAL[P2 P3]	FastPath Single	2
PMOVMASKB_mmx_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
PMOVMASKB_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
PMOVXBD_mem	XBR[P1]	FastPath Single	4
PMOVXBD_reg	XBR[P1]	FastPath Single	2
PMOVXBDQ_mem	XBR[P1]	FastPath Single	4
PMOVXBDQ_reg	XBR[P1]	FastPath Single	2
PMOVXBQ_mem	XBR[P1]	FastPath Single	4
PMOVXBQ_reg	XBR[P1]	FastPath Single	2
PMOVXBW_mem	XBR[P1]	FastPath Single	4
PMOVXBW_reg	XBR[P1]	FastPath Single	2
PMOVXDQ_mem	XBR[P1]	FastPath Single	4
PMOVXDQ_reg	XBR[P1]	FastPath Single	2
PMOVXWD_mem	XBR[P1]	FastPath Single	4
PMOVXWD_reg	XBR[P1]	FastPath Single	2
PMOVXWQ_mem	XBR[P1]	FastPath Single	4
PMOVXWQ_reg	XBR[P1]	FastPath Single	2
PMOVZBD_mem	XBR[P1]	FastPath Single	4
PMOVZBD_reg	XBR[P1]	FastPath Single	2
PMOVZBQ_mem	XBR[P1]	FastPath Single	4
PMOVZBQ_reg	XBR[P1]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PMOVZXBW_mem	XBR[P1]	FastPath Single	4
PMOVZXBW_reg	XBR[P1]	FastPath Single	2
PMOVZXDQ_mem	XBR[P1]	FastPath Single	4
PMOVZXDQ_reg	XBR[P1]	FastPath Single	2
PMOVZXWD_mem	XBR[P1]	FastPath Single	4
PMOVZXWD_reg	XBR[P1]	FastPath Single	2
PMOVZXWQ_mem	XBR[P1]	FastPath Single	4
PMOVZXWQ_reg	XBR[P1]	FastPath Single	2
PMULDQ_mem	MMA[P0]	FastPath Single	8
PMULDQ_reg	MMA[P0]	FastPath Single	4
PMULHSW_mem	MMA[P0]	FastPath Single	8
PMULHSW_mmx_mem	MMA[P0]	FastPath Single	8
PMULHSW_mmx_reg	MMA[P0]	FastPath Single	4
PMULHSW_reg	MMA[P0]	FastPath Single	4
PMULHUW_mem	MMA[P0]	FastPath Single	8
PMULHUW_mmx_mem	MMA[P0]	FastPath Single	8
PMULHUW_mmx_reg	MMA[P0]	FastPath Single	4
PMULHUW_reg	MMA[P0]	FastPath Single	4
PMULHW_mem	MMA[P0]	FastPath Single	8
PMULHW_mmx_mem	MMA[P0]	FastPath Single	8
PMULHW_mmx_reg	MMA[P0]	FastPath Single	4
PMULHW_reg	MMA[P0]	FastPath Single	4
PMULLD_mem	MMA[P0]	FastPath Single	9
PMULLD_reg	MMA[P0]	FastPath Single	5
PMULLW_mem	MMA[P0]	FastPath Single	8
PMULLW_mmx_mem	MMA[P0]	FastPath Single	8
PMULLW_mmx_reg	MMA[P0]	FastPath Single	4
PMULLW_reg	MMA[P0]	FastPath Single	4
PMULUDQ_mem	MMA[P0]	FastPath Single	8
PMULUDQ_mmx_mem	MMA[P0]	FastPath Single	8
PMULUDQ_mmx_reg	MMA[P0]	FastPath Single	4
PMULUDQ_reg	MMA[P0]	FastPath Single	4
POR_mem	MAL[P2 P3]	FastPath Single	6
POR_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
POR<mmx>_mem	MAL[P2 P3]	FastPath Single	6
POR<mmx>_reg	MAL[P2 P3]	FastPath Single	2
PSADBW_mem	MAL[P2 P3]	FastPath Double	6
PSADBW_mmx_mem	MAL[P2 P3]	FastPath Double	6
PSADBW_mmx_reg	MAL[P2 P3]	FastPath Double	2
PSADBW_reg	MAL[P2 P3]	FastPath Double	2
PSHUFB_mem	XBR[P1]	FastPath Single	7
PSHUFB_mmx_mem	XBR[P1]	FastPath Single	7
PSHUFB_mmx_reg	XBR[P1]	FastPath Single	3
PSHUFB_reg	XBR[P1]	FastPath Single	3
PSHUFD_mem	XBR[P1]	FastPath Single	6
PSHUFD_reg	XBR[P1]	FastPath Single	2
PSHUFHW_mem	XBR[P1]	FastPath Single	6
PSHUFHW_reg	XBR[P1]	FastPath Single	2
PSHUFLW_mem	XBR[P1]	FastPath Single	6
PSHUFLW_reg	XBR[P1]	FastPath Single	2
PSHUFW_mmx_mem	XBR[P1]	FastPath Single	6
PSHUFW_mmx_reg	XBR[P1]	FastPath Single	2
PSIGNB_mem	MAL[P2 P3]	FastPath Single	6
PSIGNB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSIGNB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSIGNB_reg	MAL[P2 P3]	FastPath Single	2
PSIGND_mem	MAL[P2 P3]	FastPath Single	6
PSIGND_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSIGND_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSIGND_reg	MAL[P2 P3]	FastPath Single	2
PSIGNW_mem	MAL[P2 P3]	FastPath Single	6
PSIGNW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSIGNW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSIGNW_reg	MAL[P2 P3]	FastPath Single	2
PSLLD_mem	XBR[P1]	FastPath Single	7
PSLLD_mmx_mem	XBR[P1]	FastPath Single	7
PSLLD_mmx_reg	XBR[P1]	FastPath Single	3
PSLLD_reg	XBR[P1]	FastPath Single	3

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PSLLDQ_mem	XBR[P1]	FastPath Single	6
PSLLDQ_reg	XBR[P1]	FastPath Single	2
PSLLQ_mem	XBR[P1]	FastPath Single	7
PSLLQ_mmx_mem	XBR[P1]	FastPath Single	7
PSLLQ_mmx_reg	XBR[P1]	FastPath Single	3
PSLLQ_reg	XBR[P1]	FastPath Single	3
PSLLW_mem	XBR[P1]	FastPath Single	7
PSLLW_mmx_mem	XBR[P1]	FastPath Single	7
PSLLW_mmx_reg	XBR[P1]	FastPath Single	3
PSLLW_reg	XBR[P1]	FastPath Single	3
PSRAD_mem	XBR[P1]	FastPath Single	7
PSRAD_mmx_mem	XBR[P1]	FastPath Single	7
PSRAD_mmx_reg	XBR[P1]	FastPath Single	3
PSRAD_reg	XBR[P1]	FastPath Single	3
PSRAW_mem	XBR[P1]	FastPath Single	7
PSRAW_mmx_mem	XBR[P1]	FastPath Single	7
PSRAW_mmx_reg	XBR[P1]	FastPath Single	3
PSRAW_reg	XBR[P1]	FastPath Single	3
PSRLD_mem	XBR[P1]	FastPath Single	7
PSRLD_mmx_mem	XBR[P1]	FastPath Single	7
PSRLD_mmx_reg	XBR[P1]	FastPath Single	3
PSRLD_reg	XBR[P1]	FastPath Single	3
PSRLDQ_mem	XBR[P1]	FastPath Single	6
PSRLDQ_reg	XBR[P1]	FastPath Single	2
PSRLQ_mem	XBR[P1]	FastPath Single	7
PSRLQ_mmx_mem	XBR[P1]	FastPath Single	7
PSRLQ_mmx_reg	XBR[P1]	FastPath Single	3
PSRLQ_reg	XBR[P1]	FastPath Single	3
PSRLW_mem	XBR[P1]	FastPath Single	7
PSRLW_mmx_mem	XBR[P1]	FastPath Single	7
PSRLW_mmx_reg	XBR[P1]	FastPath Single	3
PSRLW_reg	XBR[P1]	FastPath Single	3
PSUBB_mem	MAL[P2 P3]	FastPath Single	6
PSUBB_mmx_mem	MAL[P2 P3]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PSUBB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBB_reg	MAL[P2 P3]	FastPath Single	2
PSUBD_mem	MAL[P2 P3]	FastPath Single	6
PSUBD_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBD_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBD_reg	MAL[P2 P3]	FastPath Single	2
PSUBQ_mem	MAL[P2 P3]	FastPath Single	6
PSUBQ_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBQ_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBQ_reg	MAL[P2 P3]	FastPath Single	2
PSUBSB_mem	MAL[P2 P3]	FastPath Single	6
PSUBSB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBSB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBSB_reg	MAL[P2 P3]	FastPath Single	2
PSUBSW_mem	MAL[P2 P3]	FastPath Single	6
PSUBSW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBSW_reg	MAL[P2 P3]	FastPath Single	2
PSUBUSB_mem	MAL[P2 P3]	FastPath Single	6
PSUBUSB_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBUSB_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBUSB_reg	MAL[P2 P3]	FastPath Single	2
PSUBUSW_mem	MAL[P2 P3]	FastPath Single	6
PSUBUSW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBUSW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBUSW_reg	MAL[P2 P3]	FastPath Single	2
PSUBW_mem	MAL[P2 P3]	FastPath Single	6
PSUBW_mmx_mem	MAL[P2 P3]	FastPath Single	6
PSUBW_mmx_reg	MAL[P2 P3]	FastPath Single	2
PSUBW_reg	MAL[P2 P3]	FastPath Single	2
PTEST_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
PTEST_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
PUNPCKHBW_mem	XBR[P1]	FastPath Single	6
PUNPCKHBW_mmx_mem	XBR[P1]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
PUNPCKHBW_mmx_reg	XBR[P1]	FastPath Single	2
PUNPCKHBW_reg	XBR[P1]	FastPath Single	2
PUNPCKHDQ_mem	XBR[P1]	FastPath Single	6
PUNPCKHDQ_mmx_mem	XBR[P1]	FastPath Single	6
PUNPCKHDQ_mmx_reg	XBR[P1]	FastPath Single	2
PUNPCKHDQ_reg	XBR[P1]	FastPath Single	2
PUNPCKHQDQ_mem	XBR[P1]	FastPath Single	6
PUNPCKHQDQ_reg	XBR[P1]	FastPath Single	2
PUNPCKHWD_mem	XBR[P1]	FastPath Single	6
PUNPCKHWD_mmx_mem	XBR[P1]	FastPath Single	6
PUNPCKHWD_mmx_reg	XBR[P1]	FastPath Single	2
PUNPCKHWD_reg	XBR[P1]	FastPath Single	2
PUNPCKLBW_mem	XBR[P1]	FastPath Single	6
PUNPCKLBW_mmx_mem	XBR[P1]	FastPath Single	6
PUNPCKLBW_mmx_reg	XBR[P1]	FastPath Single	2
PUNPCKLBW_reg	XBR[P1]	FastPath Single	2
PUNPCKLDQ_mem	XBR[P1]	FastPath Single	6
PUNPCKLDQ_mmx_mem	XBR[P1]	FastPath Single	6
PUNPCKLDQ_mmx_reg	XBR[P1]	FastPath Single	2
PUNPCKLDQ_reg	XBR[P1]	FastPath Single	2
PUNPCKLQDQ_mem	XBR[P1]	FastPath Single	6
PUNPCKLQDQ_reg	XBR[P1]	FastPath Single	2
PUNPCKLWD_mem	XBR[P1]	FastPath Single	6
PUNPCKLWD_mmx_mem	XBR[P1]	FastPath Single	6
PUNPCKLWD_mmx_reg	XBR[P1]	FastPath Single	2
PUNPCKLWD_reg	XBR[P1]	FastPath Single	2
PXOR_mem	MAL[P2 P3]	FastPath Single	6
PXOR_mmx_mem	MAL[P2 P3]	FastPath Single	6
PXOR_mmx_reg	MAL[P2 P3]	FastPath Single	2
PXOR_reg	MAL[P2 P3]	FastPath Single	2
RCPSS_mem	FMA[P0 P1]	FastPath Single	9
RCPSS_reg	FMA[P0 P1]	FastPath Single	5
RCPSS_mem	FMA[P0 P1]	FastPath Single	9
RCPSS_reg	FMA[P0 P1]	FastPath Single	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
ROUNDPD_mem	CVT[P0]	FastPath Single	8
ROUNDPD_reg	CVT[P0]	FastPath Single	4
ROUNDPS_mem	CVT[P0]	FastPath Single	8
ROUNDPS_reg	CVT[P0]	FastPath Single	4
ROUNDSD_mem	CVT[P0]	FastPath Single	8
ROUNDSD_reg	CVT[P0]	FastPath Single	4
ROUNDSS_mem	CVT[P0]	FastPath Single	8
ROUNDSS_reg	CVT[P0]	FastPath Single	4
RSQRTPS_mem	FMA[P0 P1]	FastPath Single	9
RSQRTPS_reg	FMA[P0 P1]	FastPath Single	5
RSQRTSS_mem	FMA[P0 P1]	FastPath Single	9
RSQRTSS_reg	FMA[P0 P1]	FastPath Single	5
SHUFPD_mem	XBR[P1]	FastPath Single	6
SHUFPD_reg	XBR[P1]	FastPath Single	2
SHUFPS_mem	XBR[P1]	FastPath Single	6
SHUFPS_reg	XBR[P1]	FastPath Single	2
SQRTPD_mem	FMA[P0 P1]	FastPath Single	42
SQRTPD_reg	FMA[P0 P1]	FastPath Single	38
SQRTPS_mem	FMA[P0 P1]	FastPath Single	33
SQRTPS_reg	FMA[P0 P1]	FastPath Single	29
SQRTSD_mem	FMA[P0 P1]	FastPath Single	42
SQRTSD_reg	FMA[P0 P1]	FastPath Single	38
SQRTSS_mem	FMA[P0 P1]	FastPath Single	33
SQRTSS_reg	FMA[P0 P1]	FastPath Single	29
STMXCSR_mem	microcode	microcode	NA
SUBPD_mem	FMA[P0 P1]	FastPath Single	10
SUBPD_reg	FMA[P0 P1]	FastPath Single	5
SUBPS_mem	FMA[P0 P1]	FastPath Single	10
SUBPS_reg	FMA[P0 P1]	FastPath Single	5
SUBSD_mem	FMA[P0 P1]	FastPath Single	10
SUBSD_reg	FMA[P0 P1]	FastPath Single	5
SUBSS_mem	FMA[P0 P1]	FastPath Single	10
SUBSS_reg	FMA[P0 P1]	FastPath Single	5
UCOMISD_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
UCOMISD_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
UCOMISS_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
UCOMISS_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
UNPCKHPD_mem	XBR[P1]	FastPath Single	6
UNPCKHPD_reg	XBR[P1]	FastPath Single	2
UNPCKHPS_mem	XBR[P1]	FastPath Single	6
UNPCKHPS_reg	XBR[P1]	FastPath Single	2
UNPCKLPD_mem	XBR[P1]	FastPath Single	6
UNPCKLPD_reg	XBR[P1]	FastPath Single	2
UNPCKLPS_mem	XBR[P1]	FastPath Single	6
UNPCKLPS_reg	XBR[P1]	FastPath Single	2
VADDPD_128_mem	FMA[P0 P1]	FastPath Single	10
VADDPD_128_reg	FMA[P0 P1]	FastPath Single	5
VADDPD_256_mem	FMA[P0 P1]	FastPath Double	11
VADDPD_256_reg	FMA[P0 P1]	FastPath Double	7
VADDPS_128_mem	FMA[P0 P1]	FastPath Single	10
VADDPS_128_reg	FMA[P0 P1]	FastPath Single	5
VADDPS_256_mem	FMA[P0 P1]	FastPath Double	11
VADDPS_256_reg	FMA[P0 P1]	FastPath Double	7
VADDSD_128_mem	FMA[P0 P1]	FastPath Single	10
VADDSD_128_reg	FMA[P0 P1]	FastPath Single	5
VADDSS_128_mem	FMA[P0 P1]	FastPath Single	10
VADDSS_128_reg	FMA[P0 P1]	FastPath Single	5
VADDSUBPD_128_mem	FMA[P0 P1]	FastPath Single	10
VADDSUBPD_128_reg	FMA[P0 P1]	FastPath Single	5
VADDSUBPD_256_mem	FMA[P0 P1]	FastPath Double	11
VADDSUBPD_256_reg	FMA[P0 P1]	FastPath Double	7
VADDSUBPS_128_mem	FMA[P0 P1]	FastPath Single	10
VADDSUBPS_128_reg	FMA[P0 P1]	FastPath Single	5
VADDSUBPS_256_mem	FMA[P0 P1]	FastPath Double	11
VADDSUBPS_256_reg	FMA[P0 P1]	FastPath Double	7
VAESDEC_128_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
VAESDEC_128_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
VAESDECLAST_128_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VAESDECLAST_128_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
VAEENC_128_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
VAEENC_128_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
VAEENCLAST_128_mem	XBR[P1]/FMA[P0 P1]	FastPath Double	6/5
VAEENCLAST_128_reg	XBR[P1]/FMA[P0 P1]	FastPath Double	2/5
VAESIMC_128_mem	FMA[P0 P1]	FastPath Single	9
VAESIMC_128_reg	FMA[P0 P1]	FastPath Single	5
VAESKEYGENASSIST_128_mem	FMA[P0 P1]	FastPath Single	9
VAESKEYGENASSIST_128_reg	FMA[P0 P1]	FastPath Single	5
VANDNPD_128_mem	MAL[P2 P3]	FastPath Single	6
VANDNPD_128_reg	MAL[P2 P3]	FastPath Single	2
VANDNPD_256_mem	MAL[P2 P3]	FastPath Double	7
VANDNPD_256_reg	MAL[P2 P3]	FastPath Double	3
VANDNPS_128_mem	MAL[P2 P3]	FastPath Single	6
VANDNPS_128_reg	MAL[P2 P3]	FastPath Single	2
VANDNPS_256_mem	MAL[P2 P3]	FastPath Double	7
VANDNPS_256_reg	MAL[P2 P3]	FastPath Double	3
VANDPD_128_mem	MAL[P2 P3]	FastPath Single	6
VANDPD_128_reg	MAL[P2 P3]	FastPath Single	2
VANDPD_256_mem	MAL[P2 P3]	FastPath Double	7
VANDPD_256_reg	MAL[P2 P3]	FastPath Double	3
VANDPS_128_mem	MAL[P2 P3]	FastPath Single	6
VANDPS_128_reg	MAL[P2 P3]	FastPath Single	2
VANDPS_256_mem	MAL[P2 P3]	FastPath Double	7
VANDPS_256_reg	MAL[P2 P3]	FastPath Double	3
VBLENDPD_128_mem	MAL[P2 P3]	FastPath Single	6
VBLENDPD_128_reg	MAL[P2 P3]	FastPath Single	2
VBLENDPD_256_mem	MAL[P2 P3]	FastPath Double	7
VBLENDPD_256_reg	MAL[P2 P3]	FastPath Double	3
VBLENDPS_128_mem	MAL[P2 P3]	FastPath Single	6
VBLENDPS_128_reg	MAL[P2 P3]	FastPath Single	2
VBLENDPS_256_mem	MAL[P2 P3]	FastPath Double	7
VBLENDPS_256_reg	MAL[P2 P3]	FastPath Double	3
VBLENDVPD_128_mem	XBR[P1]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VBLENDVPD_128_reg	XBR[P1]	FastPath Single	2
VBLENDVPD_256_mem	XBR[P1]	FastPath Single	7
VBLENDVPD_256_reg	XBR[P1]	FastPath Single	3
VBLENDVPS_128_mem	XBR[P1]	FastPath Single	6
VBLENDVPS_128_reg	XBR[P1]	FastPath Single	2
VBLENDVPS_256_mem	XBR[P1]	FastPath Double	7
VBLENDVPS_256_reg	XBR[P1]	FastPath Double	3
VBROADCASTF128_256_mem	MAL[P2 P3]	FastPath Double	6
VBROADCASTSD_256_mem	MAL[P2 P3]	FastPath Double	6
VBROADCASTSS_128_mem	MAL[P2 P3]	FastPath Single	2
VBROADCASTSS_256_mem	MAL[P2 P3]	FastPath Double	6
VCMPPD_128_mem	FMA[P0 P1]	FastPath Single	6
VCMPPD_128_reg	FMA[P0 P1]	FastPath Single	2
VCMPPD_256_mem	FMA[P0 P1]	FastPath Single	5
VCMPPD_256_reg	FMA[P0 P1]	FastPath Single	2
VCMPPS_128_mem	FMA[P0 P1]	FastPath Single	5
VCMPPS_128_reg	FMA[P0 P1]	FastPath Single	2
VCMPPS_256_mem	FMA[P0 P1]	FastPath Single	5
VCMPPS_256_reg	FMA[P0 P1]	FastPath Single	2
VCMPSD_128_mem	FMA[P0 P1]	FastPath Single	5
VCMPSD_128_reg	FMA[P0 P1]	FastPath Single	2
VCMPS_128_mem	FMA[P0 P1]	FastPath Single	5
VCMPS_128_reg	FMA[P0 P1]	FastPath Single	2
VCOMISD_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
VCOMISD_128_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
VCOMISS_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
VCOMISS_128_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
VCVTDQ2PD_128_mem	XBR[P1]/CVT[P0]	FastPath Double	6/4
VCVTDQ2PD_128_reg	XBR[P1]/CVT[P0]	FastPath Double	2/4
VCVTDQ2PD_256_mem	microcode	microcode	11
VCVTDQ2PD_256_reg	microcode	microcode	7
VCVTDQ2PS_128_mem	CVT[P0]	FastPath Single	8
VCVTDQ2PS_128_reg	CVT[P0]	FastPath Single	4
VCVTDQ2PS_256_mem	microcode	microcode	7

