

BCM53134P

Multiport Ultra Low-Power Gigabit Ethernet Switch

General Description

The Broadcom® BCM53134P is an ultra low-power, highly integrated, cost-effective smart-managed Gigabit switch. The switch design is based on the field-proven, industry-leading ROBO architecture. This device combines all the functions of a high-speed switch system including packet buffers, PHY transceivers, media access controllers (MACs), address management, port-based rate control, and a nonblocking switch fabric into a single 28 nm CMOS device. Designed to be fully compliant with the IEEE 802.3 and IEEE 802.3x specifications, including the MAC-control PAUSE frame, the BCM53134P provides compatibility with all industry-standard Ethernet, Fast Ethernet, and Gigabit Ethernet (GbE) devices.

The BCM53134P has a rich feature set suitable for not only standard GbE connectivity for broadband home gateways, desktop, and laptop PCs, but also for next-generation gaming consoles, set-top boxes, networked DVD players, and home theater receivers. It is also specifically designed for next generation SOHO/SMB routers and gateways.

The BCM53134P contains four full-duplex 10/100/1000BASE-T Ethernet transceivers. In addition, the BCM53134P has two PHY-less interfaces for the CPU or a router chip, providing flexible 10/100/1000 Mb/s connectivity. One RGMII interface can be connected to a CPU entity and configured as an IMP (In-Band Management port).

The second interface port is available in an SGMII interface supporting rates of up to 2.5 Gb/s (2.5 Gb/s is for 2500BASE-X only) that can be connected to optical modules or any other device with an SGMII interface. The BCM53134P provides 70+ on-chip MIB counters to collect receive and transmit statistics for each port.

Features

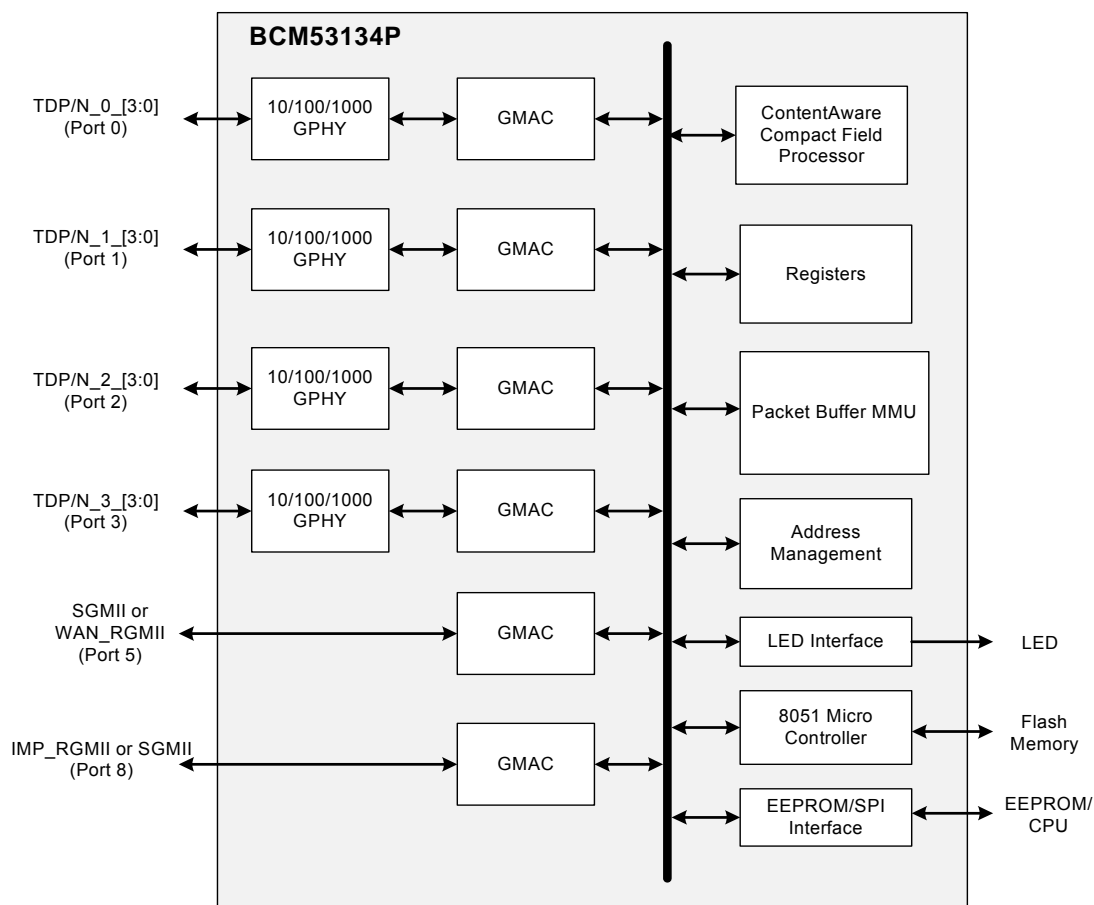
- Six 10/100/1000 media access controllers
- Four-port 10/100/1000 transceivers for TX
- One RGMII interface for an IMP for connection to a CPU/management entity without PHY
- One SGMII interface supports 1G/2.5 Gb/s (2.5 Gb/s is for 2500BASE-X only)
- Supports 2500BASE-X, 1000BASE-X, and 100BASE-FX
- IEEE 802.1p, MAC Port, TOS, and DiffServ QoS for six queues, plus two time-sensitive queues
- Port-based VLAN
- IEEE 802.1Q-based VLAN with 4K entries
- MAC-based trunking with automatic link failover
- Port-based rate control
- Port mirroring (Ingress/Egress)
- Supports IPv4 and IPv6
- Priority modification on egress
- BroadSync® HD for IEEE 802.1AS support
- Timestamp tagging at MAC interface
- Time-aware egress scheduler
- DOS attack prevention
- IGMP Snooping, MLD snooping support
- Spanning tree support (multiple spanning trees—up to eight)
- Embedded CPU (8051) processor for cable diagnostics
- CableChecker™ with unmanaged mode support
- Double tagging/QinQ
- IEEE 802.az Energy Efficient Ethernet (EEE) support
- IEEE 802.3x programmable per-port flow control and backpressure, with IEEE 802.1X support for secure user authentication
- EEPROM, MDC/MDIO, and SPI Interface.
- Serial Flash Interface for accessing embedded CPU (8051)
- 4K entry MAC address table with automatic learning and aging
- 128 KB packet buffer (1 KB = 1024 bytes)
- 128 multicast group support

Features (Continued)

- Jumbo frame support up to 9720 bytes
- 1.0V for core and 3.3V for I/O

- RGMII with option of 3.3V/2.5V or 1.8V/1.5V
- JTAG support
- 212-pin FBGA package

Figure 1: Functional Block Diagram



Note: WAN_RGMII and IMP_RGMII are either or and multiple by strap pin – "LED6_SGMII_P8_SEL".

Table of Contents

Chapter 1: Introduction	8
1.1 Overview	8
1.2 Audience	9
1.3 Data Sheet Information	9
Chapter 2: Features and Operation	10
2.1 Overview	10
2.2 Quality of Service and Scheduling	11
2.2.1 CoS Mapping	13
2.2.2 SF3 Egress Queues and Scheduler	14
2.2.2.1 Egress Transmit Queues	14
2.2.2.2 Scheduler	15
2.2.3 Scheduling	17
2.2.4 Leaky Bucket Shaper	17
2.3 Port-Based VLAN	19
2.4 IEEE 802.1Q VLAN	19
2.4.1 IEEE 802.1Q VLAN Table Organization	20
2.5 Double-Tagging	20
2.5.1 ISP Port	21
2.5.2 Customer Port	21
2.5.3 Uplink Traffic (from Customer Port to ISP)	22
2.5.4 Downlink Traffic (from ISP to Customer Port)	22
2.6 Egress VID Modification	22
2.7 Jumbo Frame Support	23
2.8 Port Trunking/Aggregation	23
2.9 WAN Port	24
2.10 Rate Control	25
2.10.1 Ingress Rate Control	25
2.10.2 Two-Bucket System	25
2.10.3 Egress Rate Control	26
2.10.4 Bucket Bit Rate	26
2.11 Protected Ports	26
2.12 Port Mirroring	26
2.12.1 Enabling Port Mirroring	27
2.12.2 Capture Port	27
2.12.3 Mirror Filtering Rules	27
2.12.3.1 Port Mask Filter	27
2.12.3.2 Packet Address Filter	27

