

Contents

Foreword xi

Preface xiii

Questions Answered by Each Chapter xiv

Naming Conventions xv

Acknowledgements xv

Chapter 1 Introduction to the Architecture 1

What is the Intel® Internet Exchange Architecture? 2

Why the IXP2XXX Network Processors? 4

What is the IXP2XXX Product Line? 6

Microengines 7

Intel® XScale™ Microarchitecture Core 8

Media and Switch Fabric Interface 9

SRAM Interface 9

DRAM Interface 10

SHaC Unit 10

PCI Interface 11

Chassis 11

What is the Intel® IXA Portability Framework? 11

Data Plane Software Structure 12

Microblock 14

Dispatch Loop and Microblock Infrastructure Library 14

Data Plane Libraries 16

Intel® XScale™ Core Component 16

- Resource Manager 16
- Operating System Services Library 17
- Control Plane PDK 17
- Typical External System Architecture 17
- Typical Data Flow through IXP2XXX Network Processors 18

Chapter 2 Packet Processing Basics 23

- Typical Hardware Platform 24
 - Switch Fabric 24
 - Management and Control Processors 25
 - I/O Blade 26
 - Trunk Blade 27
- A Standards-based Software Architecture 28
- Data Plane Processing Stages 30
 - Packet Reception 33
 - Reassembly 34
 - Classification 35
 - Forwarding 38
 - Metering and Statistics 39
 - Congestion Management 41
 - Enqueue and Dequeue 43
 - Transmit Scheduling 45
 - Simple Priority 46
 - Round Robin 46
 - Weighted Round Robin 46
 - Deficit Round Robin 46
 - Transmit Fill 47
- Packet Buffering 47

Chapter 3 External Architecture 51

- External Architecture Overview 52
- System Configurations 54
 - IXP2800 Dual Ingress/Egress Configuration 55
 - IXP2400 Dual Ingress/Egress Architecture 58
 - IXP2800 Single Ingress/Egress Configuration 60
 - IXP2400 Single Ingress/Egress Model 62
 - Chaining Configuration 64
 - A Complete IXP2XXX System 64

- IXP2XXX External Interfaces 65
 - IXP28X0 Media and Switch Fabric Interface 66
 - IXP2400 Media and Switch Fabric Interface 67
 - MSF CSIX Flow Control 68
 - SPI-4.2 Flow Control 72
 - PCI Bus 76
 - IXP28x0 RDRAM 77
 - IXP2400 DRAM 78
 - SRAM 79
 - Slow Port 80
 - Serial Port 80
 - GPIO 80

Chapter 4 Intel® IXA Microarchitecture 81

- Parallel Processing Paradigm 83
- The Chassis 85
- Microengines 91
 - Control Store 93
 - Contexts 93
 - Data Path Registers 94
 - Execution Data Path 95
 - CRC Unit 97
 - Event Signals 97
- Media and Switch Fabric Interface 98
 - MSF Receive Processing 103
 - MSF Transmit Processing 106
- Intel® XScale™ Core 108
- Intel® XScale™ Core Peripherals 110
 - Interrupt Controller 110
 - Timers 110
- PCI Controller 111
- SRAM Unit 111
- DRAM Functional Unit 113
- The SHaC 113
 - Scratchpad Memory 114
 - Hash Unit 114
 - Control and Status Register Access Proxy 114
- IXP2850 Crypto Unit 116
 - Crypto Data Flow 118
- Summary 120

Chapter 5 Parallel Processing 121

- Why Not Use a General Purpose Processor? 122
- Multi-processing and Multithreading 124
 - Multi-processing 124
 - Multithreading 127
- The Serial Stream Processing Problem 132
- Distributed Cache 138
 - General Operation of the Distributed Cache 140
 - Folding 142
 - Interleaved Threading 148
- Summary 151

Chapter 6 Managing Data Structures 153

- Using Memory 154
- Linked List Management Structures 155
 - Enqueueing a Packet 157
 - Dequeueing a Packet 161
- SRAM Q_Array Controller 166
- Ring Buffers 176
- The “Many-to-One” Problem 180
- Summary 182

Chapter 7 Mapping Tasks to Microengines 183

- Mutual Exclusion and Ordering 184
- Hyper Task Chaining 184
 - Function Chaining 185
 - Context Pipelining 191
 - Hyper Task Chaining Diff Serv Example 193
- Pool of Threads 194
 - Functional Thread Assignment 197
 - General Thread Assignment 198
 - Packet Ordering with Pool of Threads 198
 - End-to-End Packet Ordering with Pool of Threads 199
 - Managing Large Programs 201
 - Head of Line Blocking 202
 - Determining the Size of an AISR 203
 - AISR Data Structure 205
 - Implementing AISR 207
 - Maintaining Partial Packet Order with Pool of Threads 208

Non-Blocking Packet Ordering 213
Critical Sections 218
HTC or POTs? Which Approach Should You Use? 223
Combining POTs with HTC 230

Chapter 8 Performance Considerations 233

Three Performance Variables 234
Packet Arrival Rate 236
IXP2XXX Hardware Resource Budget 238
Task Resource Utilization 242
A Real-Life Example 244
 POS Receive 246
 Layer 2 PPP Decapsulation/Classify/IPv4 247
 Queue Manager 249
 CSIX Transmit Scheduler 251
 CSIX Transmit Processing 252
 OC-48 Performance Summary 253

Glossary 257

References 269

Index 271