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VCVTDQ2PS_256_reg	microcode	microcode	3
VCVTPD2DQ_128_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
VCVTPD2DQ_128_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
VCVTPD2DQ_256_mem	microcode	microcode	12
VCVTPD2DQ_256_reg	microcode	microcode	8
VCVTPD2PS_128_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
VCVTPD2PS_128_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
VCVTPD2PS_256_mem	microcode	microcode	11
VCVTPD2PS_256_reg	microcode	microcode	7
VCVTPS2DQ_128_mem	CVT[P0]	FastPath Single	8
VCVTPS2DQ_128_reg	CVT[P0]	FastPath Single	4
VCVTPS2DQ_256_mem	microcode	microcode	9
VCVTPS2DQ_256_reg	microcode	microcode	5
VCVTPS2PD_128_mem	XBR[P1]/CVT[P0]	FastPath Double	6/4
VCVTPS2PD_128_reg	XBR[P1]/CVT[P0]	FastPath Double	2/4
VCVTPS2PD_256_mem	microcode	microcode	11
VCVTPS2PD_256_reg	microcode	microcode	7
VCVTSD2SI_128_OP32_mem	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTSD2SI_128_OP32_reg	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTSD2SI_128_OP64_mem	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTSD2SI_128_OP64_reg	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTSD2SS_128_mem	CVT[P0]	FastPath Single	8
VCVTSD2SS_128_reg	CVT[P0]	FastPath Single	4
VCVTSI2SD_128_OP32_mem	CVT[P0]	FastPath Single	8
VCVTSI2SD_128_OP32_reg	CVT[P0]	FastPath Double	4
VCVTSI2SD_128_OP64_mem	CVT[P0]	FastPath Single	8
VCVTSI2SD_128_OP64_reg	CVT[P0]	FastPath Double	4
VCVTSI2SS_128_OP32_mem	CVT[P0]	FastPath Single	8
VCVTSI2SS_128_OP32_reg	CVT[P0]	FastPath Double	4
VCVTSI2SS_128_OP64_mem	CVT[P0]	FastPath Single	8
VCVTSI2SS_128_OP64_reg	CVT[P0]	FastPath Double	4
VCVTSS2SD_128_mem	CVT[P0]	FastPath Single	8
VCVTSS2SD_128_reg	CVT[P0]	FastPath Single	4
VCVTSS2SI_128_reg32_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VCVTSS2SI_128_reg32_xmm	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTSS2SI_128_reg64_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTSS2SI_128_reg64_xmm	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTTPD2DQ_128_mem	CVT[P0]/XBR[P1]	FastPath Double	8/2
VCVTTPD2DQ_128_reg	CVT[P0]/XBR[P1]	FastPath Double	4/2
VCVTTPD2DQ_256_mem	microcode	microcode	9
VCVTTPD2DQ_256_reg	microcode	microcode	9
VCVTTPS2DQ_128_mem	CVT[P0]	FastPath Single	8
VCVTTPS2DQ_128_reg	CVT[P0]	FastPath Single	4
VCVTTPS2DQ_256_mem	microcode	microcode	9
VCVTTPS2DQ_256_reg	microcode	microcode	5
VCVTTSD2SI_128_reg32_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTTSD2SI_128_reg32_xmm	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTTSD2SI_128_reg64_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTTSD2SI_128_reg64_xmm	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTTSS2SI_128_reg32_mem32	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTTSS2SI_128_reg32_xmm	CVT[P0]/STO[P3]	FastPath Double	4/2
VCVTTSS2SI_128_reg64_mem64	CVT[P0]/STO[P3]	FastPath Double	8/2
VCVTTSS2SI_128_reg64_xmm	CVT[P0]/STO[P3]	FastPath Double	4/2
VDIVPD_128_mem	FMA[P0 P1]	FastPath Single	31
VDIVPD_128_reg	FMA[P0 P1]	FastPath Single	27
VDIVPD_256_mem	FMA[P0 P1]	FastPath Double	31
VDIVPD_256_reg	FMA[P0 P1]	FastPath Double	27
VDIVPS_128_mem	FMA[P0 P1]	FastPath Single	28
VDIVPS_128_reg	FMA[P0 P1]	FastPath Single	24
VDIVPS_256_mem	FMA[P0 P1]	FastPath Double	28
VDIVPS_256_reg	FMA[P0 P1]	FastPath Double	24
VDIVSD_128_mem	FMA[P0 P1]	FastPath Single	31
VDIVSD_128_reg	FMA[P0 P1]	FastPath Single	27
VDIVSS_128_mem	FMA[P0 P1]	FastPath Single	28
VDIVSS_128_reg	FMA[P0 P1]	FastPath Single	24
VDPPD_128_mem	microcode	microcode	19
VDPPD_128_reg	microcode	microcode	15
VDPPS_128_mem	microcode	microcode	29

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VDPPS_128_reg	microcode	microcode	25
VDPPS_256_mem	microcode	microcode	31
VDPPS_256_reg	microcode	microcode	27
VEEXTRACTF128_256_mem	MAL[P2 P3]/STO[P3]	FastPath Double	2/4
VEEXTRACTF128_256_reg	MAL[P2 P3]	FastPath Single	2
VEEXTRACTPS_128_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
VEEXTRACTPS_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
VFMADD132PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD132PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD132PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD132PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD132PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD132PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD132PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD132PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD132SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD132SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD132SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD132SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD132SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD132SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD132SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD132SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD213PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD213PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD213PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD213PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD213PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD213PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD213PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD213PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD213SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD213SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD213SD_256_mem	FMA[P0 P1]	FastPath Double	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFMADD213SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD213SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD213SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD213SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD213SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD231PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD231PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD231PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD231PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD231PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD231PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD231PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD231PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD231SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD231SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD231SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD231SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADD231SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADD231SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADD231SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADD231SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADDPD_256_reg	FMA[P0 P1]	FastPath Double	7
VFMADDPS_128_mem	FMA[P0 P1]	FastPath Single	10
VFMADDPS_128_reg	FMA[P0 P1]	FastPath Single	5
VFMADDPS_256_mem	FMA[P0 P1]	FastPath Double	11
VFMADDPS_256_reg	FMA[P0 P1]	FastPath Double	7
VFMADDSD_128_mem	FMA[P0 P1]	FastPath Single	10
VFMADDSD_128_reg	FMA[P0 P1]	FastPath Single	5
VFMADDSS_128_mem	FMA[P0 P1]	FastPath Single	10
VFMADDSS_128_reg	FMA[P0 P1]	FastPath Single	5
VFMADDSUB132PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB132PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB132PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB132PD_256_reg	FMA[P0 P1]	FastPath Double	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFMADDSUB132PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB132PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB132PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB132PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB213PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB213PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB213PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB213PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB213PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB213PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB213PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB213PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB231PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB231PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB231PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB231PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB231PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB231PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMADDSUB231PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMADDSUB231PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMADDSUBPD_128_mem	FMA[P0 P1]	FastPath Single	10
VFMADDSUBPD_128_reg	FMA[P0 P1]	FastPath Single	5
VFMADDSUBPD_256_mem	FMA[P0 P1]	FastPath Double	11
VFMADDSUBPD_256_reg	FMA[P0 P1]	FastPath Double	7
VFMADDSUBPS_128_mem	FMA[P0 P1]	FastPath Single	10
VFMADDSUBPS_128_reg	FMA[P0 P1]	FastPath Single	5
VFMADDSUBPS_256_mem	FMA[P0 P1]	FastPath Double	11
VFMADDSUBPS_256_reg	FMA[P0 P1]	FastPath Double	7
VFMSUB132PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB132PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB132PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB132PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB132PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB132PS_128_reg	FMA[P0 P1]	FastPath Single	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFMSUB132PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB132PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB132SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB132SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB132SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB132SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB132SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB132SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB132SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB132SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB213PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB213PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB213PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB213PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB213PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB213PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB213PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB213PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB213SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB213SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB213SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB213SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB213SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB213SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB213SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB213SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB231PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB231PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB231PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB231PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB231PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB231PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB231PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB231PS_256_reg	FMA[P0 P1]	FastPath Double	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFMSUB231SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB231SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB231SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB231SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUB231SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUB231SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUB231SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUB231SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD132PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD132PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD132PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD132PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD132PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD132PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD132PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD132PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD213PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD213PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD213PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD213PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD213PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD213PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD213PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD213PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD231PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD231PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD231PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD231PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD231PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD231PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFMSUBADD231PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFMSUBADD231PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFMSUBADDDPD_128_mem	FMA[P0 P1]	FastPath Single	10
VFMSUBADDDPD_128_reg	FMA[P0 P1]	FastPath Single	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFMSUBADDPD_256_mem	FMA[P0 P1]	FastPath Double	11
VFMSUBADDPD_256_reg	FMA[P0 P1]	FastPath Double	7
VFMSUBADDPDPS_128_mem	FMA[P0 P1]	FastPath Single	10
VFMSUBADDPDPS_128_reg	FMA[P0 P1]	FastPath Single	5
VFMSUBADDPDPS_256_mem	FMA[P0 P1]	FastPath Double	11
VFMSUBADDPDPS_256_reg	FMA[P0 P1]	FastPath Double	7
VFMSUBPD_128_mem	FMA[P0 P1]	FastPath Single	10
VFMSUBPD_128_reg	FMA[P0 P1]	FastPath Single	5
VFMSUBPD_256_mem	FMA[P0 P1]	FastPath Double	11
VFMSUBPD_256_reg	FMA[P0 P1]	FastPath Double	7
VFMSUBPS_128_mem	FMA[P0 P1]	FastPath Single	10
VFMSUBPS_128_reg	FMA[P0 P1]	FastPath Single	5
VFMSUBPS_256_mem	FMA[P0 P1]	FastPath Double	11
VFMSUBPS_256_reg	FMA[P0 P1]	FastPath Double	7
VFMSUBSD_128_mem	FMA[P0 P1]	FastPath Single	10
VFMSUBSD_128_reg	FMA[P0 P1]	FastPath Single	5
VFMSUBSS_128_mem	FMA[P0 P1]	FastPath Single	10
VFMSUBSS_128_reg	FMA[P0 P1]	FastPath Single	5
VFNMADD132PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD132PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD132PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD132PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD132PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD132PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD132PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD132PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD132SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD132SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD132SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD132SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD132SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD132SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD132SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD132SS_256_reg	FMA[P0 P1]	FastPath Double	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFNMADD213PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD213PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD213PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD213PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD213PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD213PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD213PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD213PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD213SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD213SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD213SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD213SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD213SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD213SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD213SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD213SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD231PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD231PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD231PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD231PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD231PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD231PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD231PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD231PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD231SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD231SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD231SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD231SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADD231SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMADD231SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMADD231SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMADD231SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMADDPD_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMADDPD_128_reg	FMA[P0 P1]	FastPath Single	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFNMADDPD_256_mem	FMA[P0 P1]	FastPath Double	11
VFNMADDPD_256_reg	FMA[P0 P1]	FastPath Double	7
VFNMADDPDPS_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMADDPDPS_128_reg	FMA[P0 P1]	FastPath Single	5
VFNMADDPDPS_256_mem	FMA[P0 P1]	FastPath Double	11
VFNMADDPDPS_256_reg	FMA[P0 P1]	FastPath Double	7
VFNMADDS_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMADDS_128_reg	FMA[P0 P1]	FastPath Single	5
VFNMADDS_256_mem	FMA[P0 P1]	FastPath Single	10
VFNMADDS_256_reg	FMA[P0 P1]	FastPath Single	6
VFNMSUB132PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB132SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB132SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213PS_256_reg	FMA[P0 P1]	FastPath Double	NA

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFNMSUB213SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB213SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB213SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231PD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231PD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231PD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231PD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231PS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231PS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231PS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231PS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231SD_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231SD_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231SD_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231SD_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231SS_128_mem	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231SS_128_reg	FMA[P0 P1]	FastPath Single	NA
VFNMSUB231SS_256_mem	FMA[P0 P1]	FastPath Double	NA
VFNMSUB231SS_256_reg	FMA[P0 P1]	FastPath Double	NA
VFNMSUBPD_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMSUBPD_128_reg	FMA[P0 P1]	FastPath Single	5
VFNMSUBPD_256_mem	FMA[P0 P1]	FastPath Double	11
VFNMSUBPD_256_reg	FMA[P0 P1]	FastPath Double	7
VFNMSUBPS_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMSUBPS_128_reg	FMA[P0 P1]	FastPath Single	5
VFNMSUBPS_256_mem	FMA[P0 P1]	FastPath Double	11
VFNMSUBPS_256_reg	FMA[P0 P1]	FastPath Double	7
VFNMSUBSD_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMSUBSD_128_reg	FMA[P0 P1]	FastPath Single	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VFNMSUBSS_128_mem	FMA[P0 P1]	FastPath Single	10
VFNMSUBSS_128_reg	FMA[P0 P1]	FastPath Single	5
VFRCZPD_128_mem	microcode	microcode	NA
VFRCZPD_128_reg	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
VFRCZPD_256_mem	microcode	microcode	NA
VFRCZPD_256_reg	microcode	microcode	NA
VFRCZPS_128_mem	microcode	microcode	NA
VFRCZPS_128_reg	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
VFRCZPS_256_mem	microcode	microcode	NA
VFRCZPS_256_reg	microcode	microcode	NA
VFRCZSD_128_mem	microcode	microcode	NA
VFRCZSD_128_reg	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
VFRCZSS_128_mem	microcode	microcode	NA
VFRCZSS_128_reg	CVT[P0]/FMA[P0 P1]	FastPath Double	4/5
VHADDPD_128_mem	microcode	microcode	15
VHADDPD_128_reg	microcode	microcode	11
VHADDPD_256_mem	microcode	microcode	15
VHADDPD_256_reg	microcode	microcode	11
VHADDPS_128_mem	microcode	microcode	15
VHADDPS_128_reg	microcode	microcode	11
VHADDPS_256_mem	microcode	microcode	15
VHADDPS_256_reg	microcode	microcode	11
VHSUBPD_128_mem	microcode	microcode	15
VHSUBPD_128_reg	microcode	microcode	11
VHSUBPD_256_mem	microcode	microcode	15
VHSUBPD_256_reg	microcode	microcode	11
VHSUBPS_128_mem	microcode	microcode	15
VHSUBPS_128_reg	microcode	microcode	11
VHSUBPS_256_mem	microcode	microcode	15
VHSUBPS_256_reg	microcode	microcode	11
VINSERTF128_256_mem	MAL[P2 P3]	FastPath Double	7
VINSERTF128_256_reg	MAL[P2 P3]	FastPath Double	3
VINSERTPS_128_mem	XBR[P1]	FastPath Single	6
VINSERTPS_128_reg	XBR[P1]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VLDDQU_128_mem	None	FastPath Single	4
VLDDQU_256_mem	None	FastPath Double	5
VLDMXCSR_128_mem	microcode	microcode	NA
VMASKMOVDQU_128_reg	microcode	microcode	NA
VMASKMOVDPD_128_mem	FMA[P0 P1]	FastPath Single	2
VMASKMOVDPD_256_mem	FMA[P0 P1]	FastPath Double	3
VMASKMOVPS_128_mem	FMA[P0 P1]	FastPath Single	2
VMASKMOVPS_256_mem	FMA[P0 P1]	FastPath Double	3
VMAXPD_128_mem	FMA[P0 P1]	FastPath Single	5
VMAXPD_128_reg	FMA[P0 P1]	FastPath Single	2
VMAXPD_256_mem	FMA[P0 P1]	FastPath Double	7
VMAXPD_256_reg	FMA[P0 P1]	FastPath Double	3
VMAXPS_128_mem	FMA[P0 P1]	FastPath Single	5
VMAXPS_128_reg	FMA[P0 P1]	FastPath Single	2
VMAXPS_256_mem	FMA[P0 P1]	FastPath Double	7
VMAXPS_256_reg	FMA[P0 P1]	FastPath Double	3
VMAXSD_128_mem	FMA[P0 P1]	FastPath Single	5
VMAXSD_128_reg	FMA[P0 P1]	FastPath Single	2
VMAXSS_128_mem	FMA[P0 P1]	FastPath Single	5
VMAXSS_128_reg	FMA[P0 P1]	FastPath Single	2
VMINPD_128_mem	FMA[P0 P1]	FastPath Single	5
VMINPD_128_reg	FMA[P0 P1]	FastPath Single	2
VMINPD_256_mem	FMA[P0 P1]	FastPath Double	7
VMINPD_256_reg	FMA[P0 P1]	FastPath Double	3
VMINPS_128_mem	FMA[P0 P1]	FastPath Single	5
VMINPS_128_reg	FMA[P0 P1]	FastPath Single	2
VMINPS_256_mem	FMA[P0 P1]	FastPath Double	7
VMINPS_256_reg	FMA[P0 P1]	FastPath Double	3
VMINSD_128_mem	FMA[P0 P1]	FastPath Single	5
VMINSD_128_reg	FMA[P0 P1]	FastPath Single	2
VMINSS_128_mem	FMA[P0 P1]	FastPath Single	5
VMINSS_128_reg	FMA[P0 P1]	FastPath Single	2
VMOVAPD_128_mem	MAL[P2 P3]	FastPath Single	4
VMOVAPD_128_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VMOVAPD_256_mem	MAL[P2 P3]	FastPath Double	5
VMOVAPD_256_reg	MAL[P2 P3]	FastPath Double	3
VMOVAPS_128_mem	MAL[P2 P3]	FastPath Single	4
VMOVAPS_128_reg	MAL[P2 P3]	FastPath Single	2
VMOVAPS_256_mem	MAL[P2 P3]	FastPath Double	5
VMOVAPS_256_reg	MAL[P2 P3]	FastPath Double	3
VMOVD_128_mem32_xmm	STO[P3]	FastPath Single	4
VMOVD_128_reg32_xmm	STO[P3]	FastPath Single	2
VMOVD_128_xmm_mem32	None	FastPath Double	8
VMOVD_128_xmm_reg32	None	FastPath Double	8
VMOVDDUP_128_mem	XBR[P1]	FastPath Single	4
VMOVDDUP_128_reg	XBR[P1]	FastPath Single	2
VMOVDDUP_256_mem	XBR[P1]	FastPath Double	5
VMOVDDUP_256_reg	XBR[P1]	FastPath Double	3
VMOVDQA_128_mem	MAL[P2 P3]	FastPath Single	4
VMOVDQA_128_reg	MAL[P2 P3]	FastPath Single	2
VMOVDQA_256_mem	MAL[P2 P3]	FastPath Double	5
VMOVDQA_256_reg	MAL[P2 P3]	FastPath Double	3
VMOVDQU_128_mem	MAL[P2 P3]	FastPath Single	4
VMOVDQU_128_reg	MAL[P2 P3]	FastPath Single	2
VMOVDQU_256_mem	MAL[P2 P3]	FastPath Double	5
VMOVDQU_256_reg	MAL[P2 P3]	FastPath Double	3
VMOVHLPD_128_reg	XBR[P1]	FastPath Single	2
VMOVHPD_128_mem_reg	XBR[P1]	FastPath Double	4
VMOVHPD_128_reg_mem	XBR[P1]	FastPath Single	4
VMOVHPS_128_mem_reg	XBR[P1]	FastPath Double	4
VMOVHPS_128_reg_mem	XBR[P1]	FastPath Single	4
VMOVLHPS_128_reg	XBR[P1]	FastPath Single	2
VMOVLPD_128_mem	FMA[P0 P1]	FastPath Single	4
VMOVLPS_128_mem	FMA[P0 P1]	FastPath Single	4
VMOVMSKPD_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/4
VMOVMSKPD_256_reg	XBR[P1]/STO[P3]	FastPath Double	2/4
VMOVMSKPS_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/4
VMOVMSKPS_256_reg	XBR[P1]/STO[P3]	FastPath Double	2/4

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VMOVNTDQ_128_mem	STO[P3]	FastPath Single	4
VMOVNTDQA_128_mem	STO[P3]	FastPath Single	4
VMOVNTPD_128_mem	STO[P3]	FastPath Single	4
VMOVNTPS_128_mem	STO[P3]	FastPath Single	4
VMOVQ_128_mem_xmm	MAL[P2 P3]	FastPath Single	4
VMOVQ_128_mem_xmm	MAL[P2 P3]	FastPath Single	4
VMOVQ_128_reg64_xmm	MAL[P2 P3]	FastPath Single	2
VMOVQ_128_xmm_reg64	MAL[P2 P3]	FastPath Single	2
VMOVSD_128_mem	FMA[P0 P1]	FastPath Single	4
VMOVSD_128_reg	FMA[P0 P1]	FastPath Single	2
VMOVSHDUP_128_mem	None	FastPath Single	4
VMOVSHDUP_128_reg	XBR[P1]	FastPath Single	2
VMOVSHDUP_256_mem	XBR[P1]	FastPath Double	5
VMOVSHDUP_256_reg	XBR[P1]	FastPath Double	3
VMOVSLDUP_128_mem	None	FastPath Single	4
VMOVSLDUP_128_reg	XBR[P1]	FastPath Single	2
VMOVSLDUP_256_mem	XBR[P1]	FastPath Double	5
VMOVSLDUP_256_reg	XBR[P1]	FastPath Double	3
VMOVSS_128_mem	None	FastPath Single	4
VMOVSS_128_reg	FMA[P0 P1]	FastPath Single	2
VMOVUPD_128_mem	None	FastPath Single	4
VMOVUPD_128_reg	MAL[P2 P3]	FastPath Single	2
VMOVUPD_256_mem	None	FastPath Double	5
VMOVUPD_256_reg	MAL[P2 P3]	FastPath Double	3
VMOVUPS_128_mem	None	FastPath Single	4
VMOVUPS_128_reg	MAL[P2 P3]	FastPath Single	2
VMOVUPS_256_mem	None	FastPath Double	5
VMOVUPS_256_reg	MAL[P2 P3]	FastPath Double	3
VMPSADBW_128_mem	microcode	microcode	8
VMPSADBW_128_reg	microcode	microcode	4
VMULPD_128_mem	FMA[P0 P1]	FastPath Single	10
VMULPD_128_reg	FMA[P0 P1]	FastPath Single	5
VMULPD_256_mem	FMA[P0 P1]	FastPath Double	11
VMULPD_256_reg	FMA[P0 P1]	FastPath Double	7