2.12.3.3 Packet Divider Filter.....	28
2.13 IGMP Snooping	28
2.14 MLD Snooping.....	28
2.15 IEEE 802.1X Port-Based Security	28
2.16 DoS Attack Prevention	30
2.17 Compact Field Processor	31
2.17.1 Parser	32
2.17.1.1 L2 Framing Structure Parsing	33
2.17.1.1.1 VLAN Tagging Structure Parsing	33
2.17.1.1.2 Ethernet Framing Structure Parsing.....	34
2.17.1.2 L3 Framing Structure Parsing	35
2.17.1.2.1 IPv4 Header	35
2.17.1.2.2 IPv6 Header	36
2.17.1.2.3 IPv6 Extension Header.....	36
2.17.1.3 L4 Framing Structure Parsing.....	38
2.17.1.3.1 TCP Header	38
2.17.1.3.2 UDP Header	39
2.17.1.3.3 UDPLite Header	40
2.17.1.3.4 ICMP/IGMP Headers.....	40
2.17.1.4 User Defined Field (UDF) Extraction	40
2.17.1.4.1 UDF Offset Base Generation	40
2.17.1.4.2 UDF Specification.....	41
2.17.2 Search Key Composition	42
2.17.3 Action Resolution	48
2.17.3.1 Policy Action Definitions	48
2.17.3.2 Metering/Statistics Selection.....	50
2.18 Multiple Spanning Tree Protocol	50
2.19 Software Reset	51
2.20 BroadSync HD	51
2.20.1 Time Base and Slot Generation.....	51
2.20.2 Transmission Shaping and Scheduling.....	52
2.20.2.1 BroadSync HD Class5 Media Traffic	52
2.20.2.2 BroadSync HD Class4 Media Traffic	52
2.21 CableChecker	53
2.22 Egress PCP Remarking	54
2.23 Address Management.....	55
2.23.1 Address Table Organization	55
2.23.2 Address Learning.....	56
2.23.3 Address Resolution and Frame Forwarding	56
2.23.3.1 Unicast Addresses	57
2.23.3.2 Multicast Addresses.....	57
2.23.3.3 Reserved Multicast Addresses	59

2.23.4 Static Address Entries	59
2.23.5 Accessing the ARL Table Entries	60
2.23.5.1 Searching the ARL Table	60
2.23.6 Address Aging	60
2.23.6.1 Normal Aging	60
2.23.6.2 Fast Aging	60
2.24 Power Savings Modes	61
2.24.1 Auto Power-Down Mode	61
2.24.2 Energy Efficient Ethernet Mode	61
2.24.3 Deep Green Mode	62
2.25 Interrupt	62
Chapter 3: System Functional Blocks	63
3.1 Overview	63
3.2 Media Access Controller	63
3.2.1 Receive Function	63
3.2.2 Transmit Function	64
3.2.3 Flow Control	64
3.2.3.1 10/100 Mb/s Half-Duplex	64
3.2.3.2 10/100/1000 Mb/s Full-Duplex	64
3.3 Integrated 10/100/1000 PHY	64
3.3.1 Encoder	65
3.3.2 Decoder	65
3.3.3 Link Monitor	66
3.3.4 Digital Adaptive Equalizer	66
3.3.5 Echo Canceler	66
3.3.6 Crosstalk Canceler	66
3.3.7 Analog-to-Digital Converter	66
3.3.8 Clock Recovery/Generator	67
3.3.9 Baseline Wander Correction	67
3.3.10 Multimode TX Digital-to-Analog Converter	67
3.3.11 Stream Cipher	67
3.3.12 Wire Map and Pair Skew Correction	68
3.3.13 Automatic MDI Crossover	68
3.3.14 10/100BASE-TX Forced Mode Auto-MDIX	69
3.3.15 Resetting the PHY	69
3.3.16 PHY Address	69
3.3.17 Super Isolate Mode	69
3.3.18 Standby Power-Down Mode	70
3.3.19 Auto Power-Down Mode	70
3.3.20 External Loopback Mode	70

3.3.21 Full-Duplex Mode	71
3.3.21.1 Copper Mode	71
3.3.22 Master/Slave Configuration	72
3.3.23 Next Page Exchange	72
3.4 Frame Management	72
3.4.1 In-Band Management Port	72
3.4.2 Broadcom Tag Format for Egress Packet Transfer	74
3.4.3 Broadcom Tag Format for Ingress Packet Transfer	75
3.5 MIB Engine	75
3.5.1 MIB Counters Per Port	76
3.6 Integrated High-Performance Memory	82
3.7 Switch Controller	82
3.7.1 Buffer Management	82
3.7.2 Memory Arbitration	82
3.7.3 Transmit Output Port Queues	83
Chapter 4: System Interfaces	84
4.1 Overview	84
4.2 Copper Interface	84
4.2.1 Auto-Negotiation	84
4.2.2 Line-side (Remote) Loopback Mode	84
4.3 Frame Management Port Interface	84
4.3.1 RGMII Interface	85
4.4 WAN Interface	85
4.5 Configuration Pins	85
4.6 Programming Interfaces	85
4.6.1 SPI-Compatible Programming Interface	86
4.6.1.1 SS: Slave Select	86
4.6.1.2 SCK: Serial Clock	86
4.6.1.3 MOSI: Master Output Slave Input	86
4.6.1.4 MISO: Master Input Slave Output	86
4.6.1.5 External PHY Registers	88
4.6.1.6 Reading and Writing BCM53134P Registers Using SPI	89
4.6.1.7 Normal Read Operation	90
4.6.1.8 Fast Read Operation	94
4.6.1.9 Normal Write Operation	97
4.6.2 EEPROM Interface	100
4.6.2.1 EEPROM Format	100
4.6.3 Serial Flash Interface	102
4.6.4 MDC/MDIO Interface	102
4.6.4.1 MDC/MDIO Interface Register Programming	103

4.6.4.2 Pseudo-PHY	104
4.7 LED Interfaces	110
4.7.1 Dual Input Configuration/LED Output Function	114
4.8 Digital Voltage Regulator (LDO)	115
Chapter 5: Hardware Signal Definitions	116
5.1 I/O Signal Types	116
5.2 Signal Descriptions	117
Chapter 6: Pin Assignment	123
6.1 Pin List by Pin Number	123
6.2 Pin List by Pin Name	130
Chapter 7: Electrical Characteristics	137
7.1 Absolute Maximum Ratings	137
7.2 Recommended Operating Conditions	137
7.3 Electrical Characteristics	138
Chapter 8: Timing Characteristics	140
8.1 Reset and Clock Timing	140
8.2 RGMII Interface Timing	141
8.2.1 RGMII Output Timing (Normal Mode)	141
8.2.2 RGMII Output Timing (Delayed Mode)	142
8.2.3 RGMII Input Timing (Normal Mode)	143
8.2.4 RGMII Input Timing (Delayed Mode)	143
8.3 MDC/MDIO Timing	145
8.4 Serial LED Interface Timing	147
8.5 SPI Timings	147
8.6 JTAG Interface	149
8.7 EEPROM Timing	149
8.8 SGMII/SerDes Timing	151
8.9 2.5GbE/SerDes Timing	152
Chapter 9: Thermal Characteristics	153
9.1 Package Only	153
9.2 Package Only with Heat Sink (50 x 50 x 35 mm ³)	153
9.3 Package Only	154
9.4 Package Only with Heat Sink (19 x 19 x 5 mm ³)	154
Chapter 10: Mechanical Information	155
Chapter 11: Ordering Information	156

Chapter 1: Introduction

1.1 Overview

The BCM53134P is a single-chip, six-port Gigabit Ethernet (GbE) switch device. It provides the following:

- A six-port nonblocking 10/100/1000 Mb/s switch controller.
- Four ports with 10/100/1000BASE-T compatible transceivers.
- Six integrated Gigabit MACs (GMACs).
- One RGMII port for PHY-less connection to the management agent (available only in full-duplex mode).
- One SGMII 1G/2.5G (2.5 Gb/s is for 2500BASE-X only) interface for WAN management port (available only in full-duplex mode).
- An integrated Motorola SPI-compatible interface.
- High-performance, integrated packet buffer memory.
- An address resolution engine.
- A set of management information base (MIB) statistics registers.

The GMACs support full-duplex and half-duplex modes for 10 Mb/s and 100 Mb/s, and full-duplex for 1000 Mb/s. Flow control is supported in half-duplex mode with backpressure. In full-duplex mode, IEEE 802.3x frame-based flow control is supported. The GMACs are IEEE 802.3-compliant and support a maximum frame size of 9720 bytes.

The BCM53134P supports advanced ContentAware™ processing using a compact field processor (CFP). Up to four intelligent ContentAware processes are performed in parallel for every packet. This flexible engine uses TCAM-based architecture, which allows wildcard capabilities. Action examples include dropping, changing the forward port map, adding forward port, assigning the priority of a frame, and so on. These advanced ContentAware processes are well suited for access control lists (ACLs) and DoS prevention.