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VMULPS_128_mem	FMA[P0 P1]	FastPath Single	10
VMULPS_128_reg	FMA[P0 P1]	FastPath Single	5
VMULPS_256_mem	FMA[P0 P1]	FastPath Double	11
VMULPS_256_reg	FMA[P0 P1]	FastPath Double	7
VMULSD_128_mem	FMA[P0 P1]	FastPath Single	10
VMULSD_128_reg	FMA[P0 P1]	FastPath Single	5
VMULSS_128_mem	FMA[P0 P1]	FastPath Single	10
VMULSS_128_reg	FMA[P0 P1]	FastPath Single	5
VORPD_128_mem	MAL[P2 P3]	FastPath Single	6
VORPD_128_reg	MAL[P2 P3]	FastPath Single	2
VORPD_256_mem	MAL[P2 P3]	FastPath Double	7
VORPD_256_reg	MAL[P2 P3]	FastPath Double	3
VORPS_128_mem	MAL[P2 P3]	FastPath Single	6
VORPS_128_reg	MAL[P2 P3]	FastPath Single	2
VORPS_256_mem	MAL[P2 P3]	FastPath Double	7
VORPS_256_reg	MAL[P2 P3]	FastPath Double	3
VPABSB_128_mem	MAL[P2 P3]	FastPath Single	6
VPABSB_128_reg	MAL[P2 P3]	FastPath Single	2
VPABSD_128_mem	MAL[P2 P3]	FastPath Single	6
VPABSD_128_reg	MAL[P2 P3]	FastPath Single	2
VPABSW_128_mem	MAL[P2 P3]	FastPath Single	6
VPABSW_128_reg	MAL[P2 P3]	FastPath Single	2
VPACKSSDW_128_mem	XBR[P1]	FastPath Single	6
VPACKSSDW_128_reg	XBR[P1]	FastPath Single	2
VPACKSSWB_128_mem	XBR[P1]	FastPath Single	6
VPACKSSWB_128_reg	XBR[P1]	FastPath Single	2
VPACKUSDW_128_mem	XBR[P1]	FastPath Single	6
VPACKUSDW_128_reg	XBR[P1]	FastPath Single	2
VPACKUSWB_128_mem	XBR[P1]	FastPath Single	6
VPACKUSWB_128_reg	XBR[P1]	FastPath Single	2
VPADDB_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDB_128_reg	MAL[P2 P3]	FastPath Single	2
VPADDD_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDD_128_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPADDQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPADDSB_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDSB_128_reg	MAL[P2 P3]	FastPath Single	2
VPADDSW_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDSW_128_reg	MAL[P2 P3]	FastPath Single	2
VPADDUSB_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDUSB_128_reg	MAL[P2 P3]	FastPath Single	2
VPADDUSW_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDUSW_128_reg	MAL[P2 P3]	FastPath Single	2
VPADDW_128_mem	MAL[P2 P3]	FastPath Single	6
VPADDW_128_reg	MAL[P2 P3]	FastPath Single	2
VPALIGNR_128_mem	XBR[P1]	FastPath Single	6
VPALIGNR_128_reg	XBR[P1]	FastPath Single	2
VPAND_128_mem	MAL[P2 P3]	FastPath Single	6
VPAND_128_reg	MAL[P2 P3]	FastPath Single	2
VPANDN_128_mem	MAL[P2 P3]	FastPath Single	6
VPANDN_128_reg	MAL[P2 P3]	FastPath Single	2
VPAVGB_128_mem	MAL[P2 P3]	FastPath Single	6
VPAVGB_128_reg	MAL[P2 P3]	FastPath Single	2
VPAVGW_128_mem	MAL[P2 P3]	FastPath Single	6
VPAVGW_128_reg	MAL[P2 P3]	FastPath Single	2
VPBLENDVB_128_mem	XBR[P1]	FastPath Single	6
VPBLENDVB_128_reg	XBR[P1]	FastPath Single	2
VPBLENDW_128_mem	MAL[P2 P3]	FastPath Single	6
VPBLENDW_128_reg	MAL[P2 P3]	FastPath Single	2
VPCLMULQDQ_128_mem	microcode	microcode	16
VPCLMULQDQ_128_reg	microcode	microcode	12
VPCMOV_128_mem	XBR[P1]	FastPath Single	6
VPCMOV_128_reg	XBR[P1]	FastPath Single	2
VPCMOV_256_mem	XBR[P1]	FastPath Double	6
VPCMOV_256_reg	XBR[P1]	FastPath Double	2
VPCMPEQB_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPEQB_128_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPCMPEQD_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPEQD_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPEQQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPEQQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPEQW_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPEQW_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPESTRI_128_mem	microcode	microcode	34
VPCMPESTRI_128_reg	microcode	microcode	30
VPCMPESTRM_128_mem	microcode	microcode	NA
VPCMPESTRM_128_reg	microcode	microcode	NA
VPCMPGTB_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPGTB_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPGTD_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPGTD_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPGTQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPGTQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPGTW_128_mem	MAL[P2 P3]	FastPath Single	6
VPCMPGTW_128_reg	MAL[P2 P3]	FastPath Single	2
VPCMPISTRI_128_mem	microcode	microcode	14
VPCMPISTRI_128_xmm	microcode	microcode	10
VPCMPISTRM_128_mem	microcode	microcode	NA
VPCMPISTRM_128_xmm	microcode	microcode	NA
VPCOMB_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMB_128_reg	MAL[P2 P3]	FastPath Single	2
VPCOMD_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMD_128_reg	MAL[P2 P3]	FastPath Single	2
VPCOMQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPCOMUB_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMUB_128_reg	MAL[P2 P3]	FastPath Single	2
VPCOMUD_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMUD_128_reg	MAL[P2 P3]	FastPath Single	2
VPCOMUQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMUQ_128_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPCOMUW_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMUW_128_reg	MAL[P2 P3]	FastPath Single	2
VPCOMW_128_mem	MAL[P2 P3]	FastPath Single	6
VPCOMW_128_reg	MAL[P2 P3]	FastPath Single	2
VPERMIL2PD_128_mem	XBR[P1]	FastPath Single	7
VPERMIL2PD_128_reg	XBR[P1]	FastPath Single	3
VPERMIL2PD_256_mem	XBR[P1]	FastPath Double	8
VPERMIL2PD_256_reg	XBR[P1]	FastPath Double	4
VPERMIL2PS_128_mem	XBR[P1]	FastPath Single	7
VPERMIL2PS_128_reg	XBR[P1]	FastPath Single	3
VPERMIL2PS_256_mem	XBR[P1]	FastPath Double	8
VPERMIL2PS_256_reg	XBR[P1]	FastPath Double	4
VPERMILPD_128_mem	XBR[P1]	FastPath Single	7
VPERMILPD_128_reg	XBR[P1]	FastPath Single	3
VPERMILPD_256_mem	XBR[P1]	FastPath Double	8
VPERMILPD_256_reg	XBR[P1]	FastPath Double	4
VPERMILPS_128_mem	XBR[P1]	FastPath Single	7
VPERMILPS_128_reg	XBR[P1]	FastPath Single	3
VPERMILPS_256_mem	XBR[P1]	FastPath Double	8
VPERMILPS_256_reg	XBR[P1]	FastPath Double	4
VPEXTRB_128_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
VPEXTRB_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
VPEXTRD_128_xmm_mem64	XBR[P1]/STO[P3]	FastPath Double	6/2
VPEXTRD_128_xmm_reg64	XBR[P1]/STO[P3]	FastPath Double	2/2
VPEXTRQ_128_xmm_mem64	XBR[P1]/STO[P3]	FastPath Double	6/2
VPEXTRQ_128_xmm_reg64	XBR[P1]/STO[P3]	FastPath Double	2/2
VPEXTRW_128_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
VPEXTRW_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
VPHADDBD_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDBD_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDBQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDBQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDBW_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDBW_128_reg	MAL[P2 P3]	FastPath Single	2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPHADDD_128_mem	microcode	microcode	NA
VPHADDD_128_reg	microcode	microcode	NA
VPHADDDQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDDQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDSW_128_mem	microcode	microcode	9
VPHADDSW_128_reg	microcode	microcode	5
VPHADDUBD_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDUBD_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDUBQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDUBQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDUBW_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDUBW_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDUDQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDUDQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDUWD_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDUWD_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDUWQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDUWQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDW_128_mem	microcode	microcode	9
VPHADDW_128_reg	microcode	microcode	5
VPHADDWD_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDWD_128_reg	MAL[P2 P3]	FastPath Single	2
VPHADDWQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHADDWQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHMINPOSUW_128_mem	MAL[P2 P3]/XBR[P1]	FastPath Double	6/2
VPHMINPOSUW_128_reg	MAL[P2 P3]/XBR[P1]	FastPath Double	2/2
VPHSUBBW_128_mem	MAL[P2 P3]	FastPath Single	6
VPHSUBBW_128_reg	MAL[P2 P3]	FastPath Single	2
VPHSUBD_128_mem	microcode	microcode	9
VPHSUBD_128_reg	microcode	microcode	5
VPHSUBDQ_128_mem	MAL[P2 P3]	FastPath Single	6
VPHSUBDQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPHSUBSW_128_mem	microcode	microcode	9
VPHSUBSW_128_reg	microcode	microcode	5

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPHSUBW_128_mem	microcode	microcode	9
VPHSUBW_128_reg	microcode	microcode	5
VPHSUBWD_128_mem	MAL[P2 P3]	FastPath Single	6
VPHSUBWD_128_reg	MAL[P2 P3]	FastPath Single	2
VPINSRB_128_mem	XBR[P1]	FastPath Single	6
VPINSRB_128_reg	XBR[P1]	FastPath Single	2
VPINSRD_128_mem32_xmm	XBR[P1]	FastPath Single	6
VPINSRD_128_reg32_xmm	XBR[P1]	FastPath Single	2
VPINSRQ_128_mem64_xmm	XBR[P1]	FastPath Single	6
VPINSRQ_128_reg64_xmm	XBR[P1]	FastPath Single	2
VPINSRW_128_mem	XBR[P1]	FastPath Single	6
VPINSRW_128_reg	XBR[P1]	FastPath Double	2
VPMACSDD_128_mem	MMA[P0]	FastPath Single	9
VPMACSDD_128_reg	MMA[P0]	FastPath Single	5
VPMACSDQH_128_mem	MMA[P0]	FastPath Single	8
VPMACSDQH_128_reg	MMA[P0]	FastPath Single	4
VPMACSDQL_128_mem	MMA[P0]	FastPath Single	8
VPMACSDQL_128_reg	MMA[P0]	FastPath Single	4
VPMACSSDD_128_mem	MMA[P0]	FastPath Single	9
VPMACSSDD_128_reg	MMA[P0]	FastPath Single	5
VPMACSSDQH_128_mem	MMA[P0]	FastPath Single	8
VPMACSSDQH_128_reg	MMA[P0]	FastPath Single	4
VPMACSSDQL_128_mem	MMA[P0]	FastPath Single	8
VPMACSSDQL_128_reg	MMA[P0]	FastPath Single	4
VPMACSSWD_128_mem	MMA[P0]	FastPath Single	8
VPMACSSWD_128_reg	MMA[P0]	FastPath Single	4
VPMACSSWW_128_mem	MMA[P0]	FastPath Single	8
VPMACSSWW_128_reg	MMA[P0]	FastPath Single	4
VPMACSWD_128_mem	MMA[P0]	FastPath Single	8
VPMACSWD_128_reg	MMA[P0]	FastPath Single	4
VPMACSWW_128_mem	MMA[P0]	FastPath Single	8
VPMACSWW_128_reg	MMA[P0]	FastPath Single	4
VPMADCSSWD_128_mem	MMA[P0]	FastPath Single	8
VPMADCSSWD_128_reg	MMA[P0]	FastPath Single	4

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPMADCSWD_128_mem	MMA[P0]	FastPath Single	8
VPMADCSWD_128_reg	MMA[P0]	FastPath Single	4
VPMADDUBSW_128_mem	MMA[P0]	FastPath Single	8
VPMADDUBSW_128_reg	MMA[P0]	FastPath Single	4
VPMADDWD_128_mem	MMA[P0]	FastPath Single	8
VPMADDWD_128_reg	MMA[P0]	FastPath Single	4
VPMASB_128_mem	MAL[P2 P3]	FastPath Single	6
VPMASB_128_reg	MAL[P2 P3]	FastPath Single	2
VPMASD_128_mem	MAL[P2 P3]	FastPath Single	6
VPMASD_128_reg	MAL[P2 P3]	FastPath Single	2
VPMASW_128_mem	MAL[P2 P3]	FastPath Single	6
VPMASW_128_reg	MAL[P2 P3]	FastPath Single	2
VPMASUB_128_mem	MAL[P2 P3]	FastPath Single	6
VPMASUB_128_reg	MAL[P2 P3]	FastPath Single	2
VPMASUD_128_mem	MAL[P2 P3]	FastPath Single	6
VPMASUD_128_reg	MAL[P2 P3]	FastPath Single	2
VPMASUW_128_mem	MAL[P2 P3]	FastPath Single	6
VPMASUW_128_reg	MAL[P2 P3]	FastPath Single	2
VPMINSB_128_mem	MAL[P2 P3]	FastPath Single	6
VPMINSB_128_reg	MAL[P2 P3]	FastPath Single	2
VPMINSD_128_mem	MAL[P2 P3]	FastPath Single	6
VPMINSD_128_reg	MAL[P2 P3]	FastPath Single	2
VPMINSW_128_mem	MAL[P2 P3]	FastPath Single	6
VPMINSW_128_reg	MAL[P2 P3]	FastPath Single	2
VPMINUB_128_mem	MAL[P2 P3]	FastPath Single	6
VPMINUB_128_reg	MAL[P2 P3]	FastPath Single	2
VPMINUD_128_mem	MAL[P2 P3]	FastPath Single	6
VPMINUD_128_reg	MAL[P2 P3]	FastPath Single	2
VPMINUW_128_mem	MAL[P2 P3]	FastPath Single	6
VPMINUW_128_reg	MAL[P2 P3]	FastPath Single	2
VPMOVMASKB_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
VPMOVSB_128_mem	XBR[P1]	FastPath Single	4
VPMOVSB_128_reg	XBR[P1]	FastPath Single	2
VPMOVSBQ_128_mem	XBR[P1]	FastPath Single	4

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPMOVSBQ_128_reg	XBR[P1]	FastPath Single	2
VPMOVSBW_128_mem	XBR[P1]	FastPath Single	4
VPMOVSBW_128_reg	XBR[P1]	FastPath Single	2
VPMOVSDQ_128_mem	XBR[P1]	FastPath Single	4
VPMOVSDQ_128_reg	XBR[P1]	FastPath Single	2
VPMOVXWD_128_mem	XBR[P1]	FastPath Single	4
VPMOVXWD_128_reg	XBR[P1]	FastPath Single	2
VPMOVXWQ_128_mem	XBR[P1]	FastPath Single	4
VPMOVXWQ_128_reg	XBR[P1]	FastPath Single	2
VPMOVZBD_128_mem	XBR[P1]	FastPath Single	4
VPMOVZBD_128_reg	XBR[P1]	FastPath Single	2
VPMOVZBQ_128_mem	XBR[P1]	FastPath Single	4
VPMOVZBQ_128_reg	XBR[P1]	FastPath Single	2
VPMOVZBW_128_mem	XBR[P1]	FastPath Single	4
VPMOVZBW_128_reg	XBR[P1]	FastPath Single	2
VPMOVZDQ_128_mem	XBR[P1]	FastPath Single	4
VPMOVZDQ_128_reg	XBR[P1]	FastPath Single	2
VPMOVXWD_128_mem	XBR[P1]	FastPath Single	4
VPMOVXWD_128_reg	XBR[P1]	FastPath Single	2
VPMOVXWQ_128_mem	XBR[P1]	FastPath Single	4
VPMOVXWQ_128_reg	XBR[P1]	FastPath Single	2
VPMULDQ_128_mem	MMA[P0]	FastPath Single	8
VPMULDQ_128_reg	MMA[P0]	FastPath Single	4
VPMULHRW_128_mem	MMA[P0]	FastPath Single	8
VPMULHRW_128_reg	MMA[P0]	FastPath Single	4
VPMULHUW_128_mem	MMA[P0]	FastPath Single	8
VPMULHUW_128_reg	MMA[P0]	FastPath Single	4
VPMULHW_128_mem	MMA[P0]	FastPath Single	8
VPMULHW_128_reg	MMA[P0]	FastPath Single	4
VPMULLD_128_mem	MMA[P0]	FastPath Single	9
VPMULLD_128_reg	MMA[P0]	FastPath Single	5
VPMULLW_128_mem	MMA[P0]	FastPath Single	8
VPMULLW_128_reg	MMA[P0]	FastPath Single	4
VPMULUDQ_128_mem	MMA[P0]	FastPath Single	8

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPMULUDQ_128_reg	MMA[P0]	FastPath Single	4
VPOR_128_mem	MAL[P2 P3]	FastPath Single	6
VPOR_128_reg	MAL[P2 P3]	FastPath Single	2
VPPERM_128_mem	XBR[P1]	FastPath Single	6
VPPERM_128_reg	XBR[P1]	FastPath Single	2
VPROTB_128_mem	XBR[P1]	FastPath Single	6
VPROTB_128_reg	XBR[P1]	FastPath Single	2
VPROTD_128_mem	XBR[P1]	FastPath Single	6
VPROTD_128_reg	XBR[P1]	FastPath Single	2
VPROTQ_128_mem	XBR[P1]	FastPath Single	6
VPROTQ_128_reg	XBR[P1]	FastPath Single	2
VPROTW_128_mem	XBR[P1]	FastPath Single	6
VPROTW_128_reg	XBR[P1]	FastPath Single	2
VPSADBW_128_mem	MAL[P2 P3]	FastPath Double	6
VPSADBW_128_reg	MAL[P2 P3]	FastPath Double	2
VPSHAB_128_mem	XBR[P1]	FastPath Single	7
VPSHAB_128_reg	XBR[P1]	FastPath Single	3
VPSHAD_128_mem	XBR[P1]	FastPath Single	7
VPSHAD_128_reg	XBR[P1]	FastPath Single	3
VPSHAQ_128_mem	XBR[P1]	FastPath Single	7
VPSHAQ_128_reg	XBR[P1]	FastPath Single	3
VPSHAW_128_mem	XBR[P1]	FastPath Single	7
VPSHAW_128_reg	XBR[P1]	FastPath Single	3
VPSHLB_128_mem	XBR[P1]	FastPath Single	7
VPSHLB_128_reg	XBR[P1]	FastPath Single	3
VPSHLD_128_mem	XBR[P1]	FastPath Single	7
VPSHLD_128_reg	XBR[P1]	FastPath Single	3
VPSHLQ_128_mem	XBR[P1]	FastPath Single	7
VPSHLQ_128_reg	XBR[P1]	FastPath Single	3
VPSHLW_128_mem	XBR[P1]	FastPath Single	7
VPSHLW_128_reg	XBR[P1]	FastPath Single	3
VPSHUFB_128_mem	XBR[P1]	FastPath Single	7
VPSHUFB_128_reg	XBR[P1]	FastPath Single	3
VPSHUFD_128_mem	XBR[P1]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPSHUFD_128_reg	XBR[P1]	FastPath Single	2
VPSHUFW_128_mem	XBR[P1]	FastPath Single	6
VPSHUFW_128_reg	XBR[P1]	FastPath Single	2
VPSHUFLW_128_mem	XBR[P1]	FastPath Single	6
VPSHUFLW_128_reg	XBR[P1]	FastPath Single	2
VPSIGNB_128_mem	MAL[P2 P3]	FastPath Single	6
VPSIGNB_128_reg	MAL[P2 P3]	FastPath Single	2
VPSIGND_128_mem	MAL[P2 P3]	FastPath Single	6
VPSIGND_128_reg	MAL[P2 P3]	FastPath Single	2
VPSIGNW_128_mem	MAL[P2 P3]	FastPath Single	6
VPSIGNW_128_reg	MAL[P2 P3]	FastPath Single	2
VPSLLD_128_mem	XBR[P1]	FastPath Single	7
VPSLLD_128_reg	XBR[P1]	FastPath Single	3
VPSLLDQ_128_reg	XBR[P1]	FastPath Single	2
VPSLLQ_128_mem	XBR[P1]	FastPath Single	7
VPSLLQ_128_reg	XBR[P1]	FastPath Single	3
VPSLLW_128_mem	XBR[P1]	FastPath Single	7
VPSLLW_128_reg	XBR[P1]	FastPath Single	3
VPSRAD_128_mem	XBR[P1]	FastPath Single	7
VPSRAD_128_reg	XBR[P1]	FastPath Single	3
VPSRAW_128_mem	XBR[P1]	FastPath Single	7
VPSRAW_128_reg	XBR[P1]	FastPath Single	3
VPSRLD_128_mem	XBR[P1]	FastPath Single	7
VPSRLD_128_reg	XBR[P1]	FastPath Single	3
VPSRLDQ_128_reg	XBR[P1]	FastPath Single	2
VPSRLQ_128_mem	XBR[P1]	FastPath Single	7
VPSRLQ_128_reg	XBR[P1]	FastPath Single	3
VPSRLW_128_mem	XBR[P1]	FastPath Single	7
VPSRLW_128_reg	XBR[P1]	FastPath Single	3
VPSUBB_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBB_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBD_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBD_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBQ_128_mem	MAL[P2 P3]	FastPath Single	6