An integrated address management engine provides address learning and recognition functions at maximum frame rates. The address table provides capacity for learning up to 4K unicast addresses. Addresses are added to the table after receiving an error-free packet.

The MIB statistics registers collect receive and transmit statistics for each port and provide direct hardware support for the Ether-like MIB, MIB II (interfaces), and the first four groups of the RMON MIB. All nine groups of RMON can be supported by using additional capabilities, such as port mirroring/snooping, together with an external microcontroller to process some MIB attributes. The MIB registers can be accessed through the Serial Peripheral Interface Port by an external microcontroller.

1.2 Audience

This document is for designers interested in integrating the BCM53134P switches into their hardware designs and for others who need specific data about the physical characteristics and operation of the BCM53134P switches.

1.3 Data Sheet Information

The following notational conventions are used in this document:

- Signal names are shown in uppercase letters (such as DATA).
- A bar over a signal name indicates that it is active low (such as $\overline{\text{CE}}$).
- In register and signal descriptions, [n:m] indicates a range from bit n to bit m (such as [7:0] indicates bits 7 through 0, inclusive).
- The use of R or Reserved indicates that a bit or a field is reserved by Broadcom for future use. Typically, R is used for individual bits and Reserved is used for fields.
- Numerical modifiers such as K or M follow traditional usage (for example, 1 KB means 1,024 bytes, 100 Mb/s [referring to fast Ethernet speed] means 100,000,000 b/s, and 133 MHz means 133,000,000 Hz).

Chapter 2: Features and Operation

2.1 Overview

The BCM53134P switches include the following features:

- “Quality of Service and Scheduling” on page 11
- “Port-Based VLAN” on page 19
- “IEEE 802.1Q VLAN” on page 19
- “Double-Tagging” on page 20
- “Jumbo Frame Support” on page 23
- “Port Trunking/Aggregation” on page 23
- “WAN Port” on page 24
- “Rate Control” on page 25
- “Protected Ports” on page 26
- “Port Mirroring” on page 26
- “IGMP Snooping” on page 28
- “MLD Snooping” on page 28
- “IEEE 802.1X Port-Based Security” on page 28
- “DoS Attack Prevention” on page 30
- “Compact Field Processor” on page 31
- “Multiple Spanning Tree Protocol” on page 50
- “Software Reset” on page 51
- “BroadSync HD” on page 51
- “CableChecker” on page 53
- “Egress PCP Remarking” on page 54
- “Address Management” on page 55
- “Power Savings Modes” on page 61

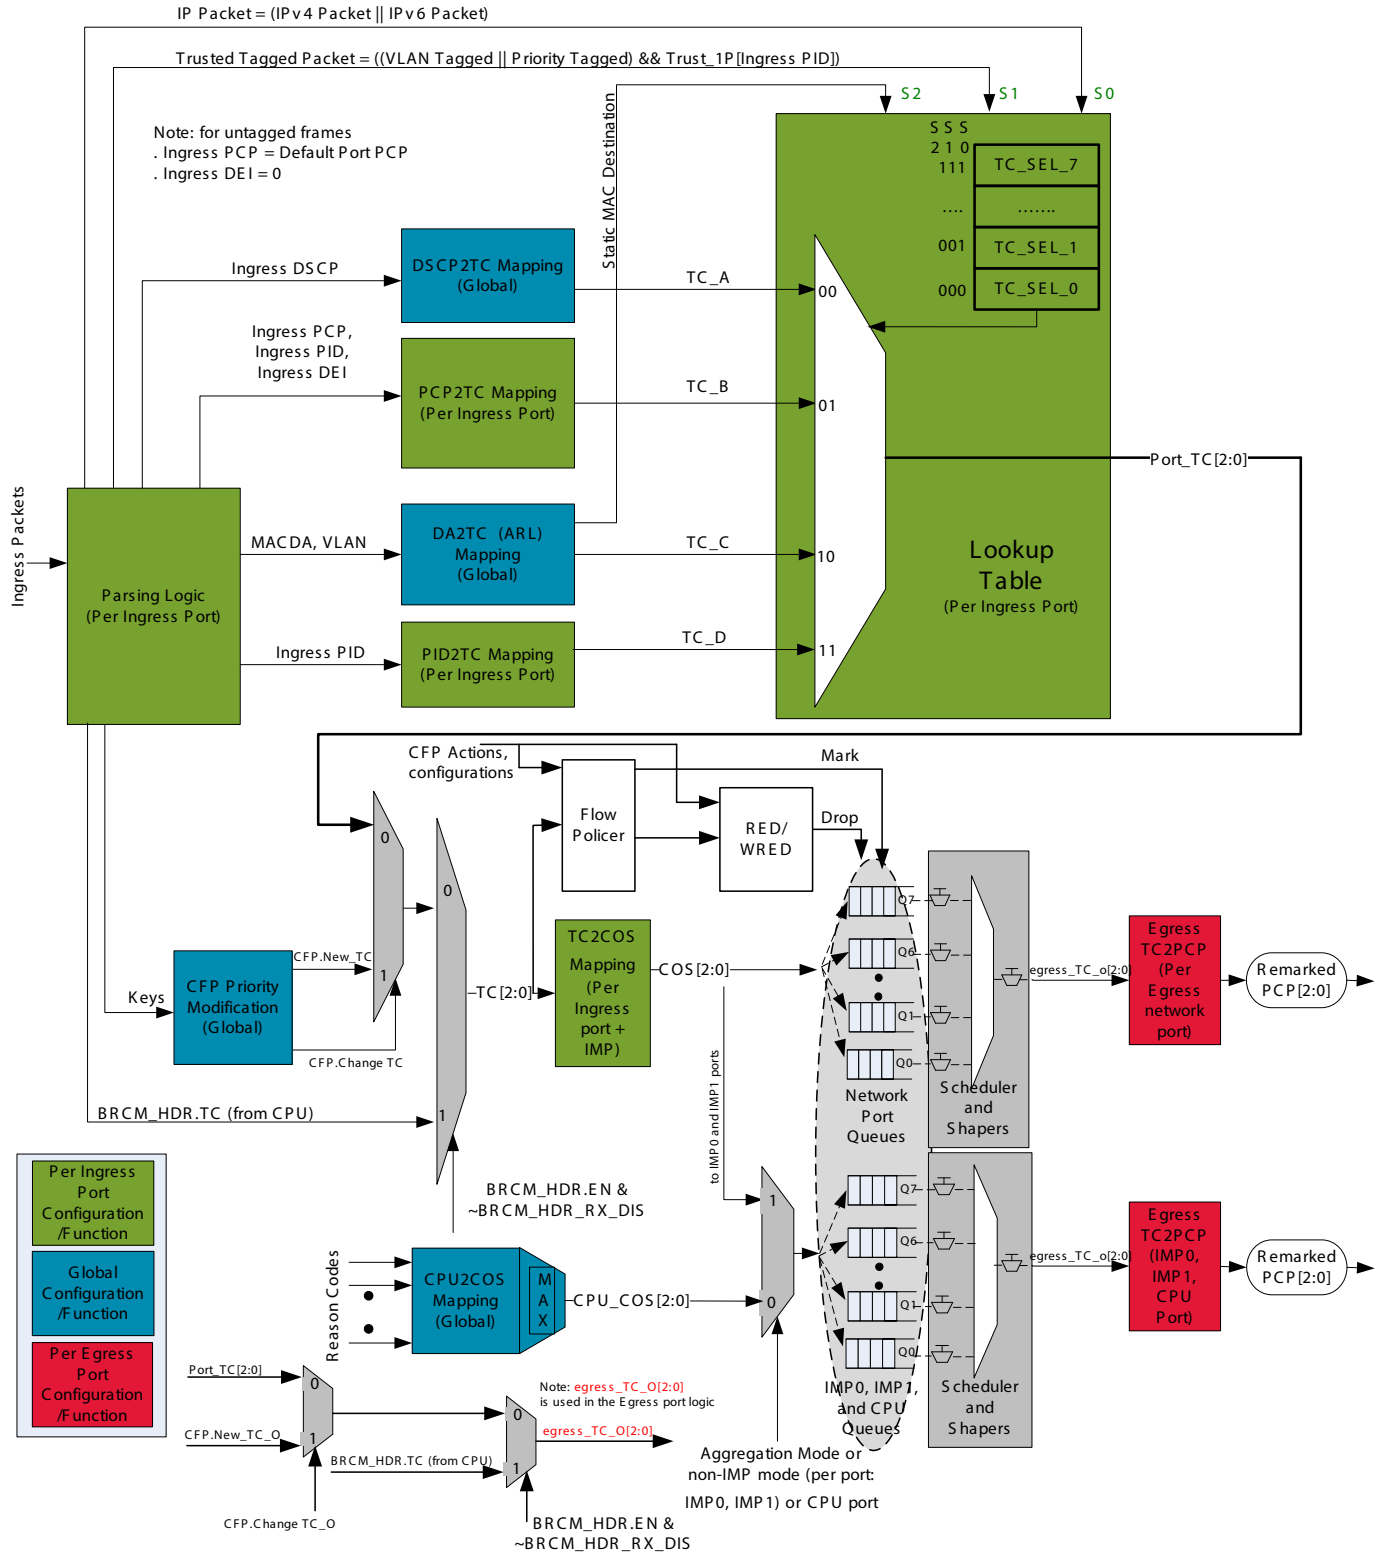
The following sections discuss each feature in detail.