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VPSUBQ_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBSB_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBSB_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBSW_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBSW_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBUSB_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBUSB_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBUSW_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBUSW_128_reg	MAL[P2 P3]	FastPath Single	2
VPSUBW_128_mem	MAL[P2 P3]	FastPath Single	6
VPSUBW_128_reg	MAL[P2 P3]	FastPath Single	2
VPTEST_128_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
VPTEST_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
VPTEST_256_mem	microcode	microcode	13
VPTEST_256_reg	microcode	microcode	9
VPUNPCKHBW_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKHBW_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKHDQ_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKHDQ_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKHQDQ_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKHQDQ_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKHWD_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKHWD_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKLBW_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKLBW_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKLDQ_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKLDQ_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKLQDQ_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKLQDQ_128_reg	XBR[P1]	FastPath Single	2
VPUNPCKLWD_128_mem	XBR[P1]	FastPath Single	6
VPUNPCKLWD_128_reg	XBR[P1]	FastPath Single	2
VPXOR_128_mem	MAL[P2 P3]	FastPath Single	5
VPXOR_128_reg	MAL[P2 P3]	FastPath Single	2
VRCPPS_128_mem	FMA[P0 P1]	FastPath Single	9

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VRCPPS_128_reg	FMA[P0 P1]	FastPath Single	5
VRCPPS_256_mem	FMA[P0 P1]	FastPath Double	10
VRCPPS_256_reg	FMA[P0 P1]	FastPath Double	5
VRCPSS_128_mem	FMA[P0 P1]	FastPath Single	9
VRCPSS_128_reg	FMA[P0 P1]	FastPath Single	5
VROUNDPD_128_mem	CVT[P0]	FastPath Single	8
VROUNDPD_128_reg	CVT[P0]	FastPath Single	4
VROUNDPD_256_mem	CVT[P0]	FastPath Double	9
VROUNDPD_256_reg	CVT[P0]	FastPath Double	5
VROUNDPS_128_mem	CVT[P0]	FastPath Single	8
VROUNDPS_128_reg	CVT[P0]	FastPath Single	4
VROUNDPS_256_mem	CVT[P0]	FastPath Double	9
VROUNDPS_256_reg	CVT[P0]	FastPath Double	5
VROUNDSD_128_mem	CVT[P0]	FastPath Single	8
VROUNDSD_128_reg	CVT[P0]	FastPath Single	4
VROUNDSS_128_mem	CVT[P0]	FastPath Single	8
VROUNDSS_128_reg	CVT[P0]	FastPath Single	4
VRSQRTPS_128_mem	FMA[P0 P1]	FastPath Single	9
VRSQRTPS_128_reg	FMA[P0 P1]	FastPath Single	5
VRSQRTPS_256_mem	FMA[P0 P1]	FastPath Double	10
VRSQRTPS_256_reg	FMA[P0 P1]	FastPath Double	5
VRSQRTSS_128_mem	FMA[P0 P1]	FastPath Single	9
VRSQRTSS_128_reg	FMA[P0 P1]	FastPath Single	5
VSHUFPD_128_mem	XBR[P1]	FastPath Single	6
VSHUFPD_128_reg	XBR[P1]	FastPath Single	2
VSHUFPD_256_mem	XBR[P1]	FastPath Double	7
VSHUFPD_256_reg	XBR[P1]	FastPath Double	3
VSHUFPS_128_mem	XBR[P1]	FastPath Single	6
VSHUFPS_128_reg	XBR[P1]	FastPath Single	2
VSHUFPS_256_mem	XBR[P1]	FastPath Double	7
VSHUFPS_256_reg	XBR[P1]	FastPath Double	3
VSQRTPD_128_mem	FMA[P0 P1]	FastPath Single	42
VSQRTPD_128_reg	FMA[P0 P1]	FastPath Single	38
VSQRTPD_256_mem	FMA[P0 P1]	FastPath Double	43

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VSQRTPD_256_reg	FMA[P0 P1]	FastPath Double	39
VSQRTPS_128_mem	FMA[P0 P1]	FastPath Single	33
VSQRTPS_128_reg	FMA[P0 P1]	FastPath Single	29
VSQRTPS_256_mem	FMA[P0 P1]	FastPath Double	34
VSQRTPS_256_reg	FMA[P0 P1]	FastPath Double	30
VSQRTSD_128_mem	FMA[P0 P1]	FastPath Single	42
VSQRTSD_128_reg	FMA[P0 P1]	FastPath Single	38
VSQRTSS_128_mem	FMA[P0 P1]	FastPath Single	33
VSQRTSS_128_reg	FMA[P0 P1]	FastPath Single	29
VSTMXCSR_128_mem	microcode	microcode	NA
VSUBPD_128_mem	FMA[P0 P1]	FastPath Single	10
VSUBPD_128_reg	FMA[P0 P1]	FastPath Single	5
VSUBPD_256_mem	FMA[P0 P1]	FastPath Double	11
VSUBPD_256_reg	FMA[P0 P1]	FastPath Double	7
VSUBPS_128_mem	FMA[P0 P1]	FastPath Single	10
VSUBPS_128_reg	FMA[P0 P1]	FastPath Single	5
VSUBPS_256_mem	FMA[P0 P1]	FastPath Double	11
VSUBPS_256_reg	FMA[P0 P1]	FastPath Double	7
VSUBSD_128_mem	FMA[P0 P1]	FastPath Single	10
VSUBSD_128_reg	FMA[P0 P1]	FastPath Single	5
VSUBSS_128_mem	FMA[P0 P1]	FastPath Single	10
VSUBSS_128_reg	FMA[P0 P1]	FastPath Single	5
VTESTPD_128_mem	XBR[P1] STO[P3]	FastPath Double	6/2
VTESTPD_128_reg	XBR[P1] STO[P3]	FastPath Double	2/2
VTESTPD_256_mem	microcode	microcode	13
VTESTPD_256_reg	microcode	microcode	9
VTESTPS_128_mem	XBR[P1]/STO[P3]	FastPath Double	6/2
VTESTPS_128_reg	XBR[P1]/STO[P3]	FastPath Double	2/2
VTESTPS_256_mem	microcode	microcode	13
VTESTPS_256_reg	microcode	microcode	9
VUCOMISD_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
VUCOMISD_128_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2
VUCOMISS_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	5/2
VUCOMISS_128_reg	FMA[P0 P1]/STO[P3]	FastPath Double	2/2

Table 12: FPU Instruction Latencies

Instruction	Pipes	Decode Type	Latencies
VUNPCKHPD_128_mem	XBR[P1]	FastPath Single	6
VUNPCKHPD_128_reg	XBR[P1]	FastPath Single	2
VUNPCKHPD_256_mem	XBR[P1]	FastPath Double	7
VUNPCKHPD_256_reg	XBR[P1]	FastPath Double	3
VUNPCKHPS_128_mem	XBR[P1]	FastPath Single	6
VUNPCKHPS_128_reg	XBR[P1]	FastPath Single	2
VUNPCKHPS_256_mem	XBR[P1]	FastPath Double	7
VUNPCKHPS_256_reg	XBR[P1]	FastPath Double	3
VUNPCKLPD_128_mem	XBR[P1]	FastPath Single	6
VUNPCKLPD_128_reg	XBR[P1]	FastPath Single	2
VUNPCKLPD_256_mem	XBR[P1]	FastPath Double	7
VUNPCKLPD_256_reg	XBR[P1]	FastPath Double	3
VUNPCKLPS_128_mem	XBR[P1]	FastPath Single	6
VUNPCKLPS_128_reg	XBR[P1]	FastPath Single	2
VUNPCKLPS_256_mem	XBR[P1]	FastPath Double	7
VUNPCKLPS_256_reg	XBR[P1]	FastPath Double	3
VXORPD_128_mem	MAL[P2 P3]	FastPath Single	6
VXORPD_128_reg	MAL[P2 P3]	FastPath Single	2
VXORPD_256_mem	MAL[P2 P3]	FastPath Double	7
VXORPD_256_reg	MAL[P2 P3]	FastPath Double	3
VXORPS_128_mem	MAL[P2 P3]	FastPath Single	6
VXORPS_128_reg	MAL[P2 P3]	FastPath Single	2
VXORPS_256_mem	MAL[P2 P3]	FastPath Double	7
VXORPS_256_reg	MAL[P2 P3]	FastPath Double	3
XORPD_mem	MAL[P2 P3]	FastPath Single	6
XORPD_reg	MAL[P2 P3]	FastPath Single	2
XORPS_mem	MAL[P2 P3]	FastPath Single	6
XORPS_reg	MAL[P2 P3]	FastPath Single	2

B.5 Amended Latency for Selected FMA Instructions

The following table shows amended latency time for selected FMA instructions, where special cases are applied in which additional latency is accumulated.

Table 13. FPU Amended Latencies

Instruction	Pipes	Decode Type	Latencies
ADDPD_reg	FMA[P0 P1]	FastPath Single	6
ADDPS_reg	FMA[P0 P1]	FastPath Single	6
ADDSD_reg	FMA[P0 P1]	FastPath Single	6
ADDSS_reg	FMA[P0 P1]	FastPath Single	6
ADDSUBPD_reg	FMA[P0 P1]	FastPath Single	6
ADDSUBPS_reg	FMA[P0 P1]	FastPath Single	6
CMPPD_mem	FMA[P0 P1]	FastPath Single	6
CMPPS_mem	FMA[P0 P1]	FastPath Single	6
CMPSD_mem	FMA[P0 P1]	FastPath Single	6
CMSSS_mem	FMA[P0 P1]	FastPath Single	6
COMISD_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
COMISS_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
FADD_reg	FMA[P0 P1]	FastPath Single	6
FADDP_reg	FMA[P0 P1]	FastPath Single	6
FCOM_mem	FMA[P0 P1]	FastPath Single	6
FCOMP_mem	FMA[P0 P1]	FastPath Single	6
FIADD_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/6
FIMUL_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/6
FISUB_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/6
FISUBR_mem	CVT[P0]/FMA[P0 P1]	FastPath Double	4/6
FLD_mem	FMA[P0 P1]	FastPath Single	6
FMADDPD_reg	FMA[P0 P1]	FastPath Single	6
FMADDPS_reg	FMA[P0 P1]	FastPath Single	6
FMADDSD_reg	FMA[P0 P1]	FastPath Single	6
FMADDSS_reg	FMA[P0 P1]	FastPath Single	6
FMSUBPD_reg	FMA[P0 P1]	FastPath Single	6
FMSUBPS_reg	FMA[P0 P1]	FastPath Single	6
FMSUBSD_reg	FMA[P0 P1]	FastPath Single	6
FMSUBSS_reg	FMA[P0 P1]	FastPath Single	6
FMUL_reg	FMA[P0 P1]	FastPath Single	6

Table 13. FPU Amended Latencies

Instruction	Pipes	Decode Type	Latencies
FMULP_reg	FMA[P0 P1]	FastPath Single	6
FNMADDPD_reg	FMA[P0 P1]	FastPath Single	6
FNMADDPS_reg	FMA[P0 P1]	FastPath Single	6
FNMADDSD_reg	FMA[P0 P1]	FastPath Single	6
FNMADDSS_reg	FMA[P0 P1]	FastPath Single	6
FNMSUBPD_reg	FMA[P0 P1]	FastPath Single	6
FNMSUBPS_reg	FMA[P0 P1]	FastPath Single	6
FNMSUBSD_reg	FMA[P0 P1]	FastPath Single	6
FNMSUBSS_reg	FMA[P0 P1]	FastPath Single	6
FSUB_reg	FMA[P0 P1]	FastPath Single	6
FSUBR_mem	FMA[P0 P1]	FastPath Single	6
FSUBR_reg	FMA[P0 P1]	FastPath Single	6
FSUBRP_reg	FMA[P0 P1]	FastPath Single	6
MAXPD_mem	FMA[P0 P1]	FastPath Single	6
MAXPS_mem	FMA[P0 P1]	FastPath Single	6
MAXSD_mem	FMA[P0 P1]	FastPath Single	6
MAXSS_mem	FMA[P0 P1]	FastPath Single	6
MINPD_mem	FMA[P0 P1]	FastPath Single	6
MINPS_mem	FMA[P0 P1]	FastPath Single	6
MINSD_mem	FMA[P0 P1]	FastPath Single	6
MINSS_mem	FMA[P0 P1]	FastPath Single	6
MULPD_reg	FMA[P0 P1]	FastPath Single	6
MULPS_reg	FMA[P0 P1]	FastPath Single	6
MULSD_reg	FMA[P0 P1]	FastPath Single	6
SUBPD_reg	FMA[P0 P1]	FastPath Single	6
SUBPS_reg	FMA[P0 P1]	FastPath Single	6
SUBSD_reg	FMA[P0 P1]	FastPath Single	6
SUBSS_reg	FMA[P0 P1]	FastPath Single	6
UCOMISD_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
UCOMISS_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
VADDPD_128_reg	FMA[P0 P1]	FastPath Single	6
VADDPS_128_reg	FMA[P0 P1]	FastPath Single	6
VADDS_128_reg	FMA[P0 P1]	FastPath Single	6
VADDSS_128_reg	FMA[P0 P1]	FastPath Single	6

Table 13. FPU Amended Latencies

Instruction	Pipes	Decode Type	Latencies
VADDSUBPD_128_reg	FMA[P0 P1]	FastPath Single	6
VADDSUBPS_128_reg	FMA[P0 P1]	FastPath Single	6
VCMPDPD_128_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_256_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_128_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_256_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_128_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_256_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_128_mem	FMA[P0 P1]	FastPath Single	6
VCMPDPD_256_mem	FMA[P0 P1]	FastPath Single	6
VCOMISD_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
VCOMISS_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
VFMAADDPD_128_reg	FMA[P0 P1]	FastPath Single	6
VFMAADDPS_128_reg	FMA[P0 P1]	FastPath Single	6
VFMAADDSD_128_reg	FMA[P0 P1]	FastPath Single	6
VFMAADDSS_128_reg	FMA[P0 P1]	FastPath Single	6
VFMAADDSUBPD_128_reg	FMA[P0 P1]	FastPath Single	6
VFMAADDSUBPS_128_reg	FMA[P0 P1]	FastPath Single	6
VFMSUBADDPD_128_reg	FMA[P0 P1]	FastPath Single	6
VFMSUBADDPS_128_reg	FMA[P0 P1]	FastPath Single	6
VFMSUBPD_128_reg	FMA[P0 P1]	FastPath Single	6
VFMSUBPS_128_reg	FMA[P0 P1]	FastPath Single	6
VFMSUBSD_128_reg	FMA[P0 P1]	FastPath Single	6
VFMSUBSS_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMAADDPD_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMAADDPS_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMAADDSD_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMAADDSS_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMSUBPD_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMSUBPS_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMSUBSD_128_reg	FMA[P0 P1]	FastPath Single	6
VFNMSUBSS_128_reg	FMA[P0 P1]	FastPath Single	6
VFRCZPD_128_reg	CVT[P0] FMA[P0 P1]	FastPath Single	4/6
VFRCZPS_128_reg	CVT[P0] FMA[P0 P1]	FastPath Single	4/6
VFRCZSD_128_reg	CVT[P0] FMA[P0 P1]	FastPath Single	4/6
VFRCZSS_128_reg	CVT[P0] FMA[P0 P1]	FastPath Single	4/6

Table 13. FPU Amended Latencies

Instruction	Pipes	Decode Type	Latencies
VMAXPD_128_mem	FMA[P0 P1]	FastPath Single	6
VMAXPS_128_mem	FMA[P0 P1]	FastPath Single	6
VMAXSD_128_mem	FMA[P0 P1]	FastPath Single	6
VMAXSS_128_mem	FMA[P0 P1]	FastPath Single	6
VMINPD_128_mem	FMA[P0 P1]	FastPath Single	6
VMINPS_128_mem	FMA[P0 P1]	FastPath Single	6
VMINSD_128_mem	FMA[P0 P1]	FastPath Single	6
VMINSS_128_mem	FMA[P0 P1]	FastPath Single	6
VMULPD_128_reg	FMA[P0 P1]	FastPath Single	6
VMULPS_128_reg	FMA[P0 P1]	FastPath Single	6
VMULSD_128_reg	FMA[P0 P1]	FastPath Single	6
VMULSS_128_reg	FMA[P0 P1]	FastPath Single	6
VPXOR_128_mem	MAL[P2 P3]	FastPath Single	6
VSUBPD_128_reg	FMA[P0 P1]	FastPath Single	6
VSUBPS_128_reg	FMA[P0 P1]	FastPath Single	6
VSUBSD_128_reg	FMA[P0 P1]	FastPath Single	6
VSUBSS_128_reg	FMA[P0 P1]	FastPath Single	6
VUCOMISD_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2
VUCOMISS_128_mem	FMA[P0 P1]/STO[P3]	FastPath Double	6/2

Appendix C Tools and APIs for AMD Family 15h ccNUMA Multiprocessor Systems

The following sections discuss tools and APIs available to support AMD Family 15h ccNUMA multiprocessor systems.

C.1 Thread/Process Scheduling, Memory Affinity

This following sections discuss tools and APIs available for assigning thread/process and memory affinity under various operating systems.

C.1.1 Support Under Linux[®]

Linux provides command-line utilities to explicitly set process/thread and memory affinity to both nodes and cores on a node. Additionally, **libnuma**, a shared library, is provided for more precise affinity control from within applications.

C.1.1.1 Controlling Process and Thread Affinity

The Linux command-line utilities offer high-level affinity control options. The **numactl** utility is a command line tool for running a process with a specific *node* affinity.

For example, to run the `foobar` program on the cores of node 0, enter the following at the command prompt:

```
numactl --cpunodebind=0 foobar
```

Application and kernel developers can use the **libnuma** shared library, which can be linked to programs and offers a stable API for setting thread affinity to a given node or set of nodes. Interested developers should consult the Linux **man** pages for details on the various functions available.

On a multicore processor, a process or thread affined to a particular node using the tools or API discussed above may still migrate back and forth between the cores of that node. This migration may or may not affect performance.

The **taskset** utility is a command-line tool for setting the process affinity for a specified program to any core. For example, to run the `foobar` program on the first two cores of node 0, enter the following on the command line:

```
taskset -c 0,1 foobar
```

In SuSE Linux Enterprise Server 10/10.1, the **numactl** utility can be used instead of **taskset** to set process affinity to any core. To repeat the previous example using **numactl**:

```
numactl --physcpubind=0,1 foobar
```

Linux provides several functions by which to set the thread affinity to any core or set of cores:

- **pthread_attr_setaffinity_np()** and **pthread_create()** are provided as a part of the older **nptl** library; they can be used to set the affinity parameter and then create a thread using that affinity.
- **sched_setaffinity()** system call and **schedutils** scheduler utilities.

In “Scheduling Single and Multithreaded Applications on Multiprocessor Systems” on page 191, we recommend scheduling multi-threaded applications in which each thread operates on independent data with one thread per node if possible. For example, if a program create four threads and runs on a 4-node system, the **numactl** command line preface for this program might be

```
numactl --cpunodebind=0,1,2,3 <program>
```

However, this command preface does not restrict any movement of threads to any core in the system, and thus does not achieve the desired affect. Using the **--physcpubind** option to restrict threads to specific cores would possibly yield

```
numactl --physcpubind=0, 8, 16, 24
```

Assuming four compute units or eight cores per node in this example, this would limit program execution to only the first core on each of four nodes. However, we are left with the possibility that any of the four threads could migrate between these four cores, possibly causing remote memory accesses and lower performance. Thus, for this scenario, it is better for each thread to set its own affinity mask, using, for example, **sched_setaffinity()** and specifying a single core.

C.1.1.2 Controlling Memory Affinity

Both **numactl** and **libnuma** library functions can be used to set memory affinity[5]. Memory affinity set by tools like **numactl** applies to all the data accessed by the entire program (including child processes). Memory affinity set by **libnuma** or other library functions can be made to apply only to specific data as determined by the program.

Both **numactl** and the **libnuma** API can be used to set a preferred memory affinity instead of forcibly binding it. In this case the binding specified is a hint to the OS; the OS may choose not to adhere to it.

At a high level, normal first touch binding, explicit binding and preferred binding are all available as memory policies on Linux.