2.2 Quality of Service and Scheduling

The Quality of Service (QoS) feature provides up to eight internal queues per port to support eight different traffic classes (TCs). Traffic class is an internal representation of the priority of an incoming packet inside the device. The traffic class assignment can be programmed so that the user can assign incoming packets to higher/lower TC priorities through TC Mapping. Then, each TC is mapped to one of eight internal class-of-service (CoS) egress queues through TC-to-CoS process. Packets assigned (mapped) to a higher priority output queue in the switch experience less delay than packets with a lower priority under congested conditions. This can be important in minimizing latency for delay-sensitive traffic.

[Figure 2 on page 12](#) shows how the BCM53134P determines the CoS and performs Priority Code Point (PCP) remarking in packets.

Figure 2: CoS and Egress Remarking Flow



The TC is selected from one of the following sources based on software configuration of an eight-entry lookup table corresponding to the ingress port on which the packet was received.

- DSCP-to-TC mapping table (DSCP2TC) (global function)
- The TC of a packet received from an Ethernet (or IMP) port is assigned the TC configured for the corresponding IP TOS/DSCP. When DSCP is disabled, or when the incoming packet is not of IPv4/v6 type, the TC that results from this mapping is 000.
- IEEE 802.1p PCP-to-TC mapping table (PCP2TC) (per ingress port function)
- The TC of a packet received from an Ethernet (or IMP) port is assigned the TC configured for the corresponding IEEE 802.1p priority code point (PCP). When IEEE 802.1p tagging is disabled or when the incoming packet is not tagged, the TC that results from this mapping is 000. The PCP of an ingress-tagged or priority-tagged packet is the same as the PCP field in the outermost VLAN header in the packet. The PCP of an untagged packet is the same as the default PCP register value corresponding to the port on which the packet was received.
- TC from ARL table (DA2TC) (global function)
When using MACDA-based QoS, destination addresses and VLAN IDs are used to index the ARL table, as described in [“Address Management” on page 55](#). The matching ARL entry contains a 3-bit TC field, as shown. These bits set the MACDA-based TC for the frame and TC can also be looked up from ARL table static entries. The MACDA-based TC is assigned to the TC bits depending upon the result. The TC bits for a learned ARL entry default to 0.
- Port-to-TC mapping table (PID2TC) (per ingress port function)
The TC of a packet received from an Ethernet (or IMP) port is assigned the TC configured for the corresponding port. The mapping mechanism is enabled and disabled using Port ID to TC Mapping Register (Page 30h: Address 48h–4Bh) programming. When disabled, the TC that results from this mapping is 000.

The Lookup Table is configured by software for each port separately. The TC generated from the CFP (one of the possible CFP actions is changing TC) can override the TC value generated from the four methods described above (DSCP2TC, PCP2TC, DA2TC, or PID2TC).

The Lookup Table is indexed by the following internal flags:

- IP Packet. The flag indicates that the packet is either an IPv4 or IPv6 packet.
- Trusted Tagged Packet. The flag indicates that the packet is either VLAN-tagged or priority-tagged, and was received on a port that is configured as a trusted port.
- Static MAC Destination. The flag indicates that the MAC destination address matched a static entry in the ARL table.

The TC_SEL_X is an 8-entry x 2-bits-per-entry table for every ingress port. One of the 8 entries (TC_SEL_7..TC_SEL_0) is selected by a 3-bit address {S2, S1, S0}, where:

- S2 = Static MAC Destination
- S1 = Trusted Tagged Packet
- S0 = IP Packet

The two bits (which are configurable by user) that are stored at the indexed entry in the table are then used to select one of the following sources of TC (before it is optionally overridden by the CFP or the TC field in the Broadcom header):

- 00: TC_A
- 01: TC_B
- 10: TC_C
- 11: TC_D

2.2.1 CoS Mapping

All packets should be configured to be mapped to appropriate TCs and those TCs should be mapped to appropriate egress queues through the TC2COS table.

A packet may be sent to the CPU through the IMP0. An additional CPU2COS mapping table may be used for the packet in nonaggregation mode to determine the CoS of those packets. The index to the mapping table entry is determined by a MAX function that selects the maximum value of reason codes. A reason code indicates the reason why a packet is sent to the CPU.

NOTE: In addition to determining the CPU queue in nonaggregation mode, the reason code may also help software process the packet in all modes. In aggregation mode, however, the CoS is determined by the same TC2COS hardware mapping function that is used for network ports.

2.2.2 SF3 Egress Queues and Scheduler

2.2.2.1 Egress Transmit Queues

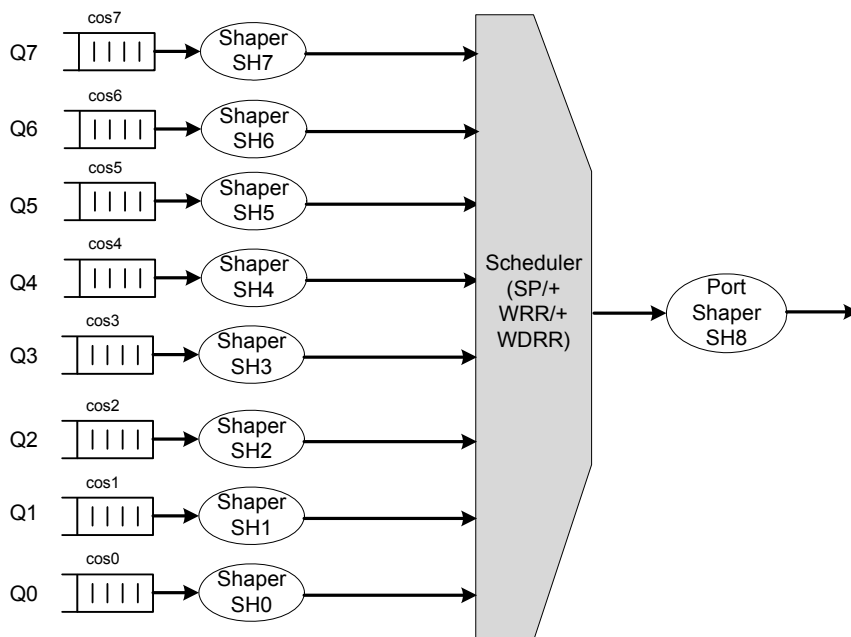
Each Ethernet egress port has eight transmit queues (CoS0–CoS7). Each CoS queue has its own dedicated counter to measure the buffer occupancy of the queue for congestion management purposes. Every Ethernet (ingress) port has its own set of counters to measure the buffer occupancy and the arrival rate related to the traffic received from the port.

The IMP (egress) port and port 5 also serve eight transmit queues. When the IMP port (or port 5) is set in non-aggregation mode (for example, IMP port is configured as a management port to CPU), the CoS (output queue) is decided based on the reasons for forwarding the packets to the CPU. When the IMP port (port 5) is set in aggregation mode (for example, IMP port is configured as a regular data uplink port), the CoS is decided from the TC based the normal packet classification flow.

For the rest of Ethernet egress ports, all incoming frames are assigned to an egress transmit queue depending on their assigned TC. Each egress transmit queue is a list that specifies an order for packet transmission. The corresponding egress port transmits packets from each of the queues according to a programmable algorithm, with the higher TC queues being given greater access than the lower TC queues. Queue 0 is the lowest-TC queue.

The queues and scheduler/shaper for each port are shown in the [Figure 3 on page 14](#).

Figure 3: Queues and Scheduler/Shaper Diagram



2.2.2.2 Scheduler

The scheduling element can be configured to be one of the following operating modes:

1. Strict Priority.
2. Weighted Round Robin (WRR): packet-based scheduling.
3. Weighted Deficit Round Robin (WDRR): byte-based scheduling.
4. Or a mix of SP, WRR, and WDRR.

Table 1 lists various configuration options of the scheduler.

Table 1: Scheduler Configuration Selections

	CoS7	CoS6	CoS5	CoS4	CoS3	CoS2	CoS1	CoS0
Option 1	SP	SP	SP	SP	SP	SP	SP	SP
Option 2	SP	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR
Option 3	SP	SP	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR
Option 4	SP	SP	SP	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR
Option 5	SP	SP	SP	SP	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR
Option 6	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR	WDRR/WRR

The operating mode of a scheduling queue/port can be configured independent of the operating mode of any other scheduling queue/port in the device. One of the two round robin scheduling algorithms (WRR or WDRR) is selected through a per-port configuration register.

Within the SP group, the precedence takes the ascending order of the CoS number. The higher the CoS number, the higher the precedence for scheduling.

Within the WDRR group, the queues are selected based on the following round-robin rules:

- The share can be weighted by assigning each queue its corresponding Weight in granularity of 1/256. If CoS2 is assigned a weight of $x/256$, CoS1 is assigned a weight of $y/256$, and CoS0 is assigned a weight of $z/256$, the ratio of transmitted packet bytes between CoS2:CoS1:CoS0 should be $x:y:z$ in each Scheduling Round, if there is always a packet waiting during the round.
- A Scheduling Round is defined as a period within which each queue gets its fair share of packet transmission resource in granularity of number of packets or number of sets of 256 bytes transmitted.
 - At the beginning of each scheduling round:
 - Every queue's share is added to its respective accumulated credit, and
 - Empty queues are noted and are not considered for scheduling in the scheduling round.
 - During each scheduling round, queues with positive credits in respective queue shaper are serviced in order of CoS (CoS7, CoS6, CoS5,...CoS0):
 - Empty queues as well as queues that do not have positive credits in the respective queue shaper are skipped.
 - When a queue is serviced, the number of packets transmitted from the queue depends on one of the two modes of operations configured by software:

Burst Mode: one or more packets are transmitted from the queue when its fair share of packet transmission resources is used up (credit becomes negative), or the queue becomes empty before the current scheduling round for the queue ends. End of current scheduling round for a queue means the queue will not be serviced again in the current scheduling round.

Non-Burst Mode: one packet is transmitted from each queue in the CoS order (CoS7...0), until the fair share of packet transmission resources of a queue is used up (credit becomes negative), or the queue becomes empty before the current scheduling round for the queue ends. End of current scheduling round for a queue means the queue will not be serviced again in the current scheduling round.

- Also, the negative accumulated credit of a non-empty queue from the current round is carried forward from the current round to the beginning of the next round.
- If a queue becomes empty before its fair share is used up, the current scheduling round for the queue ends will mean that the queue will not be serviced again in the current scheduling round. Also, its left-over credit is not carried over to the next scheduling round, which means the deficit counter of an empty queue is set to zero (to avoid carrying over accumulated credit history of empty queues). The action of setting the accumulated credit to zero is independent of whether the queue became negative after serving the last packet in the queue, or it was still positive after it served the last packet in the queue.
- The scheduler algorithm goes back to Step “At the beginning of each scheduling round” to start the next round.

The weights of a WRR/WDRR scheduler can be configured independent of weights of any other WRR/WDRR scheduler in the device.

The scheduler works by first servicing the SP queues. Queues have an intrinsic priority from high to low. That is, CoS7 has higher priority than CoS6. The second scheduling option utilizes WRR/WDRR. The WRR/WDRR scheduling discipline interleaves packets from queues based on a configured weight for the queues. That is, when a queue is selected for service the depth of the queue is sampled. Packets are transmitted from the queue in the scheduling round until either the queue becomes empty or the credit counter runs out of credits. Once these numbers of packets are serviced, then the next queue is serviced. Queues are serviced in a round-robin fashion (WRR/WDRR) until this process is completed. Queues that do not have packets to send in a round (either empty or do not have positive credits in respective queue shapers), are skipped and not serviced. A queue is serviced only once in a round. If a packet arrives in a queue just after the scheduler decided to skip a queue (because it was empty, became empty) in the current round, the packet will not be serviced in the current round by the scheduler.

For example, assume the queue weights for CoS3–CoS0 are 4, 3, 2, 1 (or 1024, 768, 512, and 256 bytes) respectively, and the scheduler is configured in burst mode of operation. Note that small weights have been used in this example for simplification. In real applications, the minimum (WDRR weight * 256) should be \geq MTU for correct behavior of the WDRR scheduler. Because of programming error, if the weight of any DRR queue is configured such that (WDRR weight * 256) < MTU, the behavior of the DRR scheduler is not predictable.

Assume an example of a scheduling round is similar to [Figure 3 on page 14](#) where the frames are queued only in CoS3, CoS1, and CoS0. CoS2 is empty through the entire round in this example. The egress packet stream created by the WDRR scheduler is depicted where the number in a rectangle indicates the length of the packet in bytes. In the example, the WDRR scheduler creates an egress packet stream by interleaving packets from CoS3, CoS1, and CoS0 (the non-empty queues with accumulated weight > 0) in a fashion that depends on the WDRR weights assigned to the respective CoS.

In the example, assume a scheduling round started with zero accumulated credit for each of the four CoS queues. The WDRR scheme will add credits to the four queues at the beginning of the round. Hence, there will be 1024 bytes accumulated credits for CoS3, 0 for CoS2, 512 for CoS1, and 256 for CoS0. CoS2 accumulated credits at the beginning of the round is 0 since it was either empty, or its shaper was blocked, or the shaper was enabled but did not have positive credits. In the current round, the WDRR scheduler will service the queues in the following order:

1. The scheduler will service CoS3 first because it is the highest order CoS queue that has positive credits in the current round. In the example, CoS3 has 494 bytes in four packets whereas its accumulated credits are 1024 bytes. Hence, it became empty when it still had 530 positive accumulated credits. However, since the queue became empty before it

could use up all its credits, the WDRR scheme will carry forward 0 accumulated credits for CoS3 to the next round. The scheduler will not service CoS3 any more in the current round even if a new packet is stored in the queue while it is servicing any remaining queues in the current round.

2. It will skip CoS2 next because it has 0 accumulated credits. The scheduler will carry forward 0 accumulated credits to the next round for CoS2. The scheduler will not service CoS2 any more in the current round even if a new packet is stored in the queue while it is servicing the remaining queues in the current round.
3. It will service CoS1 next because it has the next highest order CoS queue. CoS1 has 780 bytes in four packets, which is more than 512 bytes of accumulated credits. Hence, the WDRR scheduler will stop servicing the queue and carry forward -278 accumulated credits for CoS1 to the next round after the first three packets from the queue are transmitted. The scheduler will not service CoS1 any more in the current round even if a new packet is stored in the queue while it is servicing the remaining queues in the current round.
4. Finally, it will serve CoS0 and then end the current round. CoS0 has 204 bytes in two packets, which is a little less than 256 bytes of accumulated credits for CoS1. Since the queue will become empty before its positive credits are used up, the WDRR scheduler will carry forward 0 credits for CoS0 to the next round. The scheduler will not service CoS0 any more in the current round even if a packet arrives before the next round starts.

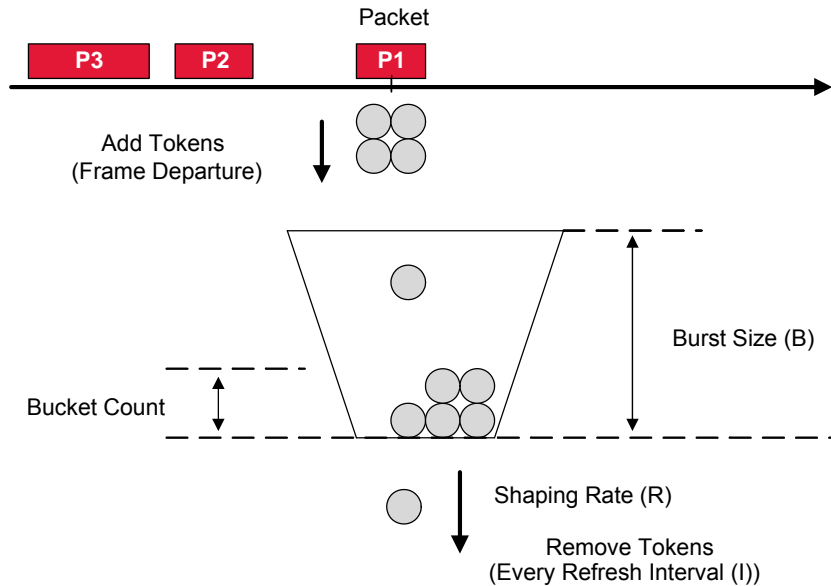
2.2.3 Scheduling

There are three shapers, SH2, SH1, and SH0, and they can be separately configured to be either a credit-based AVB shaper that is compliant with the IEEE 802.1qav standard, or a standard Leaky Bucket rate limiter.