By default, when none of the tools/API is used, Linux uses the first touch binding policy for all data. Once memory is bound, either by the OS, or by using the tools/API, the memory will normally remain resident on that node for its lifetime.

C.1.2 Support under Microsoft® Windows®

In the Microsoft Windows environment, the function to bind a thread on particular core or cores is **SetThreadAffinityMask()**. The function to run all threads in a process on particular core or cores is **SetProcessAffinityMask()**[8].

The function to set memory affinity for a thread is **VirtualAllocEx()**[9]. This function gives the developer the choice to bind memory immediately on allocation or to defer binding until first touch.

The `start /affinity xxx` command can be used to confine all of a process's threads to a specified subset of cores in the system. The memory that these threads allocate or touch will also be confined to that subset of cores. In addition, several Microsoft Enterprise products provide NUMA support and configurability, such as SQL Server 2008 [10] and IIS [11].

If an application relies on heaps in Windows, we recommend using a low fragmentation heap (LFH) and using a local heap instead of a global heap[12][13].

By default, Windows uses the first touch binding policy for all data. Once memory is bound to a node, it normally resides on that node for its lifetime.

C.1.3 Hardware Support for System Topology Discovery

AMD Family 15 processors support new hardware features to facilitate the discovery of hardware topology and configuration about the possible implementation choices of cache size, number of compute units per node sharing an L3 cache, nodes per processor package, etc. In some cases, software may require more detailed knowledge of these characteristics than that provided by the APIs mentioned above. This information can help software implement the optimizations described in Chapter 11.

This information can be accessed by use of the CPUID instruction, which can access the contents of several new registers, as described in the *BIOS and Kernel Developer's Guide (BKDG) for AMD Family 15h Processors*, order# 42301.

C.1.4 Support for Stability

There may be instances of software runtime execution which cannot tolerate the variability aspects of Application Power Management (APM) boost support. For example, when configuring a workload which makes a APIC reservation request, it may be undesirable to possibly allow the processor to deliver variable performance that cannot be sustained. To disable Application Power Management boost support do the following:

- For CPUID 8000_0007 detect that APM boost is enabled
- APM is enabled if all of the following conditions are true:
 - MSRC001_0015[CpbDis] == 0 for all cores.
 - D18F4x15C[ApmMasterEn] == 1.
 - D18F4x15C[BoostSrc] == 1.
 - D18F4x15C[NumBoostStates] != 0.

Change the ApmMasterEn by setting bit 7 to 0 which will disable power management.

C.2 Tools and APIs for Memory Node Interleaving

This section discusses tools and APIs available for performing node interleaving under various operating systems.

C.2.1 Support under Linux[®]

Linux provides several ways for an application to use memory node interleaving [5].

- **numactl** is a command line tool, which is used for node interleaving all memory accessed by a program across a set of chosen nodes.

For example, to interleave all memory accessed by program `foobar` on nodes 0 and 1, use:

```
numactl --interleave=0x03 foobar
```

- **libnuma** offers several functions a program can use to interleave a given memory region across a set of chosen nodes.

Linux only supports the round robin node interleaving policy.

C.2.2 Support under Solaris[™]

Solaris offers an API called **madvise**, which can be used with the **MADV_ACCESS_MANY** flag to tell the OS to use a memory policy that causes the OS to bind memory randomly across the nodes. This offers behavior similar to the round robin node interleaving of memory offered by Linux.

This random policy is the default memory placement policy used by Solaris for shared memory.

C.2.3 Support under Microsoft[®] Windows[®]

Microsoft Windows does not offer node interleaving.

C.2.4 Memory Node Interleaving Configuration in the BIOS

AMD family 15h ccNUMA multiprocessor systems can be configured in the BIOS to interleave all memory across all nodes on a page basis (4KB for regular pages and 2M for large pages). Enabling node interleaving in the BIOS overrides the use of any tools and causes the OS to interleave all available system memory across all nodes in a round robin manner.

C.3 OpenMP

OpenMP is an Application Programming Interface that provides shared memory parallel programming constructs for C/C++ and Fortran. OpenMP can be used for developing applications to run on a variety of multicore machines. OpenMP is supported by a wide variety of C and Fortran

compilers and is available for both Linux and Windows programming environments. For more information, see www.openmp.org.

Appendix D NUMA Optimizations for I/O Devices

D.1 AMD64 System Topologies

AMD family 15h and AMD Opteron™ systems range from single-node desktop to two, four, and eight node servers with the potential for even larger systems. Each node in the AMD family 15h system consists of four compute units attached to an integrated memory controller and up to four HyperTransport™ links. I/O devices connect to the system over non-coherent HyperTransport™ links. The term *non-coherent I/O* refers to a configuration in which the processor does not cache memory residing on an I/O device and in which the I/O device does not cache system memory shared by the processors.

The integrated Northbridge converts requests issued on the non-coherent HyperTransport link into coherent requests before forwarding them into the coherent fabric. Since memory is sharable and coherent to all of the processors in the system, it is possible that the latest copy of the requested memory location is not in memory, but is located in one of the processor caches, if that memory location is cached and that cache-line has subsequently been modified by one of the processor cores. In this case, the processor cache holding the data must recognize that it is the owner of the data and must return this data to the I/O device. This is referred to as *probing* the caches. This implies that the latency to obtain the data will be directly affected by the location of the latest data at the time the I/O device requests it.

Physical I/O devices are typically connected to PCI, PCI-X, or PCI Express interfaces that are bridged through a chipset component to non-coherent Hypertransport links. Alternatively, some I/O devices may attach directly to non-coherent Hypertansport links by means of custom components or an HTX interface. There may be multiple non-coherent Hypertransport links in a system that provide sufficient bandwidth and fan-out to numerous PCI busses and I/O devices. These multiple links may be attached to multiple different processor nodes. Thus, I/O devices can be considered to have NUMA properties and have an associated "home" node. The optimizations and recommendations described in this appendix are intended to leverage the NUMA properties associated with I/O devices to improve performance.

D.2 Optimization Strategy

The OS can help manage device NUMA topology information to supplement existing OS NUMA support by means of affinity for device-driver buffers and resources (such as interrupt-pin assignments, interrupt service routines (ISR), deferred procedure calls (DPC)), when device drivers are loaded. The application I/O-thread is the portion of application code that calls the device's I/O API, which in turn will call the I/O device-driver. On AMD family 15h systems, applications that

interface with an I/O device (through API/device drivers) usually perform best when the following conditions are true:

- Direct memory access (DMA) transfers to/from the I/O device should access memory on the node to which the I/O device is attached. From the software point of view, this is equivalent to locating I/O buffers on the node where the associated I/O device is attached.
Latencies of local memory reads by I/O devices can be ~10% to 25% lower than reads that access remote or non-node local memory.
- The device uses a significant amount of memory-mapped I/O (MMIO) from the processor directly to the device. Memory-mapped I/O is much more efficient than programmed I/O (PIO) from the point of view of the processor. PIO instructions force a strict ordering of memory reads and writes, which can cause a significant reduction in instruction throughput on a processor core. When I/O locations are memory-mapped, the semantics of their access are such that writes may be buffered and combined, greatly reducing any ordering restrictions.
- The actual code modules for the application and device driver are located on the node to which the I/O device is attached, i.e., the linear address of the code is physically mapped to memory on that node.

To take advantage of these performance factors, the following practices are recommended:

- Locate a driver's specific I/O device in the system and allocate memory to the node where the non-coherent Hypertransport link is located.
- Specify interrupt (ISR and DPC) affinity to a specific processor core and node and relocate driver code to the node where the driver, ISR, and DPC will execute.
- Stream data to buffers on the node where the I/O device is accessed, by means of non-coherent Hypertransport links.

The conditions listed above increase performance primarily due to:

- Latency between the I/O device to the NUMA-closest memory and processor cache(s) is lower.
- OS scheduler opportunities—the code-flow sequence from application code to I/O API to kernel dispatcher to device driver code remains on one processor core (for OSs that run a device from specific processor cores) and is uninterrupted.
- The code fetched for the application and device driver is mapped to the closest memory. The coherent HyperTransport link no longer needs to fetch code blocks from far nodes, freeing coherent HyperTransport bandwidth for other traffic.

NUMA-aware applications and drivers will ensure that your software will run with the highest performance possible across the many varying system topologies, from the single node desktop with a single noncoherent HyperTransport link to a large scale server with multiple noncoherent HyperTransport links.

Determining Number Of Nodes in AMD Family 15h Processor Systems in User-Mode

To implement the aforementioned optimization strategy, user mode programs can take advantage of the NUMA API support functions provided by various OSs to determine the number of nodes in the system, schedule threads, and allocate memory. While existing NUMA APIs that provide processor-to-memory NUMA support are well-established, support functions to enumerate the system topology for I/O devices are less mature.

For example, while Microsoft Vista and Windows Server 8 operating systems now provide support for determining a "home node" for an I/O device by means of the `SetupDiGetDeviceProperty()` and `GetNumaProximityNode()` functions, obtaining the same information from Linux is less straightforward, but available through `/sys/devices/pci*/*/local_cpus`.

A user-mode program must therefore create/allocate buffers and must depend to a large degree upon the operating system to allocate buffers optimally for the system into which the device is plugged. Currently most NUMA configuration information comes from the device-drivers and OS.

Applications should strive to allocate memory and schedule threads in a consistent manner (assigning threads to specific processors on a node) by using the NUMA API and allowing device drivers to do whatever is optimal—such as remapping linear to physical static I/O-buffers closer to the I/O device, copying buffers, etc., depending on the specific device and its latency characteristics. By grouping the threads that perform I/O to specific compute units on a node, it becomes easier to dynamically switch the entire device NUMA configuration to another compute unit or set of compute units on another node to verify whether performance increases or decreases. This may be necessary in cases where there is insufficient OS support to determine the NUMA node to which a particular I/O device is attached. However, the memory buffers do not switch automatically, even if the thread switches. The buffers remain on the nodes to which they were allocated and must be reallocated, if desired.

The examples that follow use Microsoft APIs. Developers who are creating applications running under Linux™, Solaris and other operating systems should consult the NUMA API documentation specific to their target environment. The guidelines in following section are recommended to compliment operating system mechanisms to correctly establish the optimal device NUMA configuration for a device.

D.3 Identifying Nodes that Have Noncoherent HyperTransport™ I/O Links

This section describes some of the mechanisms that can be used to determine the I/O device configuration with respect to NUMA nodes. In cases where the functionality provided by `SetupDiGetDeviceProperty()` and `GetNumaProximityNode()` is not available, the techniques described here can be used to determine which nodes possess non-coherent Hypertransport links and which PCI buses are down-stream from them. This information can be used to determine a logical "home" node for each I/O device.

The following sections include examples of how to associate an I/O device to a node by using the device's known PCI Bus number and by using the memory-mapped I/O address of the device that is provided in the device's base address register (BAR).

In general, AMD Family 15h processor-based systems can have up to eight nodes. Every I/O-device in the system must be connected to a node by means of a non-coherent Hypertransport link. Each node can be identified in the PCI configuration space using the PCI device ID on PCI bus 0 starting with device 24 (18h) function 0h and counting up to device 31 (1Fh) function 0h. Systems can have up to eight nodes; each node appears as one PCI device. Nodes with noncoherent HyperTransport links can be identified by reading the initialized values of the link connected bit (bit 0) and the noncoherent bit (bit 2) of the HyperTransport link type register for each HyperTransport link. There is one link type register for each link. Family 15h processors support up to four Hypertransport links; for each node, the registers are located at function 0h, registers 98h, B8h, D8h, F8h (links 0 through 3, respectively). The layout of the link type registers is shown in Figure 11.

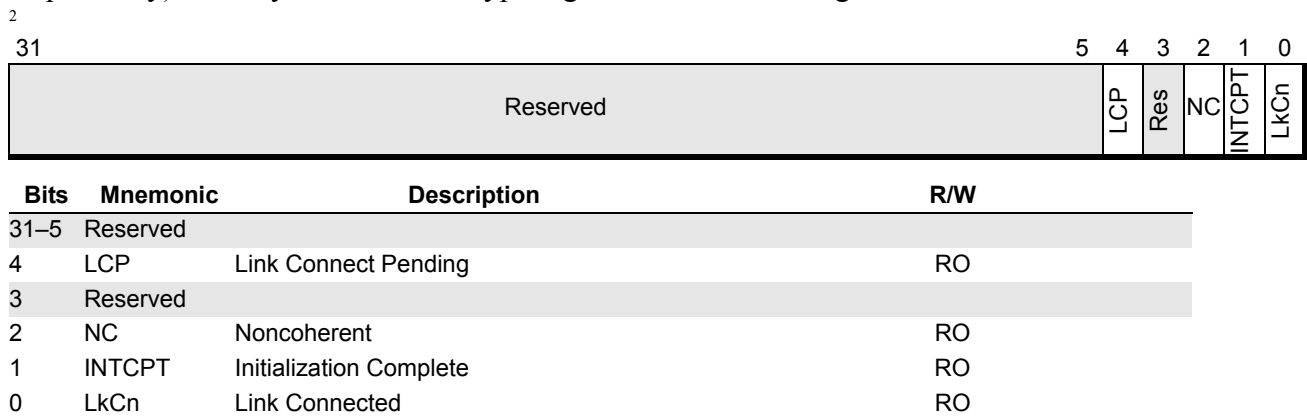


Figure 11. Link Type Registers F0x[F8, D8, B8, 98]

Only One Node in System Has Noncoherent HyperTransport I/O Links

If only one node in the system contains noncoherent HyperTransport link (or links), then record that node number; this node contains all I/O-Devices in the system. Memory buffers that are accessed by the I/O device(s) for DMA should be allocated on this node for lowest device-to-memory latency. Optimally, the I/O driver code should also run on a processor core on this node, resulting in the lowest latency for MMIO writes/reads to the device.

More Than One Node in System Has Noncoherent HyperTransport I/O Links

If more than one node contains a noncoherent HyperTransport link (or links), then the driver will need to associate the specific I/O device to a specific node in the system. As mentioned above, it may be possible on systems running Microsoft operating systems to obtain this information by means of the `SetupDiGetDeviceProperty()` and `GetNumaProximityNode()`, or from Linux systems through

`/sys/devices/pci*/*/local_cpus`. However, if these mechanisms are not available, it is also possible to use either of the following methods:

- Using the device's PCI-configuration base address register (BAR), software determines the node to which the I/O-Device is attached by comparing its base address to the contents of the MMIO routing table.
- Using the device's bus-number, the software obtains the node to which the I/O device is attached, based on the PCI bus to which PCI configuration cycles are steered.

Be sure to check the operating-system's API for other possible methods as well.

Determining the Location of the I/O-Device Using PCI-Configuration Base Address Registers

The first method uses the MMIO base and limit address registers. Figures 12 through 14 below show the location of the components of the base and limit addresses in PCI configuration space registers. The base and limit addresses are formed by reading these registers and masking off and merging the appropriate bits. There are a total of 12 address ranges determined by the base and limit addresses. Each address is formed by combining bits 31:8 from the low register with bits 7:0 from the high register (for the base) or bits 23:16 from the high register (for the limit). The result of concatenation of these two bit fields is bit field 47:16 of the address. Each address is aligned on a 64K byte boundary. See the BKDG for more details on PCI configuration space registers.

PCI/PCI-X/PCIe devices have a base address register (BAR) that allows the system BIOS and OS to map the device into the system address space. The BAR is either a 64-bit octword-aligned address in the eight bytes starting at offset 15h, or a 32-bit octword-aligned address in the four bytes starting at offset 15h. Bits 2:1 of the byte at offset 15h distinguish the size of BAR as follows (see the *PCI Local Bus Specification* for more details.):

Table 14. Size of Base Address Register

Value	Base Address Register Size	Device Location
00	32-bit decoder	Device is located anywhere in lower 4GB address space.
01	32-bit decoder	Device is located in lower 1MB address space.
10	64-bit decoder	Device is located anywhere in 2^{64} address space.
11	32-bit decoder	Reserved.

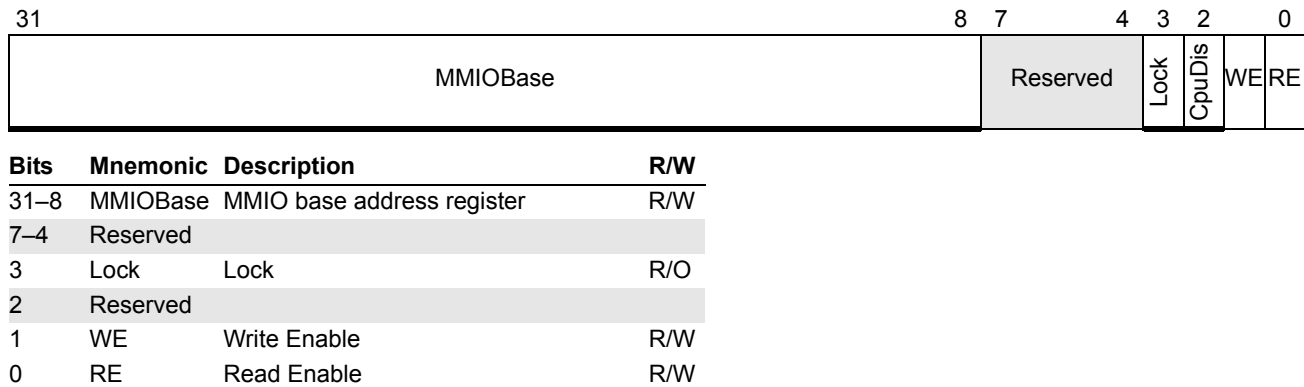


Figure 12. MMIO Base Low Address Registers F1x[B8h, B0h, A8h, A0h, 98h, 90h, 88h, 80h]

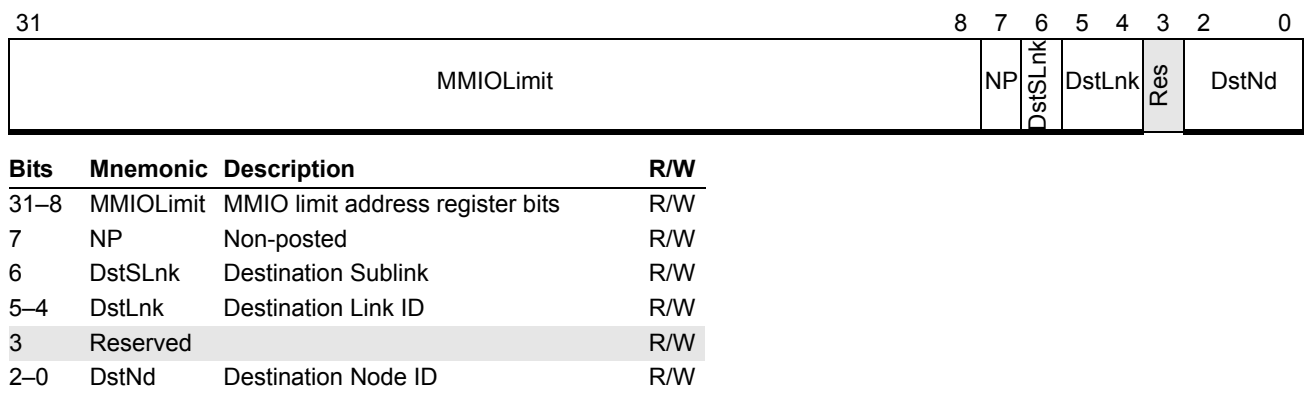


Figure 13. MMIO Limit Low Address Registers F1x[1BCh, 1B4h, 1ACh, 1A4h, 9Ch, 94h, 8Ch, 84h]

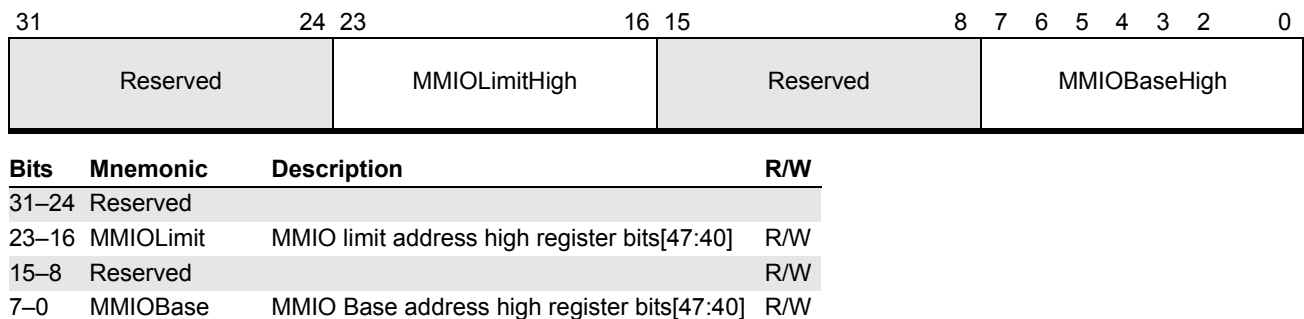


Figure 14. MMIO Base/Limit High Address Registers F1x[1CCh, 1C8h, 1C4h, 1C0h, 19Ch, 198h, 194h, 190h, 18Ch, 188h]

Most hardware is designed so that the BAR register contents can be read using MMIO or PCI configuration reads, if the driver can obtain the bus, device, and function of its device. Driver-code reads the BAR field and uses the address to find the node in the system that acts as the device bridge to decode this address in order to steer MMIO cycles into the appropriate down-stream bus (PCI, PCIe, PCI-X).