2.2.4 Leaky Bucket Shaper

Initially, the bucket is filled with tokens with a burst size (B). In every refresh interval, certain amounts of tokens are removed from the bucket, which effectively determines the shaping rate. The scheduler services a packet based on the state information of the shaper. If the bucket count is less than or equal to the burst size, the packet is in-profile. Otherwise, the packet is out-of-profile. After servicing a packet, tokens are added into the bucket by an amount of tokens equivalent to the size of the packet. In the BroadSync HD (AVB) shaping mode, the following additional token updates are required to reduce the burstiness of the EAV traffic. After the bucket counter is updated, if the queue is empty and the bucket counter is less than the burst size (B), reset the bucket counter to be equal to burst size (B). See [Figure 4](#).

Figure 4: Leaky Bucket Shaper



The leaky bucket shaper has the following parameters:

- Refresh Interval (I): Defines the how often the tokens are removed from the buckets.
- Shaping Rate (R): The rate at which the shaper limits service.
- Burst Size (B): The maximum number of tokens that can be added into the bucket.
- AVB Shaping Mode: Used to select AVB versus Normal shaping mode.

In every refresh interval, T tokens are removed from the bucket. T is the number of tokens. Associated with a token is a token size (S). The shaper operates in byte-based mode, and the token size (S) is the number of bytes per token.

The relation between Shaping Rate (R), Refresh Interval (I) and Token Size (S) can be represented by formula:

$$R = T \times (S/I)$$

NOTE: The refresh interval is fixed at 7.8125 μ s.

It is recommended that software configure the two shapers in the following way when applicable:

- **SH2** should be configured as:
 - An AVB shaper for an AVB port with Class A traffic.
 - A non-AVB shaper in all other applications.
- **SH1** should be configured as:
 - An AVB shaper for an AVB port with Class B traffic.
 - A non-AVB shaper in all other applications.
- **SH0** should be configured as a non-AVB shaper in all applications.

The Leaky Bucket threshold of a shaper can be configured in the range of 64 bytes–16 MB, with a resolution of 64 bytes. When the shaper operates as an AVB shaper, the number of tokens in the shaper is saturated by hardware (made equal to the configured threshold) when there is no packet at the shaper input. For example, if SH2 is configured as an AVB shaper and Q5 is empty, then hardware forces the accumulated credits in SH2 to be the same as the threshold. Each shaper output rate can be configured in the range of 64 Kb/s–1 Gb/s, with a resolution of 64 Kb/s.

2.3 Port-Based VLAN

The port-based virtual LAN (VLAN) feature partitions the switching ports into virtual private domains designated on a per-port basis. Data switching outside of the port's private domain is not allowed. The BCM53134P provide flexible VLAN configuration for each ingress (receiving) port.

The port-based VLAN feature works as a filter, filtering out traffic destined to nonprivate domain ports. For each received packet, the ARL resolves the DA and obtains a forwarding vector (list of ports to which the frame will be forwarded). The ARL then applies the VLAN filter to the forwarding vector, effectively masking out the nonprivate domain ports. The frame is forwarded only to those ports that meet the ARL table criteria, as well as the port-based VLAN criteria.

2.4 IEEE 802.1Q VLAN

The BCM53134P support IEEE 802.1Q VLAN and up to approximately 4096 VLAN table entries that reside in the internal embedded memory. Once the VLAN table is programmed and maintained by the microcontroller, the BCM53134P autonomously handle all operations of the protocol. These actions include the stripping or adding of the IEEE 802.1Q tag, depending on the requirements of the individual transmitting port. It also performs all the necessary VLAN lookups in addition to MAC L2 lookups.

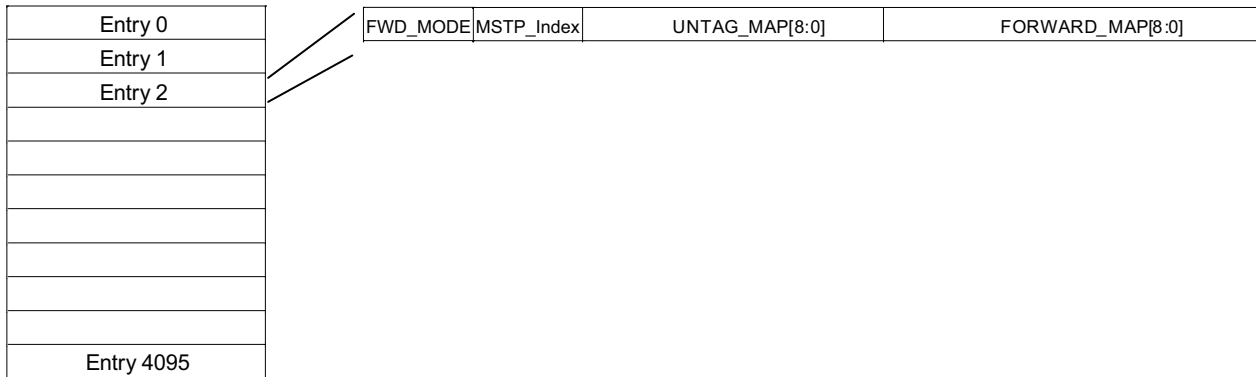
2.4.1 IEEE 802.1Q VLAN Table Organization

Each VLAN table entry, also referred to as a VLAN ID, an untag map, and a forward map:

- The untag map controls whether the egress packet is tagged or untagged.
- The forward map defines the membership within a VLAN domain.
- The FWD_MODE indicates whether the packet forwarding should be based on VLAN membership or on ARL flow.

The untag map and forward map include bit-wise representation of all the ports.

Figure 5: VLAN Table Organization




NOTE: If the MII port is configured as a management port, then the tag is not stripped even if the untag bit is set.

2.5 Double-Tagging

The BCM53134P provide the double tagging feature, which is useful for ISP applications. When the ISP aggregates incoming traffic from each individual customer, the extra tag (double tag) can provide an additional layer of tagging to the existing IEEE 802.1Q VLAN. The ISP tag (extra tag) is a way of separating individual customers from other customers. Using the IEEE 802.1Q VLAN tag, the individual customer's traffic can be identified on a per-port basis.

When the double-tagging feature is enabled (register Page 34h, Address 05h, bit[3:2]) and the Enable IEEE 802.1Q (register Page 34h, Address 00h, bit 7), users can expect two VLAN tags in a frame: the tag close to MAC_SA is the ISP tag, and the one following is the customer tag as shown in [Figure 6](#).

Figure 6: ISP Tag Diagram



The switch uses the ISP tag for ARL and VLAN table accesses and the customer tag as an IEEE 802.1Q tag. There is a per-chip programmable register Double Tagging TPID register for ISP tag (default = 9100'h). All ISP tags will be qualified by this Tag Protocol ID (TPID) value.

When the double-tagging feature is enabled, all switch ports are separated into two groups: ISP ports and customer ports. The BCM53134P performs the normalization process for all ingress frames for both Intelligent Double Tag (IDT) and Double Tag (DT) modes, whether from the ISP port or customer port. The normalization process is to insert an ISP tag, customer tag, or ISP + customer tag (depending on whether the ingress frame is without tags or with one tag) to allow all ingress frames with double tags. However, if the ingress frames are with double tags (ISP + customer tag), and the ISP tag TPID matches the TPID specified in the Double Tagging TPID register, it does not perform the normalization process. The ISP ports are defined in the ISP Port Selection Portmap register. When the port(s) corresponding bit(s) are set, those port(s) should be connected to the ISP, and otherwise connected to customers. Each switch device can have multiple ports assigned as ISP ports, and each ISP is uniquely identified using different VLAN forward maps or the port-based VLAN feature.

2.5.1 ISP Port

It is possible for the ISP port to receive three different types of frames: untagged, ISP-tagged, and ISP+Customer-tagged frames.

When the double-tagging feature is enabled and the received frame is untagged (or the TPID does not match with ISP TPID specified in Double Tagging TPID register, the default ISP tag and customer tag are added, and VLAN ID of ISP tag receives it from the port default VID. The frames are forwarded according to the VLAN table. However, if the Port-Based VLAN Control register is enabled, the egress ports specified in the port-VLAN control register override the VLAN table settings. If the received frame is ISP tagged (TPID matches with the ISP tag VLAN ID specified in the double-tagging TPID register), the default customer tag (8100 + default PVID) is added, the ISP VID is used to access the ARL table, and the ISP tag can be stripped on the way out according to the untagged bit setting in the VLAN table. In addition, ISP port frame can forward to the destination port directly based on forward port map of VLAN table by setting FWD_MODE bit to 1 of VLAN Table Entry register.

The VLAN ID is generated from the ISP tag, and TC is generated from the ingress frame outer tag.

2.5.2 Customer Port

It is also possible for the Customer port to receive two different types of frames: untagged and Customer-tagged frames.