The following procedure uses the physical address of the device as an inclusive field to find the correct node hosting the device. This procedure compares addresses at a granularity of 64K, which is the granularity of allocation for MMIO regions.

- Step 1.** Create 64-bit or 32-bit integer from the device's BAR Registers based on bits 2:1.
Mask off bottom 4 bits (3:0) to create `INT64 deviceBaseAddress`.
- Step 2.** Select any node in the system (node 0, etc) and check each of the 8 memory-mapped I/O address map regions to determine the node that will decode the I/O-Device. The memory-mapped I/O address map registers decode MMIO, if the physical address is greater than the base field and is less than or equal to the limit field programmed into these registers.
- Step 3.** Initialize loop counter to 12.
- Step 4.** Move 0x80h to variable `startIndex` and 0x180 to variable `highIndex` (Bus 0, Device N (N=AMD Node Number, 18h, 19h, etc), Function 1, Registers 0x80 and 0x180, respectively).
- Step 5.** Read PCI configuration register `startIndex` into variable `(int64) nodeMMIOaddress`.
- Step 6.** Clear the lower 8 bits (bits 7:0) of `nodeMMIOaddress`.
- Step 7.** Left-shift `nodeMMIOaddress` by 8.
- Step 7a.** Read PCI configuration register `highIndex` into `int64` variable `mmioHigh`. Save a copy in `int64` variable `orgMMIOhigh`.
- Step 7b.** Left shift `mmioHigh` by 40. Then clear the most significant bits 63:48 to zero.
- Step 7c.** Place the logical OR of `mmioHigh` and `nodeMMIOaddress` into `nodeMMIOaddress`.
- Step 8.** Now get the node's MMIO limit by reading register `startIndex + 4` into `int64` variable `mmioLimit`. Save a copy of `mmioLimit` to temporary variable `(int64) orgMMIOlimit`.
- Step 9.** Clear `mmioLimit[7:0]`. (See Step 7.)
- Step 10.** Left-shift `mmioLimit` value by 8, then OR in bottom 16-bits to all '1's (0xFFFF). This is because the hardware effectively uses all '1's in lower 16-bits of the address.
- Step 10a.** Left shift `orgMMIOhigh` by 24, and then clear bits 63:48 and 39:0 to zero.
- Step 10b.** Place the logical OR of `orgMMIOhigh` and `mmioLimit` into `mmioLimit`.

- Step 11.** Compare `deviceBaseAddress`, to see if it is greater than or equal to `nodeMMIOaddress`. If so, we then need to check that `deviceBaseAddress` falls within the MMIO limit of the node, do this by checking that `deviceBaseAddress` is lower than or equal to `mmioLimit`. If this is true, we're done, save the node ID number by extracting bits 2:0 in `orgMMIOlimit`.
- Step 12.** If Step 11 did not provide match, increment `startIndex` by 8 (i.e. 80h becomes 88h, which is the next Memory-Mapped I/O Base Low register), increment `highIndex` by 4 (i.e. 180h becomes 184h, which is the next Memory-Mapped Base/Limit High I/O register), decrement the loop counter and loop back to Step 5.

Determining Where the I/O-Device Is Located Using PCI-Bus Number of Device

Some operating systems provide a way for the driver to get system resources assigned to the device by use of a resource descriptor. For example Windows kernel-mode drivers can use the `CM_PARTIAL_RESOURCE_DESCRIPTOR` to obtain the bus number, interrupt-pin assignment, etc. Based on this information, the node upstream from the device can be determined by querying the configuration map registers (see Figure 15 on page 326).

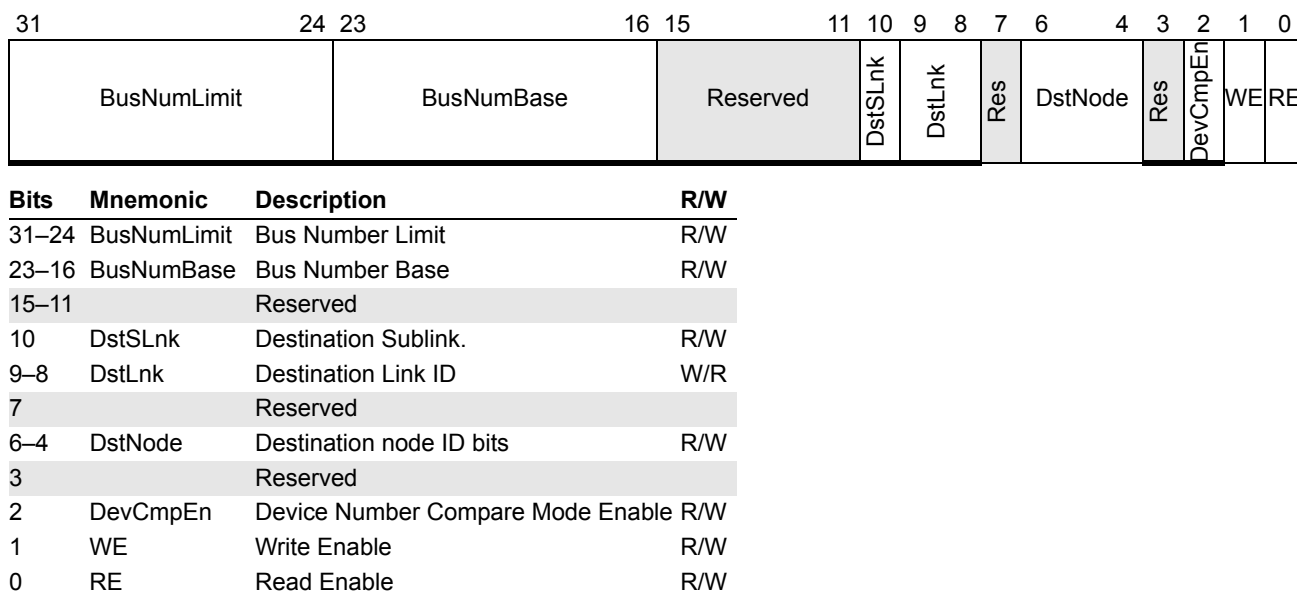


Figure 15. Configuration Map Registers F1x[E0h, E4h, E8h, ECh]

First, get the PCI bus number of the device.

- Step 1.** Get Bus number. Create 32-bit integer. Save bus as `DevBusNum`.
- Step 2.** Check the bus number against the contents of the four configuration base/limit registers to determine the node with which the bus is associated. These registers reside in the Northbridge PCI configuration space at offsets E0h through ECh. Each register maps a contiguous range of PCI bus numbers to a particular node, allowing up to four separate bus

ranges to be mapped to as many as four different nodes. The registers are replicated on each node so it is necessary to check the registers of only one node.

The remaining steps constitute a loop to access all four registers and implement step 2. Use the configuration map registers (Figure 15) to determine the node to which the device is attached.

- Step 3.** Initialize loop index `cm_node` to `E0h`.
- Step 4.** Read Northbridge PCI configuration space offset `cm_reg` and test the device compare mode enable bit (bit 2 `DevCmpEn`).
- Step 5.** If the `DevCmpEn` bit is set, use the PCI bus Device number in place of the PCI bus number for this check (save device number as `DevBusNum`).
- Step 6.** Read each configuration base/limit register and compare `DevBusNum` against the range defined by the `BusNumBase` and `BusNumLimit` fields until the desired bus range is found.
- Step 7.** Get the node number from the `DstNode` field of that register. The I/O device is attached to this node and the device's buffers should be allocated on this node. The range comparison succeeds if `DevBusNum` is greater than or equal to `BusNumBase` and less than or equal to `BusNumLimit`.
- Step 8.** If not successful, increment `cm_node` by 4 and return to step 3.

D.4 Access of PCI Configuration Register

Kernel mode drivers can use the operating system's low-level port access functions to read PCI configuration registers in the AMD Family 15h processor and integrated host bridge. These registers specify the system topology—the nodes on which each device resides.

A brief description of how to generate PCI configuration space reads is described below. Consult the *BIOS and Kernel Developer's Guide for AMD Family 15h Processors, order# 42301*, for a more detailed description of PCI configuration space. The following scheme uses a configuration index port (configuration address register `0CF8h`) as shown in Figure 16 and a configuration data port (configuration data register `0CFCh`) as shown in Figure 17.

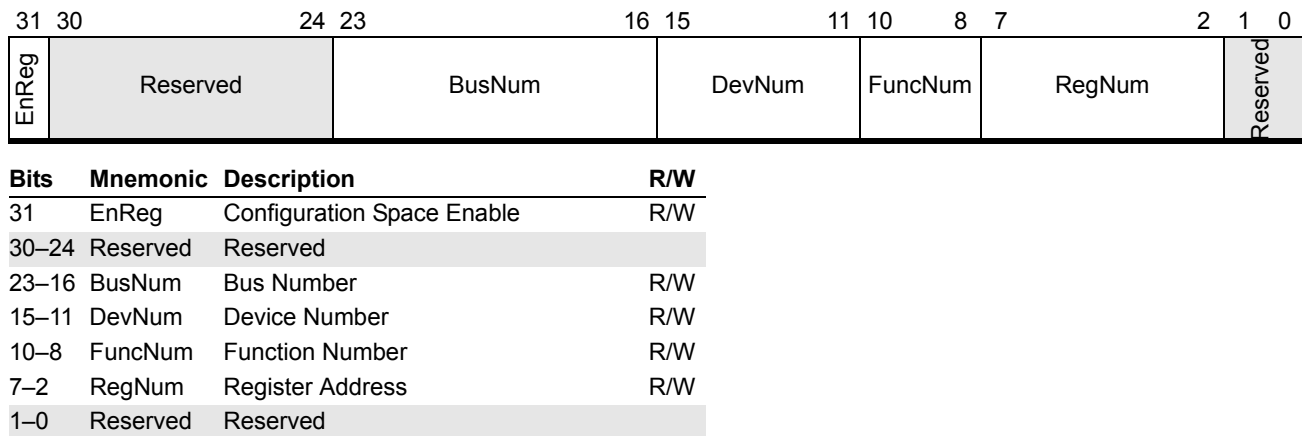


Figure 16. Configuration Address Register (0CF8h)




Figure 17. Configuration Data Register (0CFCh)

Use of Index and Data Ports

For thread-safety reasons, kernel mode function drivers should avoid performing the I/O directly to ports 0CF8h and 0CFCh in order to ensure that only a single thread is using the index and data ports exclusively.

The following pseudocode shows how (after ensuring exclusive access), an OS support routine or command-line debugger, performs I/O to read PCI configuration registers in the AMD Processor/Northbridge.

```
BUS 0x0, Device x18h (24), Function 0x1, register 0x60h // Node ID Register:
```

```
unsigned int busNum;
unsigned int devNum;
unsigned int funcNum;
unsigned int regNum;
unsigned int pci_registerSelect;
unsigned int pci_configData;
```

```
// Setup desired bus, device, and function number.
```



```

busNum=0x0;
devNum=0x18;
funcNum=0x1;
regNum=0x60;

// Setup the register with bus, device, and function.
// Also set the enable register read bit 31 0x80000000.

pci_register=(pci_registerSelect | 0x80000000);
pci_register=(pci_registerSelect | (busNum << 16));
pci_register=(pci_registerSelect | devNum << 11);
pci_register=(pci_registerSelect | (funcNum <<8));
pci_register=(pci_registerSelect | regNum);

// setup for PCI-configuration Read

#define PCI_CONFIGURATION_ADDRESS 0xCF8
#define PCI_CONFIGURATION_DATA 0xCFC

__asm{
    mov     edx, PCI_CONFIGURATION_ADDRESS;
    mov     eax, pci_registerSelect;
    out     edx, eax; //32-bit write
    mov     edx, PCI_CONFIGURATION_DATA
    in      eax, edx// read data
    mov     pci_configData, eax
    } // endasm here

```

When the system is operating under the runtime environment, the ideal thread-safe method by which to access the device-specific PCI configuration space is to use the operating system's PCI-bus driver. The methods to do this depend on the operating system.

D.5 I/O Thread Scheduling

Optimization

❖ Keep the I/O thread scheduled on the same node as the I/O device buffers, and allocate the I/O device buffers on the node where the I/O device is located.

Rationale

By assigning processor affinity, threads can be maintained on the node closest to the I/O device buffers. Operating system functions, such as `VirtualAllocExNuma()` can be used to specify the node to which to allocate the I/O device buffer-memory. The node selected should be the same node where the I/O device can be found down-stream via the noncoherent Hypertransport link. Use the mechanisms described above to determine the appropriate node. Set the processor affinity mask for the threads that will access these I/O buffers to indicate the cores of this node using operating system functions, such as `SetThreadAffinityMask()`.

D.6 Using Write-Only Buffers for Device Consumption

Optimization

❖ Use streaming-stores to optimize data for I/O-device consumption, if the application is using a write-only buffer.

Rationale

A processor writes to write-only output buffers, but does not read them. When using a write-only output buffer to be consumed by a device:

- Create the buffer on the node where the I/O device buffers are established. (This should be the node which has the down-stream path to the I/O device, as recommended in D.5 on page 329.)
- Use streaming-store instructions to avoid moving the buffer into the L1 cache of the writing CPU.
- Use the non-temporal streaming store instructions such as:

MOVNTI—Stream to memory-integer

MOVNTPS—Stream to memory-packed scalar floating-point

- Consult compiler intrinsic support to avoid assembly-language, such as:

```
void _mm_stream_ps(float * p , __m128 a ); // Uses MOVNTPS
```

When streaming the data by use of non-temporal instructions, data is write-combined on the node sending the data and then is forwarded to the node where the I/O-device buffer exists (see Appendix A, “Implementation of Write-Combining.”). Streaming the data has two advantages:

- First, there is no coherent HyperTransport read traffic from memory into the L2 cache.
- Second the device read/write latency to the buffer can be lower if the buffer is closer to device.

D.7 Using Interrupt Affinity

Optimization

❖ Make sure that interrupts from a device are serviced by a processor core that is on the same node as the device.

Rationale

Interrupt affinity is maintained by assuring that interrupt service routines (ISRs) from a device are run on a core that is on the same node as the device. There are various ways this enhances performance:

- Interrupt affinity can improve the service routine's cache locality—code and data have a better chance of being resident in the AMD family 15h processor's L1, L2, or L3 cache.

- Handling the interrupt on the node where the device is located lowers the latency (number of hops) for any MMIO reads or writes that the service routine may make to the device.
- Performance can be further increased if memory, including device buffers, that the service routine accesses are allocated on the same node.
- Use the techniques described in sections D.3 and D.4 to determine to which node a particular device is attached.

Setting up Interrupt Affinity

There are several ways to set up interrupt affinity.

- Use an OS-provided API to specify which processor cores should run the interrupt service routine (ISR). Some OSs can allow the device's ISR to run on a set of primary processor cores. Select the cores on the node to which the I/O device is connected downstream. Under Microsoft® Windows™, the driver can use the `WdfInterruptSetPolicy()` function in conjunction with the `IoConnectInterrupt()` function to accomplish this (consult Microsoft driver development documentation for full details). Device buffers that the ISR accesses should be allocated on this node as well.
- Devices that are message-signaled interrupt (MSI or MSI-X) capable can specify interrupt affinity in the MSI message. Specify processor cores on the node into which the I/O device is attached. Message-signaled interrupts offer many performance improvements over legacy PCI/PCI-X line-based interrupts. Less interrupt sharing occurs, which decreases the latency required to service the interrupt.
- Specify which processor cores should run the deferred procedure call (DPC) for further processing. If possible, queue the DPC that will be run after the ISR on the node into which the I/O-Device is attached. Map the buffers used by the DPC to the node closest to the I/O-Device. For example, the Microsoft Windows operating systems provides one DPC queue for each processor. Drivers can control the queue to which the operating system assigns the DPC. By default, when the driver calls `KeInsertQueueDpc()` or `IoRequestDpc()`, the DPC is queued on the currently active processor. In addition, drivers can specify the processor queue by calling `KeSetTargetProcessorDpc()` before calling `KeInsertQueueDpc()` or `IoRequestDpc()`.

Appendix E Remarks on the RDTSC(P) Instruction

The RDTSC and RDTSCP instructions are used to load the value of the time stamp counter (TSC) into the EDX:EAX register pair. These two instructions differ as follows: the RDTSC instruction may execute speculatively and out of order with respect to other instructions (except other RDTSC instructions), while the RDTSCP instruction does not. The RDTSCP also identifies the processor core on which it is executed. When a code sequence ending in an RDTSC instruction is executed, there is no guarantee that all prior instructions in the code sequence have been retired at the time when the TSC is read.

On the other hand, the RDTSCP instruction waits for all the previous instructions to be retired before reading the TSC, thus producing the expected TSC value. For this reason, it is recommended that RDTSCP be used to measure the clock cycles consumed by a hotspot function. Both RDTSC and RDTSCP are executed in program order with respect to other RDTSC(P) instructions.

If RDTSC is used, it should be accompanied by a separate serializing instruction (such as a CPUID instruction). In AMD Family 15h processors, the MFENCE instruction, which is not intercepted in virtualized environments, can be used in place of the CPUID instruction as a serializing instruction.

In the previous generation multi-core processors, each core has its own timestamp counter locked to its core. Starting with AMD Family 15h processors, there exists a single clock source in the NorthBridge for all timestamp counters in a processor and these counters are incremented in lockstep. This enables the cycle counter to provide monotonically increasing values at a constant rate even when the cores are in power saving modes. This behavior of RDTSC(P) is indicated if EDX bit 8 is set to 1, as returned by CPUID function 8000_0007h. Note that an operating system can write different values to each core's TSC and can establish or correct a core-to-core skew, after which the TSCs all advance in lockstep with each other and thus maintain a constant core-to-core skew. The precision with which software can synchronize the TSCs across cores is dependent on the approach used, as well as platform factors, such as the consistency of inter-core communication latency through shared memory. This precision is typically limited to a few tens of cycles. In particular, the skew may exceed inter-core communication latencies such that inter-core observation of TSC values may not show strict monotonic behavior—a TSC value acquired from another core after the local TSC is read may have a lower value. Successive TSC reads within a core, however, give monotonically increasing values.

The HWCR[TscFreqSel] bit is set by the BIOS to scale the TSC frequency to the P(0) frequency of the CPU. To calculate the elapsed wall clock time from the values returned by two RDTSC(P) instructions, use the following formula.

Elapsed Wall Clock Time (in seconds) = (Second RDTSC result – First RDTSC result) / CPU's P(0) frequency.

The following example explains how to use the RDTSC(P) instruction to measure the clock cycles consumed by a hot spot function:

```
unsigned long long RDTSCP()
{
    unsigned long long tsc;
    asm volatile (".byte 0x0f, 0x01, 0xf9" : "=A" (tsc) : : "%ecx");
    return tsc;
}

#define MFENCE() asm volatile("mfence")

unsigned long long compute_mfence_overhead(int N)
{
    int i;
    unsigned long long tsc, next_tsc;
    tsc = RDTSCP();
    for (i = 0; i < N; i++)
        MFENCE();
    next_tsc = RDTSCP();
    return ((next_tsc - tsc)/N);
}

unsigned long long time_hotspot(int N)
{
    unsigned long long tsc, next_tsc, avg_hotspot_time;
    int i;

    /* start the timer */
    tsc = RDTSCP();
    /* N is the number of iterations. The higher the value of N,
       the more accurate the avg_hotspot_time (except for OS context switches.*/
    for (i = 0; i < N; i++)
    {
        /* MFENCE is used to serialize the control flow
           between iterations. */
        MFENCE();
        hotspot(); /* HotSpot function */
    }
    next_tsc = RDTSCP();

    avg_hotspot_time = (next_tsc - tsc)/N - compute_mfence_overhead(N);
    return (avg_hotspot_time);
}
```

Appendix F Guide to Instruction-Based Sampling on AMD Family 15h Processors

Instruction-Based Sampling (IBS) is a performance monitoring technique that provides precise information about AMD64 instruction fetch behavior and about the execution of operations that are issued from AMD64 instructions. This information can be used to analyze and improve the performance of programs executing on AMD Family 15h Processors.

IBS provides four important advantages over conventional performance counter sampling:

- Hardware events are attributed precisely to the instructions that cause the events. Conventional performance counter sampling is not precise, making it difficult, if not impossible, to attribute events to specific instructions. This limits the ability to pin-point performance issues at the instruction and source code levels.
- A wide range of events are monitored and collected with each IBS sample. Either multiple sampling runs or counter multiplexing must be used to collect the same range of information with conventional performance counter sampling.
- The virtual and physical addresses of load/store operands are collected. Profiling tools can use this information to associate specific data structures with the x86 instructions performing load/store operations.
- Latency is measured for key performance parameters such as data cache miss latency.

The precision afforded by IBS also enables automated optimization techniques (e.g., profile-directed optimization) which require detailed, precise information about instruction-level program behavior.

F.1 Background

Some familiarity with the microarchitecture of AMD Family 15h processors is required to understand how instruction-based sampling works and to interpret the data produced by IBS. Important information on the microarchitecture of AMD Family 15h Processors can be found in “Microarchitecture of AMD Family 15h Processors” on page 29 of this volume. This section summarizes a few important points.

The *BIOS and Kernel Developers Guide (BKDG) for AMD Family 15h Processors*, order# 42301, provides many specific details about IBS (events, model specific registers, etc.) This appendix is intended to complement the information in the BKDG. The BKDG should be regarded as the definitive resource about IBS features.