When the double-tagging feature is enabled, all the ingress frames perform the normalization process to insert an ISP tag or ISP + Customer tag (depending whether the ingress frame is without tags or with one tag) to allow all ingress frames with double tags. The VLAN ID of ISP tag receives it from the port default VID.

The VLAN ID is generated from the ISP tag, and the TC is generated from the ingress frame outer tag.

NOTE: It is illegal to strip out the ISP tag on the ISP egress port by using the untagged bit setting in the VLAN table.

NOTE: Only the VLAN tagged or untagged packets are expected for the ingress of the customer ports. The customer does not add the ISP tags.

There are two possible traffic scenarios:

- One scenario is from a customer port to an ISP port.
- The second scenario is from an ISP port to a customer port.

2.5.3 Uplink Traffic (from Customer Port to ISP)

Data traffic is traffic received from the customer port without tags or a customer tag, and the frame is destined for an ISP port. The customer ingress port performs a normalization process to allow ingress frames with double tags (ISP + Customer tag), and the ISP tag VID is based on the port default VID tag.

However, if the ingress frame is with an IEEE 802.1p tag, the VID of IEEE 802.1p tag is changed by the VID of the port default VID tag after the customer port normalization process. The TC do not change.

Control traffic frames can be forwarded to the CPU first and then the CPU forwards to the ISP port if the switch management mode is enabled and if the RESV_MCAST_FLOOD bit=0 in the Global VLAN Control 4 register. In this case, the control frame adds an ISP tag by ingress port and forwards to the CPU. The CPU can then forward it to the ISP port with or without the ISP tag by using the egress-direct feature.

2.5.4 Downlink Traffic (from ISP to Customer Port)

Data traffic frame received from the ISP port may or may not have an ISP tag attached. When the received frame does not have an ISP tag and customer tag, the ISP ingress port does a normalization process to insert double tags (ISP + Customer tag), and the ISP tag VID is based on the port default VID tag. All ARL and VID table access should be based on the new tag. The traffic is then forwarded to the customer port through proper VLAN configuration. Usually, the software configures so the customer Egress port continuously removes the ISP tag. However, it is based on how the untagged map is configured.

Moreover, if the ingress frame is with an IEEE 802.1p tag, the VID of the IEEE 802.1p tag is changed by the VID of the port default VID tag after the ISP port normalization process. The TC will not change.

The Control traffic is forwarded to the CPU when the switch management mode is enabled and if RESV_MCAST_FLOOD bit=0 in the Global VLAN Control 4 register. The BCM53134P can also support multiple ISP port configurations by enabling the FWD_MODE bit of the VLAN Table Entry register.

There are also two ways to separate traffic that belongs to two different ISP customers:

1. Each group (ISP and customer) is assigned to the same VLAN group, so that traffic does not leak to other ISP.
2. Use the Port-based VLAN to separate traffic that belongs to a different ISP.

2.6 Egress VID Modification

The BCM53134P provides the egress VID remarking feature. The VID field of the outer tag or inner tag can be modified or removed. This feature is port-dependent; each port can be set up differently. This feature is available only when Intelligent Double Tag (IDT) mode is enabled (Page 34h, Address 05h, Bit[3:2]).

The BCM53134P performs the normalization process for all ingress packets when double-tag enabled (DT_Mode or IDT_Mode).

At each egress port, the double-tag of a normalized packet is subject to VID modification before the packet is transmitted by the corresponding port.

The egress VID modification instruction is generated through a mapping table indexed by the Egress Port ID and the Classification ID generated from the CFP. Each egress port supports 256 entries to the VID mapping table (register Page 34h, Address 40h).

Table 2: VID Modification Instruction Mapping

Index	VID Instruction for Outer Tag	VID Instruction for Inner Tag
{Egress Port ID, Classification ID}	<p>00 = Indicates the outer tag should be sent as the internally normalized VLAN tag.</p> <p>01 = Indicates the outer tag should be sent as the one when the packet is received from its ingress port. If the packet is received originally with an outer VLAN tag, it should be sent out through the egress port with the same outer VLAN tag, otherwise it should be sent out without the outer VLAN tag.</p> <p>10 = Remove outer tag before the packet is transmitted.</p> <p>11 = VID modified on outer tag before the packet is transmitted.</p>	<p>00 = Indicates the inner tag should be sent as the internally normalized VLAN tag.</p> <p>01 = Indicates the inner tag should be sent as the one when the packet is received from its ingress port. If the packet is received originally with an inner VLAN tag, it should be sent out through the egress port with the same outer VLAN tag, otherwise it should be sent out without the inner VLAN tag.</p> <p>10 = Remove inner tag before the packet is transmitted.</p> <p>11 = VID modified on inner tag before the packet is transmitted.</p>


2.7 Jumbo Frame Support

The BCM53134P can receive and transmit frames of extended length on ports linked at gigabit speed. Referred to as jumbo frames, these packets are longer than the standard maximum size, but shorter than 9720 bytes. Jumbo packets can be received or forwarded only to 1000BASE-T–linked ports that are jumbo-frame enabled. Up to 38 buffer memory pages are required for storing the longest allowed jumbo frame. While there is no physical limitation to the number of ports that can be jumbo enabled, it is recommended that no more than two be enabled simultaneously to ensure system performance. There is no performance penalty for enabling additional jumbo ports beyond the potential strain on memory resources that can occur due to accumulated jumbo packets at multiple ports.

2.8 Port Trunking/Aggregation

The BCM53134P supports MAC-based trunking. The trunking feature allows up to four ports to be grouped together as a single-link connection between two switch devices. This increases the effective bandwidth through a link and provides redundancy. The BCM53134P allow up to two trunk groups. Trunks are composed of predetermined ports and can be enabled using the Trunking Group 0 register. Ports within a trunk group must be of the same linked speed. By performing a dynamic hashing algorithm on the MAC address, each packet destined for the trunk is forwarded to one of the valid ports within the trunk group. This method has several key advantages. By dynamically performing this function, the traffic patterns can be more balanced across the ports within a trunk. In addition, the MAC-based algorithm provides dynamic failover. If a port within a trunking group fails, the other port within the trunk automatically assumes all traffic designated for the trunk. It allows for a seamless, automatic redundancy scheme. This hashing function can be performed on the DA, SA, or DA/SA.

Figure 7: Trunking

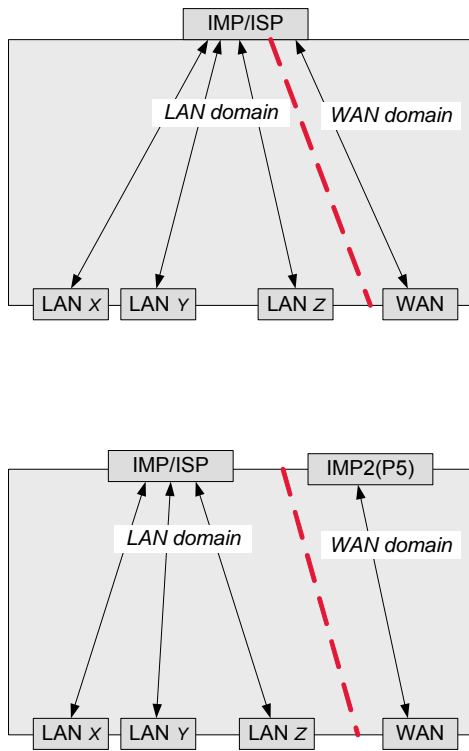


2.9 WAN Port

The BCM53134P offers a programmable WAN port feature: a WAN Port Select register (page 00h, address 26h). Select a port as a WAN port to forward all of that port's traffic to only the CPU port. The non-WAN port traffic from all other local ports does not flood to the WAN port.

Figure 8 shows the WAN and LAN domain separation when the WAN port is selected.

Figure 8: WAN and LAN Domain Separation




2.10 Rate Control

2.10.1 Ingress Rate Control

Forwarding broadcast traffic consumes switch resources, which can negatively impact the forwarding of other traffic. The rate-based broadcast storm suppression mechanism is used to protect regular traffic from an overabundance of broadcast or multicast traffic. This feature monitors the rate of ingressed traffic of programmable packet types. If the rates of these packet types exceed the programmable maximum rate, the packets are dropped.

The broadcast storm suppression mechanism works on a credit-based rate system that figuratively uses a bucket to track the bandwidth of each port (see Figure 9). Credit is continually added to the bucket at a programmable bucket bit rate. Credit is decremented from the bucket whenever one of the programmable packet types is ingressed at the port. If no packets are ingressed for a considerable length of time, the bucket credit continues to increase up to a programmable-maximum bucket size. If a heavy burst of traffic is suddenly ingressed at the port, the bucket credit becomes drained. When the bucket is emptied, incoming traffic is constrained to the bucket bit rate (the rate at which credit is added to the bucket). At this point, excess packets either are dropped or deterred using flow control, depending upon the Suppression Drop mode.