F.2 Overview

The successive pipeline phases of fetch, decode, dispatch and execution decouple the fetching of AMD64 instructions from their eventual conversion and execution as macro-ops. The separation between fetch and execution is reflected in the IBS mechanism. Instruction-based sampling consists of two parts:

- **Fetch sampling**—collects and reports performance data on AMD64 instruction fetch behavior. Fetch sampling provides information about instruction TLB and instruction cache behavior for fetched instruction bytes.
- **Op sampling**—collects and reports performance data on the execution of instruction operations (ops). Op sampling produces retirement cycle counts that are common to all sampled operations and execution-related data that are specific to the kind of operation (e.g., branch or return) that was sampled, including load and store operations.

A *fetch* is an access to the instruction cache that results in bytes being delivered to the decode instruction buffer and can contain multiple instructions.

Op sampling records and reports the address of the AMD64 instruction from which the op was generated and issued. This allows profilers and other supporting software tools to associate op performance data with a “parent” AMD64 instruction.

Fetch sampling and op sampling are independent and may be separately enabled. Fetch and op sampling may also be enabled at the same time. When fetch and op sampling are enabled at the same time, some additional interference will result due to the larger number of samples taken, effect on pipeline behavior, cache effects, etc. This behavior is a natural consequence of sampling. Fetch and op sampling each have their own model specific registers (MSRs) to control sampling and to report results.

The same overall process is used to take a fetch sample or an op sample. Generically, an IBS sample is taken in the following way:

1. Software loads a maximum instruction fetch (or op) selection count into the appropriate IBS control MSR.
2. Software enables IBS mode in the appropriate control MSR.
3. The periodic selection counter is automatically incremented by the hardware. For fetch sampling, the periodic fetch counter is incremented for each completed fetch. For op sampling, the periodic op counter is incremented each processor cycle (cycles-based selection mode) or the periodic op counter is incremented each time an op is dispatched (dispatched op-based selection.)
4. When the selection counter reaches the maximum selection count, an instruction fetch (or op) is selected and the instruction fetch (or op) is tagged.
5. As the tagged instruction fetch (or op) is processed by the hardware, events that occur due to the tagged instruction fetch (or tagged op) are recorded by the hardware (e.g., did it cause a cache miss, branch mispredict, etc.).

6. If the tagged instruction fetch (or tagged op) finishes, an interrupt is raised and all of the collected information is passed to an interrupt service routine (ISR) through the IBS MSRs.
7. The ISR saves the information producing an IBS sample. After collecting the sample, the ISR clears the current count and goes to step 2.

Software using IBS needs to take a few differences between fetch and op sampling into account in order to handle both kinds of performance data. The differences between fetch and op sampling are described in the following sections.

F.3 IBS fetch sampling

This section describes IBS fetch sampling in more depth. IBS fetch sampling captures information about the process of fetching instruction bytes.

An “attempted fetch” is a request to load instruction bytes from a specific virtual memory address. (The term “linear” is sometimes used in place of “virtual” when referring to virtual addresses.) A “completed fetch” is an attempted fetch (a request) that eventually delivers instruction bytes to the decoders. An “aborted fetch” is a request that does not complete.

F.3.1 Taking an IBS fetch sample

IBS fetch sampling is controlled by values that are configured into specific IBS MSRs (see the BKDG for details). The periodic fetch counter and the maximum fetch count value control the sampling process. The IBS fetch sampling mechanism counts completed fetches in the periodic fetch counter in order to select the next attempted fetch to tag and monitor. When the periodic fetch counter reaches the maximum fetch count, the next attempted fetch is tagged and monitored. Thus, the maximum fetch select count determines how often an attempted fetch is selected and tagged. This quantity is often called the “sampling period.” An interrupt is generated when the tagged fetch completes or aborts. IBS information is reported for both completed and aborted fetches.

Both the periodic fetch counter and the maximum fetch count are 20-bit values. Software can set the high order 16 bits of the maximum fetch count; the low order 4 bits are always set to zero. The periodic fetch counter is reset when the maximum fetch count is reached. Under software control, the low order 4 bits of the periodic fetch counter can be set to a pseudo-random 4-bit value. Randomization helps to prevent the sampling process from “syncing up” with tight loops that are being executed—a periodicity effect that could affect the accuracy of performance statistics. Software may also choose to randomize the maximum fetch count by generating its own randomized maximum fetch count.

IBS fetch sampling data are returned in MSRs that are read by the interrupt service routine. The following information about the tagged fetch is returned:

- The virtual fetch address (always valid),
- The corresponding physical fetch address (valid when the MSR flag `IbsPhysAddrValid` is set),

- Event flags indicating completion, instruction cache (IC) miss, L1 instruction translation lookaside buffer (ITLB) miss or L2 ITLB miss,
- Size of the page translation in the L1 ITLB (a 2-bit field that is valid when the MSR flag `IbsPhysAddrValid` is set), and
- The instruction fetch latency in cycles.

Software may augment the IBS fetch data with other information such as a timestamp, the process identifier, etc. The virtual fetch address directly relates the information in the sample to the fetched instruction bytes, which cannot be done with counter sampling because conventional counter sampling records the interrupt restart IP instead of the instruction that actually caused the events. Out-of-order execution further complicates the attribution of an IP to events since many instructions may be in-flight. The IBS virtual and physical fetch addresses precisely identify the instruction that caused the events reported by the hardware—the key advantage of IBS.

Table 15 table summarizes the hardware event flags and values that are available for analysis.

Table 15. IBS Hardware Event Flags

Flag/Field	Purpose/meaning
<code>IbsPhyAddrValid</code>	Physical address and page size are valid
<code>IbsL1TlbMiss</code>	Fetch initially missed in L1 ITLB
<code>IbsL2TlbMiss</code>	Fetch initially missed in L2 ITLB
<code>IbsL1TlbPgSz</code>	Size of page translation in L1 ITLB
<code>IbsIcMiss</code>	Fetch initially missed in the instruction cache
<code>IbsFetchComp</code>	Fetch completed
<code>IbsFetchLat</code>	Cycles from fetch initiated to completed/aborted

The flag/field names used in this table correspond to the flags and fields in the IBS model-specific registers. Please see the *BIOS and Kernel Developer's Guide for AMD Family 15h Processors*, order# 42301, for the location, size and position of these flags and fields.

The IBS fetch latency value (`IbsFetchLat`) reports the number of cycles between the instant the fetch was initiated and the instant an instruction was delivered to the decoder (completion) or until the instant when the fetch was aborted. The translation page size (`IbsL1TlbPgSz`) is a 2-bit field that indicates the size of the page that was used during virtual to physical address translation as performed in the L1 ITLB.

F.3.2 Interpreting IBS fetch data

Before discussing the interpretation of IBS fetch data, it helps to have a little background on the process of fetching instruction bytes.

An attempted fetch must first be processed by the instruction translation lookaside buffers (ITLB) to convert the virtual address to a physical address. AMD Family 15h Processors use a two-level ITLB structure consisting of a level 1 (L1) ITLB and a level 2 (L2) ITLB. If page translation information is

found in the L1 ITLB, the virtual address is translated using that information. If translation information is not in the L1 ITLB (an L1 miss) and it is available in the L2 ITLB, the L1 ITLB is reloaded with the information from the L2 ITLB and address translation completes. If neither ITLB contains the page translation information (an L1 ITLB miss and an L2 ITLB miss), then the information is loaded from memory-resident page tables.

Once the physical address is available, the instruction cache (IC) is accessed using the physical address. If the instruction bytes for that address are present in the IC, it is returned and is delivered to the decoders. If the instruction bytes are not present in the IC (an IC miss), it must be obtained from either the L2 cache, L3 cache, or system memory. In any case, the instruction bytes are delivered to the decoders once it is available.

As stated earlier, an attempted fetch is said to complete when its instruction bytes are delivered to the decoders. An attempted fetch that did not complete is an aborted fetch. An attempted fetch may abort at any point in the process of fetching instruction bytes. A fetch may abort due to a control transfer misprediction from an earlier fetch.

Instruction fetch is a highly speculative activity. Some completed fetches could also be on the wrong path, but the redirection does not arrive until after the instruction has left the IC. Thus, some completed fetches are speculative and the corresponding instructions may not be executed or retired. The hardware event flags and values in the previous table are made available to software when an IBS fetch sampling interrupt is generated. The miss event flags indicate whether the attempted fetch initially missed the L1 ITLB, L2 ITLB, or IC. Address translation and IC access may eventually succeed. These flags show the hardware condition on the first attempt at translation/access. Eight combinations of the event flags, as shown in the table below, are produced by the hardware. In Table 16, "TlbMiss" means "IbsL1TlbMiss or IbsL2TlbMiss," the logical OR of the IBS event flags IbsL1TlbMiss and IbsL2TlbMiss.

Table 16. Event Flag Combinations

TlbMiss	IbsPhyAddrValid	IbsFetchComp	IbsIcMiss	Interpetation
0	0	0	0	Killed by redirect before ITLB/IC access
0	1	0	1	L1 ITLB hit, IC miss, likely redirect during IC fill
0	1	1	0	L1 ITLB hit, IC hit
0	1	1	1	L1 ITLB hit, IC miss
1	0	0	0	ITLB miss, likely redirect during L1 ITLB reload
1	1	0	1	ITLB miss, IC miss, likely redirect during IC fill
1	1	1	0	ITLB miss, IC hit
1	1	1	1	ITLB miss, IC miss

The first case, in which IbsL1TlbMiss, IbsL2TlbMiss, IbsPhyAddrValid, IbsFetchComp, and IbsIcMiss are all clear, does not provide any useful information since the attempted fetch is killed very early before ITLB or IC access. We refer to such fetches as “killed fetches.”

The other three cases in which a fetch does not complete are also likely due to a redirection. These four cases are the result of incorrect branch speculation. While completed fetches have “full” information, the IC and ITLB information for aborted fetches is just as important. Instruction cache and TLB accesses can be both constructively and destructively influenced by earlier wrong path accesses.

F.4 IBS op sampling

This section describes IBS op sampling in more detail. IBS op sampling provides information about the execution of ops.

F.4.1 Taking an IBS op sample

IBS op sampling is controlled by values that are configured into specific IBS MSRs. (See the BKDG for details.) The periodic op selection counter and the maximum count value control the sampling process. The current periodic op selection count is maintained in an 27-bit counter. Software can read and write this counter in order to save and restore its current value, and to preset the count.

A 27-bit maximum op selection count value determines how often a micro-op is sampled (the sampling period.) The high order 23 bits of the maximum op selection count are configured by software. The internal, low order 4 bits of the maximum op selection count are always zero.

IBS provides two op selection and tagging modes: cycles-based selection and dispatch op-based selection. The mode is selected by software through the IBS op control MSR.

- Cycles-based selection. IBS counts processor cycles in order to select and tag an op for sampling. The current periodic op selection count is incremented each processor cycle. When the current count reaches the maximum op selection value, a one-of-four round-robin counter selects an op in the next dispatch line. If the op selected by the round-robin counter is invalid, the next younger op is tagged.
- Dispatched op-based selection. IBS counts dispatched ops in order to select and tag an op for sampling. When the current periodic op selection count reaches the maximum op selection value, the op is tagged.

In both modes, an interrupt is generated when a tagged op is retired. The IBS event and latency values are then read from the MSRs by the interrupt service routine. The interrupt service routine may combine the IBS op data with other information (such as a timestamp and process ID) forming a complete software sample to be saved for post-processing.

Software must randomize the sampling period. Coincident periodicity may occur between the sampling process and the workload. Periodicity affects op selection, producing profiles that exhibit aliasing effects (over- or under-sampling of certain ops/instructions). Randomization of the low order 4 bits is especially important. This must be accomplished by writing a pseudo-random value to the periodic op selection counter since the low order 4 bits of the maximum op selection count are always zero.

Ops may stall in the pipeline stage in which they are tagged. When cycles-based selection is configured, a stalled op is more likely to be tagged. This behavior affects the statistical distribution of ops in the resulting program profile. Dispatched op-based selection is the preferred mode because the statistical distribution is not affected by stalls and the distribution of samples more accurately reflects op and instruction execution frequency.

IBS only returns data for tagged ops that retire. However, a tagged op may be flushed before retirement. In this case, IBS data for the flushed op is discarded (i.e., the sample is dropped.). The number of dropped samples due to a flushed tagged op can be counted by a performance monitoring event (see the BKDG). After a tagged op is flushed, the current count is set to a pseudo-randomized value and a new op is tagged when the current count again reaches the maximum count value.

F.4.2 Interpreting IBS op data

An op can be classified into one of several broad categories according to the major operation which it performs: arithmetic, logic, shift, etc. In general, IBS treats ops as *undifferentiated*, that is, the category or function is not explicitly identified by IBS. However, two categories of ops are explicitly identified: *branch* and *resync*. A branch op implements AMD64 branch semantics and includes unconditional jumps, conditional jumps, subroutine call and return. One subtype of branch op is also explicitly identified by IBS: *return*. A return op implements AMD64 return semantics. A resync op is only found in certain VectorPath instructions and causes a complete pipeline flush. Branch, return and resync are explicitly identified since interesting information about program control flow can be obtained by monitoring their behavior.

In addition to performing a major function (such as arithmetic, branch, etc.), an op may initiate a memory read, memory write, or a read and write to the same memory address. IBS explicitly identifies those ops which perform a *load* (memory read) and/or *store* (memory write) operation. When interpreting IBS data, please note:

- Some ops can perform a “load-operate-store” sequence to the same address and are identified by IBS as performing both a load *and* a store operation.
- Some branch ops perform a load operation and will be identified by IBS as performing a load.

The exact type of the sampled op is specified by one or more bits in the IBS MSRs that return sample information:

- The IbsOpBrnRet (where the “Ret” suffix stands for “retired”) and IbsOpReturn bits in the IbsOpData MSR indicate whether the op was a branch or return.
- The IbsOpBrnResync bit in the IbsOpData MSR indicates whether the op was a resync.

If none of these bits are set, the op is undifferentiated. Undifferentiated ops are still important as they provide information about program execution. The IbsLdOp and IbsStOp bits in the IbsOpData3 MSR indicate whether the op performed a load operation or a store operation, respectively.

Three values are reported for all ops:

- The virtual address of the parent AMD64 instruction from which the tagged op was issued (IbsOpRip), which is valid if the IbsRipInvalid bit is clear,
- The tag-to-retire count in cycles (IbsTagToRetCtr), and
- The completion-to-retire count in cycles (IbsCompToRetCtr).

These values are returned in model-specific registers.

The virtual address of the parent AMD64 instruction can be used to associate the IBS op sample with the AMD64 instruction from which the op was issued. More than one op may be issued from a single AMD64 instruction. All such ops have the same virtual instruction address (RIP) as the parent AMD64 instruction.

Tag-to-retire count and completion-to-retire count are retire-related cycle counts.

Tag-to-retire Count

The tag-to-retire count is the number of cycles between the instant the op is tagged to the instant the op is retired. The op is tagged when it leaves the decode unit.

Instructions can stall after they are decoded due to a lack of resources, such as reservation station entries. These cycles are included in the tag-to-retire count.

The tag-to-retire time includes the time spent waiting for operands, time spent waiting to issue after operands are available, time spent in an execution unit, and time spent waiting for all the younger ops in the scheduling window to retire.

Completion-to-retire count

The completion-to-retire count is the number of cycles between the instant the op completed and the instant the op was retired. An operation is complete when it has finished execution. The completion-to-retire count indicates how long retirement was delayed after completion.

The difference between the completion-to-retire count and the tag-to-retire count is the number of cycles that occur between tagging and completion.

F.4.3 Interpreting IBS branch/return/resync op data

Information about branch, return and resync ops are reported in the the IbsOpData MSR. The event flags and counts returned by the IbsOpData MSR are summarized in Table 17.

Table 17. IbsOpData MSR Event Flags and Counts

Flag/field	Purpose/meaning
IbsOpBrnRet	Op was a retired branch
IbsOpBrnMisp	Op was a branch that mispredicted
IbsOpBrnTaken	Op was a branch that was taken
IbsOpReturn	Op was a return
IbsOpMispReturn	Op was a return that mispredicted
IbsOpBrnResync	Op was a resync
IbsTagToRetCtr	Cycles from op tagging to retirement
IbsCompToRetCtr	Cycles from op completion to retirement

As noted earlier, the IbsTagToRetCtr and IbsCompToRetCtr fields are valid for all op samples, not just branch, return and resync ops.

A branch operation is a change in program control flow (or micro-code control flow for IbsOpBrnMisp and IbsOpBrnTaken). Information is reported only for retired branches since IBS data is only reported for retired ops. (Information for flushed ops is not reported.) Mispredicted branches retire then kill all younger ops after them and redirect the front end of the pipeline.

The IbsOpBrnRet flag indicates whether the tagged op was an operation with AMD64 branch semantics (set) or not (clear.) If the IbsOpBrnFlag is set then the IbsOpBrnMisp and IbsOpBrnTaken flags indicate the execution status of the op as shown in Table 18.

Table 18. Execution Status Indicated by IbsOpBrnMisp and IbsOpBrnTaken Flags

IbsOpBrnMisp	IbsOpBrnTaken	Execution status
0	0	Was not mispredicted and was not taken
0	1	Was not mispredicted and was taken
1	0	Was mispredicted and was not taken
1	1	Was mispredicted and was taken

The IbsOpBrnMisp and IbsOpBrnTaken bits can be viewed as a property of the tagged branch op.

The IbsOpReturn and IbsOpMispReturn flags can also be considered to be properties of the tagged AMD64 or VectorPath branch op. The IbsOpReturn flag is set when the tagged op was, specifically, a return op. The IbsOpMispReturn flag indicates whether the return op was mispredicted or not. A resync op is not predicted and is always taken.

Table 19. Execution Status Indicated by IbsOpReturn and IbsOpMispReturn Flags

IbsOpReturn	IbsOpMispReturn	Execution status
0	N/A	Branch op was not a return
1	0	Was a correctly predicted return op
1	1	Was a mispredicted return op

IbsOpBrnMisp and IbsOpBrnTaken are valid for all branch ops including micro-code ops and return ops. The IbsOpReturn and IbsOpMispReturn flags merely provide additional information for return ops.

IBS reports the branch target address in a separate MSR. The branch target address is valid if it is non-zero. The target address may be used to build a dynamic control graph for frequently executed code, including control edges that cannot be determined through static analysis (for example, edges due to indirect jumps).

F.4.4 Interpreting IBS Load/Store Data

If the sampled op accesses memory, information about the load and/or store operation is returned in four model specific registers:

- The IbsOpData2 and IbsOpData3 registers contain event flag and latency information accumulated in the Northbridge and load/store unit, respectively.
- The IbsDcLinAd and IbsDcPhysAd registers contain the virtual (linear) and physical address of the memory operand, i.e., the address of the memory location read and/or written.

The address of the memory operand can be used to associate the load or store operation with a data structure in memory. The virtual and physical addresses are valid when the IbsDcLinAddrValid and IbsDcPhyAddrValid bits are set, respectively, in the IbsOpData3 register. Some ops that are issued from VectorPath instructions use a physical address directly. Thus, it is possible to have an IBS sample with a valid physical address and an *invalid* virtual address.

The flags and fields in the IbsOpData3 MSR provide basic information about any memory access initiated by the sampled op. The IbsLdOp and IbsStOp fields indicate whether a load and/or store were initiated by the sampled op. If a load operation initially misses in the data cache (as indicated by IbsDcMiss), the IbsDcMissLat field returns the number of clock cycles from when the miss was detected until data was delivered to the core. This field is not valid for store operations.

IbsDcStBnkCon and IbsDcLdBnkCon fields are set when a memory op cannot access the cache due to a bank conflict, resulting in a delay of the op.

The data cache miss latency (IbsDcMissLat Dcache) is only valid for *loads* that miss in the data cache. The timed latency interval for the IbsDcMissLat is calculated from the data cache miss to data cache write.

Table 20 summarizes the information in the IbsOpData3 register.

Table 20. IbsOpData3 Register Information

Flag/field	Purpose/Meaning
IbsLdOp	Tagged op initiated a load operation
IbsStOp	Tagged op initiated a store operation
IbsDcL1tlbMiss	Translation info not initially present in L1 DTLB
IbsDcL2tlbMiss	Translation info not initially present in L2 DTLB

Table 20. lbsOpData3 Register Information

lbsDcL1tlbHit2M	Translation info was eventually present in a 2M page entry in L1 DTLB
lbsDcL1tlbHit1G	Translation info was eventually present in a 1G page entry in L1 DTLB
lbsDcL2tlbHit2M	Translation info was initially present in a 2M page entry in L2 DTLB
lbsDcMiss	Load/store initially missed in the data cache
lbsDcMisAcc	Load/store crossed a 128-bit address boundary (misaligned)
lbsDcLdBnkCon	Load/store had a bank conflict with a load.
lbsDcStToLdFwd	Data was forwarded from a store to the tagged load
lbsDcStToLdCan	Forwarding from a store to the tagged load was cancelled
lbsDcUcMemAcc	Load/store accessed uncacheable memory
lbsDcWcMemAcc	Load/store accessed write combining memory
lbsDcLockedOp	Load/store was a locked operation
lbsDcMabHit	Load/store hit on an allocated MAB entry
lbsDcLinAddrValid	Virtual (linear) address valid
lbsDcPhyAddrValid	Physical address valid
lbsDcMissLat	Data cache miss latency (load only)

F.4.5 Interpreting IBS load/store Northbridge data

The memory hierarchy in the AMD Family 15h Processor consists of a shared L1 instruction cache, a dedicated cluster L1 data cache, a shared core-pair L2 cache, a shared L3 cache and system memory. The L3 cache is shared among the cores within a multi-core processor. System memory is supported by a non-uniform memory access (NUMA) architecture in which some portion of physical memory is local to the processor while the remaining portions of physical memory are remote. Access to remote memory is implemented through the AMD Direct Connect Architecture via coherent HyperTransport™ links.

If a core cannot satisfy a load or store operation from L1 data cache or L2 cache, it communicates a request to the Northbridge through its System Request Interface. The Northbridge is shared across cores. The Northbridge consists of:

- A System Request Interface (SRI) to each core,
- A shared L3 cache (if present),
- A memory controller (MCT) to handle communication with local memory,
- One or more HyperTransport (HT) link interfaces, and
- A crossbar (XBAR) to handle communication between the SRI, MCT and HT links.

Further, the Northbridge performs address space routing. There are four main types of address space routing: system memory (DRAM), Memory-mapped IO (MMIO), IO space, and configuration space. (See the BIOS and Kernel Developer's Guide for AMD Family 15h Processors, order# 42301, for

more information about this and other aspects of the Northbridge.) The Northbridge also handles communication between a core and its local Advanced Programmable Interrupt Controller (APIC.)

When the Northbridge receives a request for data through the SRI, the data will be retrieved from one of several data sources depending upon the physical location of the data and possibly its coherency state. Relative to the processor making the request, data may be returned from:

- L1 or L2 cache local to the core(s) making the request
- Local L3 cache
- Remote L1/L2/L3 cache (after traversing a coherent HT link)
- Local system memory (via the local MCT)
- Remote system memory (after traversing a coherent HT link and via the remote MCT)
- Local MMIO, configuration space, or APIC
- Remote MMIO or configuration space

Information about Northbridge activity is gathered and returned when IBS op sampling is enabled and a load operation misses in both the L1 data cache and the L2 cache. The IbsOpData2 register returns information from the Northbridge. Data in this register is valid when a load misses in both the L1 data cache and the L2 cache. The fields in the IbsOpData2 register are summarized in Table 21.

Table 21. IbsOpData2 Register Fields

Flag/field	Purpose/Meaning
NbIbsReqCacheHitSt	Modified state (0), Owned state (1)
NbIbsReqDstProc	Request serviced by local (0) or remote (1) memory
NbIbsReqSrc	Data source (see Table 22)

It is important to emphasize that Northbridge data are only valid for load operations. Store operations may retire before they read data into the local cache. Thus, a subset of IBS information is either invalid or unreliable for store operations because a store operation may have retired and caused a sampling interrupt before store-related Northbridge events even occur. This behavior affects the validity of IBS data cache miss latency, which is valid only for load operations. Software developers should filter out NB and IBS data cache miss latency for store operations and report data only for load operations.

The NbIbsReqDstProc bit indicates whether the request was serviced locally or by a remote processor. Local service is typically faster. The NbIbsReqSrc field indicates the data source which satisfied the request, as described in Table 22. The NbIbsReqCacheHitSt indicates the cache state (modified or owned) when the data source type is “Cache.”

Table 22. Northbridge Request Data Source Field

NblbsReqSrc	Northbridge Request Data Source
0x0	No valid status
0x1	Data returned from local L3 cache
0x2	Data returned from local CPU cache in another core or remote L1/L2/L3 cache
0x3	Data returned from DRAM
0x4	Reserved
0x5	Reserved
0x6	Reserved
0x7	Data returned from MMIO/configuration space/PCI/APIC

IBS Northbridge event data may be interpreted according to Table 23.

Table 23. IBS Northbridge Event Data

NblbsReqSrc	NblbsReqDstProc	Meaning
0	0	Northbridge data is invalid
0	1	Northbridge data is invalid
1	0	Request served from local L3 cache
1	1	N/A
2	0	Request served from L1 or L2 of a local core
2	1	Request served from L1 or L2 of a remote core or a remote L3 cache
3	0	Request served from local DRAM
3	1	Request served from remote DRAM
7	0	Request served from local MMIO/Config/PCI/APIC
7	1	Request served from remote MMIO/Config/PCI

F.5 Software-based analysis

The approach to event computation and reporting described in this section is based upon concepts and methods that are used to compute and report results obtained through conventional performance counter sampling. This approach has two main advantages:

- Software developers and other end users are already familiar with performance counter sampling and the hardware events that are measured and reported.
- Profilers and other software tools can exploit existing data structures and methods to post-process IBS sample data and to correlate events with instructions, source lines, functions, modules, threads and processes.

The main disadvantage is that the full flexibility of analysis afforded by IBS data is not fully realized. Unfortunately, description of alternative approaches is beyond the scope of this paper.

F.5.1 Derived events and post-processing

The method described here converts IBS sample data into a set of derived event counts. A *derived event* is a specific, useful hardware condition that can be determined through a combination of one or more IBS event flags or values (such as the translation page size, tag-to-retire count, etc.) Example derived events include IBS instruction cache miss, IBS mispredicted branch and IBS data cache miss.

Caution: *End users should be discouraged from making direct comparison between IBS derived events and the performance counter events with the same or similar names because the sampling method and populations are different. Performance counter sampling is triggered by an event configured for a counter. IBS fetch sampling is triggered by completed fetches and IBS op sampling is triggered by processor cycles. Also, PMC execution events may be triggered by any op while IBS op events are counted only for retired ops.*

The process of converting IBS sample data to derived event counts is straightforward. An IBS fetch sample may be represented as shown in Table 24.

Table 24. An IBS Fetch Sample

Virtual fetch address	lbslcMiss	lbsL1TlbMiss	lbsL2TlbMiss	lbsComp	...
0x04000020	1	0		1	

Each column represents a hardware flag or value returned with an IBS sample. The virtual fetch address can be correlated back to an instruction in a software process using the same well-known techniques employed in conventional performance counter sampling. To visualize the process of computing derived events, consider the ten IBS fetch samples arranged as a 2-D table, where each row is a sample:

Table 25. 2-D Table of IBS Fetch Samples

Virtual fetch address	lbslcMiss	lbsL1TlbMiss	lbsL2TlbMiss	lbsComp	...
0x04000020	1	1	0	1	
0x04000040	0	0	0	1	
0x04000020	0	0	0	1	
0x04000040	0	0	0	1	
0x04000020	0	0	0	0	
0x04000040	0	0	0	1	
0x04000040	0	0	0	1	
0x04000020	0	0	0	1	
0x04000044	0	0	0	1	
0x04000080	1	0	0	1	

To compute the three derived events,

- IBS instruction cache miss,
- IBS L1 ITLB miss, L2 ITLB hit,
- IBS fetch completed,

scan sequentially through the table from top to bottom and count the number of occurrences of the hardware conditions associated with the three events. (In practice, post-processing software must keep a running count for all derived events.) In the case of IBS instruction cache miss, for example, there are two samples with the `IbsIcMiss = 1` condition, so the total count reported for this event is two. The number of IBS fetch completed events is nine.

The IBS L1 ITLB miss, L2 ITLB hit event requires the use of a slightly more complicated condition, (`IbsL1TlbMiss & ~IbsL2TlbMiss`), but the counting procedure is the same. The count reported for this event is one, since only one IBS fetch sample satisfies the condition.

Derived events can be placed into a histogram to obtain a program profile. Figure 18 shows the histogram for the IBS fetch completed derived event.




Figure 18. Histogram for the IBS Fetch Completed Derived Event

The following sections define derived events for IBS fetch and op data.

F.5.2 Derived events for IBS fetch data

Performance analysis tools (e.g., a profiler) can combine event flags to derive new events. Table 26 illustrates the kinds of events that can be derived from the basic flags and fields provided in an IBS fetch sample.

Table 26. New Events Derived from Combined Event Flags

ID	Name	Derivation
F000	IBS fetch samples	Number of all IBS fetch samples
F001	IBS fetch killed	Number of killed IBS fetch samples
F002	IBS fetch attempted	Number of non-killed IBS fetch samples

Table 26. New Events Derived from Combined Event Flags (Continued)

F003	IBS fetch completed	IbsFetchComp
F004	IBS fetch aborted	~IbsFetchComp
F005	IBS L1 ITLB hit	~IbsL1TlbMiss & IbsPhyAddrValid
F006	IBS L1 ITLB miss, L2 ITLB hit	IbsL1TlbMiss & ~IbsL2TlbMiss
F007	IBS L1 ITLB miss, L2 ITLB miss	IbsL1TlbMiss & IbsL2TlbMiss
F008	IBS instruction cache miss	IbsIcMiss
F009	IBS instruction cache hit	IbsFetchComp && ~IbsIcMiss
F00A	IBS 4K page translation	IbsL1TlbPgSz=0 & IbsPhyAddrValid
F00B	IBS 2M page translation	IbsL1TlbPgSz=1 & IbsPhyAddrValid
F00C	IBS 1G page translation	IbsL1TlbPgSz=2 & IbsPhyAddrValid
F00D	Reserved	
F00E	IBS fetch latency	IbsfetchLat

The ID numbers in Table 26 are the event identifiers that the AMD CodeAnalyst™ profiler uses to identify a derived event.

The first five derived events break down the IBS fetch samples by five broad categories.

- **IBS fetch samples** is the number of all IBS fetch samples that were taken.
- **IBS fetch killed** is the number of all IBS fetch samples that were killed before ITLB/IC access. IbsL1TlbMiss, IbsL2TlbMiss, IbsPhyAddrValid, IbsFetchComp and IbsIcMiss are all clear.
- **IBS fetch attempted** is the number of IBS fetch samples minus the number of IBS fetch killed samples.
- **IBS fetch completed** is the number of attempted fetches that completed, i.e., delivered instruction bytes to the decoder.
- **IBS fetch aborted** is the number of attempted fetches that did not complete.

It should be noted that the notion of an attempted fetch here excludes killed fetches. Killed fetches do not provide useful analytical information and are filtered out. Killed fetches are not included in the IBS fetch aborted derived event.

For the remaining derived events, an occurrence of an event is tallied if it meets the specified condition in the Derivation column.

The IbsPhysAddrValid bit is needed to form the derived event “IBS L1 ITLB hit” because the condition (IbsL1TlbMiss=0 & IbsL2TlbMiss=0) alone is not sufficient to determine whether the attempted fetch hit in both the L1 and L2 ITLB or that the attempted fetch completed the initial ITLB access. When IbsPhysAddrValid is set, it indicates that address translation completed and produced a physical address. This bit must also be used to detect valid page translation information.

F.5.3 Derived Events for all Ops

A derived event is an event that is formed using a combination of IBS event flags and field values. The quantity of each derived event is computed and reported to end users by profiling software.

There are three derived events that are defined for all ops regardless of type:

Table 27. Derived Events for All Ops

ID	Name	Derivation
F100	IBS all op samples	Number of all IBS op samples
F101	IBS tag to retire cycles	Sum of all tag to retire cycles
F102	ibs completion to retire cycles	Sum of all completion to retire cycles

The **IBS all op samples** derived event is a count of all IBS op samples taken without regard to op type (i.e., undifferentiated, branch/return/resync.) Similarly, the **IBS tag to retire cycles** and **IBS completion to retire cycles** are computed across all IBS op samples for a given IP.

F.5.4 Derived events for IBS branch/return/resync ops

The following derived events measure the behavior of branch, return and resync ops:

Table 28. Derived Events to Measure Branch, Return and Resync Ops

ID	Name	Derivation
F103	IBS branch op	lbsOpBrnRet
F104	IBS mispredicted branch op	lbsOpBrnRet & lbsOpBrnMisp
F105	IBS taken branch op	lbsOpBrnRet & lbsOpBrnTaken
F106	IBS mispredicted taken branch op	lbsOpBrnRet & lbsOpBrnTaken & lbsOpBrnMisp
F107	IBS return op	lbsOpReturn
F108	IBS mispredicted return op	lbsOpReturn & lbsOpMispReturn
F109	IBS resync op	lbsOpBrnResync

All events in Table 28 (F103-F109) are applicable only for ops that follow AMD64 branch semantics and therefore do not include micro-code branches.

F.5.5 Derived events for IBS load/store operations

Table 29 summarizes derived events for ops that perform load and/or store operations. With the exception of the first three derived events, the condition (IbsLdOp | IbsStOp) is assumed.

Table 29. Derived Events for Ops That Perform Load and/or Store Operations

ID	Name	Derivation
F200	IBS All Load/Store Ops	IbsLdOp IbsStOp
F201	IBS Load Ops	IbsLdOp
F202	IBS Store Ops	IbsStOp
F203	IBS L1 DTLB Hit	~IbsDcL1tlbMiss & IbsDcLinAddrValid
F204	IBS L1 DTLB Miss L2 DTLB Hit	IbsDcL1tlbMiss & ~IbsDcL2tlbMiss
F205	IBS L1 DTLB Miss L2 DTLB Miss	IbsDcL1tlbMiss & IbsDcL2tlbMiss
F206	IBS DC Miss	IbsDcMiss
F207	IBS DC Hit	~IbsDcMiss
F208	IBS Misaligned Access	IbsDcMisAcc
F209	IBS Bank Conflict On Load Op	IbsDcLdBnkCon
F20A	Reserved	
F20B	IBS Store to Load Forwarded	IbsDcStToLdFwd
F20C	IBSStore to Load Forwarding Cancelled	IbsDcStToLdCan
F20D	IBS UC memory access	IbsDcUcMemAcc
F20E	IBS WC memory access	IbsDcWcMemAcc
F20F	IBS locked operation	IbsDcLockedOp
F210	IBS MAB hit	IbsDcMabHit
F211	IBS L1 DTLB 4K page	~IbsDcL1tlbHit2M & ~IbsDcL1tlbHit1G & IbsDcLinAddrValid
F212	IBS L1 DTLB 2M page	IbsDcL1tlbHit2M & IbsDcLinAddrValid
F213	IBS L1 DTLB 1G page	IbsDcL1tlbHit1G & IbsDcLinAddrValid
F214	Reserved	
F215	IBS L2 DTLB 4K page	~IbsDcL2tlbMiss & IbsDcL1tlbMiss & ~IbsDcL1tlbHit2M & IbsDcLinAddrValid
F216	IBS L2 DTLB 2M page	~IbsDcL2tlbMiss & IbsDcL1tlbMiss & IbsDcL1tlbHit2M & IbsDcLinAddrValid
F217	Reserved	
F218	Reserved	
F219	IBS DC miss load latency	IbsDcMissLat when IbsLdOp & IbsDcMiss

The **IBS all load/store ops** derived event is a count of all IBS op samples that involve either a load and/or store operation. The **IBS Load Ops** and **IBS Store Ops** events break out the number of load and store operations performed by all sampled ops.

Detection of L1 and L2 DTLB miss events are more easily decoded than the similar events in IBS fetch samples, since all sampled ops are retired ops. Retired ops do not include speculative activity. All address translations must eventually complete in the L1 DTLB. Thus, the IBS translation page size flags for the L1 DTLB are always set according to the size of the completed translation. A 4K page L1 DTLB translation occurs when both the IbsDcL1tlbHit2M and IbsDcL1tlbHit1G bits are clear. The L2 DTLB translation page size is valid when IbsDcL1tlbMiss is set. A 4K page L2 DTLB

translation occurs when both the IbsDcL2tlbMiss and IbsDcL2tlbHit2M are clear. Checking the IbsDcLinAddrValid bit is necessary to be sure that an address translation was attempted.

The **IBS DC miss load latency** is only valid for load operations. Miss load latency should only be tallied by software when the IbsLdOp bit is set.

The remaining derived events are simply the number of IBS op samples for which an event bit is set (or clear), e.g., **IBS MAB hit** is tallied when the IbsDcMabHit bit is set.

F.6 Derived Events for Northbridge Activity

The IBS Northbridge derived events measure the number of local and remote accesses (without regard to data source) and measure the number of local and remote accesses by data source. The event derivations in Table 30 assume that the overall Northbridge IBS data validity condition:

$$\text{IbsLdOp} \ \& \ \text{IbsDcMiss} \ \& \ (\text{NbIbsReqSrc} = 0)$$

is true.

Table 30. IBS Northbridge Derived Events

ID	Name	Derivation
F240	IBS NB local	$\sim \text{NbIbsReqDstProc}$
F241	IBS NB remote	NbIbsReqDstProc
F242	IBS NB local L3	$\text{NbIbsReqSrc}=0x1 \ \& \ \sim \text{NbIbsReqDstProc}$
F243	IBS NB local L1/L2 (intercore)	$\text{NbIbsReqSrc}=0x2 \ \& \ \sim \text{NbIbsReqDstProc}$
F244	IBS NB remote L1/L2/L3 cache	$\text{NbIbsReqSrc}=0x2 \ \& \ \text{NbIbsReqDstProc}$
F245	IBS NB local DRAM	$\text{NbIbsReqSrc}=0x3 \ \& \ \sim \text{NbIbsReqDstProc}$
F246	IBS NB remote DRAM	$\text{NbIbsReqSrc}=0x3 \ \& \ \text{NbIbsReqDstProc}$
F247	IBS NB local other	$\text{NbIbsReqSrc}=0x7 \ \& \ \sim \text{NbIbsReqDstProc}$
F248	IBS NB remote other	$\text{NbIbsReqSrc}=0x7 \ \& \ \text{NbIbsReqDstProc}$
F249	IBS NB cache M state	$\text{NbIbsReqSrc}=0x2 \ \& \ \sim \text{NbIbsReqCacheHitSt}$
F24A	IBS NB cache O state	$\text{NbIbsReqSrc}=0x2 \ \& \ \text{NbIbsReqCacheHitSt}$
F24B	IBS NB local latency	IbsDcMissLat when $\sim \text{NbIbsReqDstProc}$
F24C	IBS NB remote latency	IbsDcMissLat when NbIbsReqDstProc

These derived events correspond to the seven kinds of Northbridge activity described in Section 4.5.

Index

A

address-generation interlocks 139
arrays 45

B

base address register 323
boolean operators 48
branches
 compound branch conditions 48
 dependent on random data 125
 optimizing density of 123
 prediction 34
 replace with computation in 3DNow! code 128

C

C language 48
 array notation versus pointers 45
 structures 63
cache
 64-byte cache line 116
caches
 probing 319
CALL and RETURN instructions 126
code padding using neutral code fillers 92
code segment (CS) base, nonzero 128
Configuration Address Register (0CF8h) 328
Configuration Data Register (0CFCh) 328
const type qualifier 56
CPUID 208

D

data cache 33, 203
data organization 203
data-parallel threading 202
device NUMA 319
device's bus-number 323
DirectPath
 DirectPath over VectorPath instructions 79
displacements, 8-bit sign-extended 92
division 143, 144, 145, 146, 160
 replace division with multiplication, integer 64, 143

F

false data sharing 204
far control-transfer instructions 130
floating-point

- division and square roots 68
- execution unit 38
- scheduler 37
- variables and expressions are type float 44

H

- HWCR 333
- HyperTransport assist 41
- HyperTransport™ 319

I

- I/O Thread Scheduling 329
- if statement 49, 58
- immediates, 8-bit sign-extended 92
- IMUL instruction 147
- index and data ports 328
- inline functions 138, 152
- inline REP string with low counts 150
- instruction
 - cache 32
 - integer execution 35
 - short encodings 84
- instruction fetch and decoding 34
- integer
 - arithmetic, 64-bit 152
 - division 64
 - execution unit 36
 - operand, consider sign 67
 - scheduler 36
 - use 32-bit data types for integer code 66
- IO-device consumption 330

L

- L2 cache controller 33
- LEA instruction 82, 90
- LEAVE instruction 89
- load/store 38, 53
- load-execute instructions 79
 - floating-point instructions 81, 82
 - integer instructions 80
- local functions 59
- local variables 65
- LOCK 208
- locked instructions 208
- LOOP instruction 130
- loops
 - generic loop hoisting 56
 - minimize pointer arithmetic 139
 - partial loop unrolling 135
 - REP string with low variable counts 151
 - unroll small loops 47
 - unrolling loops 133

M

memory
 pushing memory data 141
memory barriers 207
memory-mapped I/O (MMIO) 320
MFENCE 208, 334
MMIO base and limit address registers 323
MMX instructions
 PREFETCHNTA/T0/T1/T2 instructions 108
Move/Compute 186
MOVZX and MOVSX instructions 139
multiplication
 by constant 147
multi-threading 202
muxing constructs 129

N

non-coherent I/O 319
Northbridge 319

O

operands
 largest possible operand size, repeated string 150

P

parallelism 59
pointers
 dereferenced arguments 65
 use array-style code instead 45
population-count function 166
prefetch
 determining distance 111
 multiple 109, 110
PREFETCH and PREFETCHW instructions 105, 108
processor
 microarchitecture 30
prototypes 56

R

RDTSC 333
RDTSCP 19, 333
register reads and writes, partial 84
REP prefix 150

S

scheduling 133
serializing instruction 333
SFENCE 210
SHLD instruction 90
SHR instruction 90

SSE 167
SSE2 167
stack
 alignment considerations 119
store-to-load forwarding 52, 53, 99, 102, 103
stream processing 206
string Instructions 150
string instructions 150
subexpressions, explicitly extract common 62
superscalar processor 31
switch statement 58

T

task decomposition 202
task-parallel threading 202
time stamp counter 19, 333
timestamp counters 333
TSC 19, 333

U

unit-stride access 106, 112, 113

V

virtualized environments 333

W

WB memory 208
write combining 114, 227, 228, 229

X

XCHG 209
XOR instruction 151