Figure 9: Bucket Flow



2.10.2 Two-Bucket System

For added flexibility, the BCM53134P employs two buckets to track the rate of ingressed packets. Each of the two buckets (Bucket 0 and Bucket 1) can be programmed to monitor different packet types. For example, Bucket 0 could monitor broadcast packets, while Bucket 1 monitors multicast packets. Multiple packet types can be monitored by each bucket, and a packet type can be monitored by both buckets.

The rates of each bucket can be individually programmed. For example, the broadcast packets of Bucket 0 could have a maximum rate of 3 Mb/s, whereas the multicast packets of Bucket 1 could be allowed up to 80 Mb/s. The size of each bucket can be programmed. This determines the maximum credit that can accumulate in each bucket. The rate count and bucket size can be individually programmed for each port, providing another level of flexibility. Suppression control can be enabled or disabled on a per-port basis. This system allows the user to control dual packet-type rates on a per-port basis.

2.10.3 Egress Rate Control

The BCM53134P monitors the rate of egress traffic per port. Unlike the ingress traffic rate control, the egress rate control provides only the per-port rate control regardless of traffic types. This feature uses only one bucket to track the rate of egressed packets. The egress rate control feature supports only absolute bit rate mode (Bit Rate Mode = 0), and the bucket bit rate calculation is shown in [Table 3](#).

2.10.4 Bucket Bit Rate

The relative ingress rates of each bucket can be programmed on a per-port basis. Each port has a programmable rate count value for Bucket 0 and Bucket 1. Additionally, the bit rate mode is programmed on a chip basis. If this bit is 1, the packet rate is automatically scaled according to the port link speed. Ports operating at 1000 Mb/s would be allotted a 100 times higher ingress rate than ports linked at 10 Mb/s. Together, the rate count value and the bit rate mode determine the bucket bit rate, which is a reflection of how quickly data can be ingressed (Kb/s) at the given port for a given bucket. The rate count values are specified in [Table 3](#). Values outside these ranges are not valid entries.

Table 3: Bucket Bit Rate

Rate Count (RC)	Bit Rate Mode	Link Speed	Bucket Bit Rate Equation	Approximate Computed Bucket Bit Rate Values (as a function of RC)
1–28	0	Any	$(RC \times 8 \times 1M)/125$	64 KB, 128 KB, 192 KB,..., 1.792 MB
29–127	0	Any	$(RC - 27) \times 1M$	2 MB, 3 MB, 4 MB,..., 100 MB
128–240	0	Any	$(RC - 115) \times 1M \times 8$	104 MB, 112 MB, 120 MB,..., 1000 MB
1–125	1	10 Mb/s	$(RC \times 8 \times 1M)/100$	0.08 MB, 0.16 MB, 0.24 MB,..., 10 MB
1–125	1	100 Mb/s	$(RC \times 8 \times 1M)/10$	0.8 MB, 1.6 MB, 2.4 MB,..., 100 MB
1–125	1	1000 Mb/s	$RC \times 8 \times 1M$	8 MB, 16 MB, 24 MB,..., 1000 MB

NOTE: 1M represents 1×10^6 .

2.11 Protected Ports

The Protected Ports feature allows certain ports to be designated as protected. All other ports are unprotected. Traffic between protected port group members is blocked. However, protected ports are able to send traffic to unprotected ports. Unprotected ports can send traffic to any port. Several applications that can benefit from protected ports:

- **Aggregator:** For example, all the available ports are designated as protected ports except a single aggregator port. No traffic incoming to the protected ports is sent within the protected ports group. Any flooded traffic is forwarded only to the aggregator port.
- To prevent nonsecured ports from monitoring important information on a server port, the server port and nonsecured ports are designated as protected. The nonsecured ports will not be able to receive traffic from the server port.


2.12 Port Mirroring

The BCM53134P support Port Mirroring, allowing ingress and/or egress traffic to be monitored by a single port designated as the mirror capture port. The BCM53134P can be configured to mirror the ingress traffic and/or egress traffic of any other port (s). Mirroring multiple ports is possible, but can create congestion at the mirror capture port. Several filters are used to decrease congestion.

2.12.1 Enabling Port Mirroring

Port Mirroring is enabled by setting the Mirror Enable bit.

Figure 10: Mirror Filter Flow



2.12.2 Capture Port

The capture port is capable of monitoring other specified ports. Frames transmitted and received at the other ports are forwarded to the Capture port according to the mirror filtering rules discussed below.

2.12.3 Mirror Filtering Rules

Mirror filtering rules consist of a set of three filter operations (Port Mask, Packet Address, and Packet Divider) that are applied to traffic ingressed and/or egressed at a switch port.

2.12.3.1 Port Mask Filter

The IN_MIRROR_MASK bits define the receive ports that are monitored. The OUT_MIRROR_MASK bits define the transmit ports that are monitored.

Any number of ingress/egress ports can be programmed to be mirrored, but bandwidth restrictions on the one-mirror capture port should be taken into account to avoid congestion or packet loss.

2.12.3.2 Packet Address Filter

The type of filtering that is applied to frames received on the mirrored ports is configurable. The IN_MIRROR_FILTER bits select among the following where x is the 48-bit MAC address. Likewise, the type of filtering that is applied to frames transmitted on the egressed mirrored ports:

- Mirror all received frames
- Mirror received frames with DA = x
- Mirror received frames with SA = x

2.12.3.3 Packet Divider Filter

The IN_DIV_EN bit allows further statistical sampling. When IN_DIV_EN = 1, the receive frames passing the initial filter are divided by the value IN_MIRROR_DIV, which is a 10-bit value. Only one out of every n frames is forwarded to the mirror capture port, where $n = \text{IN_MIRROR_DIV} + 1$. This allows the following additional capabilities:

- Mirror every n^{th} received frame
- Mirror every n^{th} received frame with DA = x
- Mirror every n^{th} received frame with SA = x

NOTE: When multiple ingress ports have been enabled in the IN_MIRROR_MASK, the cumulative total packet count received from all ingress ports is divided by the value of IN_MIRROR_DIV to deliver the nth receive frame to the mirror capture port. Egressed frames are governed by the OUT_MIRROR_MASK bit and the OUT_MIRROR_DIV bit.

2.13 IGMP Snooping

The BCM53134P supports IP layer IGMP Snooping, which includes IGMP unknown, query, report, and leave message.

A frame with a value of 2 in the IP header protocol field and IGMP frames are forwarded to the CPU port. The management CPU can then determine, from the IGMP control packets which port should participate in the multigroup session. The management CPU proactively programs the multicast address in the ARL table or the multiport address entries. If the IGMP_UKN_FWD_EN, IGMP_QRY_FWD_EN, IGMP_RPTLVE_FWD_EN is enabled, IGMP frames will be trapped to the CPU port only.

2.14 MLD Snooping

The BCM53134P supports IP layer MLD Snooping, which includes MLD query, report, and done message. For of the query and report/done message types, there are four options available: discard, forward normally, forward to CPU, or forward normally and copy to CPU. The CPU is then expected to interpret these messages and configure the address table accordingly.

2.15 IEEE 802.1X Port-Based Security

IEEE 802.1X is a port-based authentication protocol. By receiving and extracting special frames, the CPU can control whether the ingress and egress ports should forward packets or not. If a user port wants service from another port (authenticator), it must get approved by the authenticator. EAPOL is the protocol used by the authentication process. The BCM53134P detects EAPOL frames by checking the destination address of the frame. The Destination addresses should be either a multicast address as defined in IEEE 802.1X (01-80-C2-00-00-03) or a user-predefined MAC (unicast or multicast) address. Once EAPOL frames are detected, the frames are forwarded to the CPU so it can send the frames to the authenticator server. Eventually, the CPU determines whether the requestor is qualified or not based on its MAC_Source addresses, and frames are either accepted or dropped. The per-port EAP can be programmed in the register.

BCM53134P provides three modes for implementing the IEEE 802.1X feature. Each mode can be selected by setting the appropriate bits in the register.

- The Basic Mode (when EAP Mode = 00'b) is the standard mode. The EAP_BLK_MODE bit is set before authentication to block all of the incoming packets. Upon authentication, the EAP_BLK_MODE bit is cleared to allow all the incoming packets. In this mode, the Source Address of incoming packets is not checked.
- The second mode is Extended Mode (when EAP Mode = 10'b), where an extra filtering mechanism is implemented after the port is authenticated. If the Source MAC address is unknown, the incoming packets would be dropped and the unknown SA is not learned. However if the incoming packet is IEEE 802.1X packet, or special frames, the incoming

packets is forwarded. The definition of the Unknown SA in this case is when the switch cannot match the incoming Source MAC address to any of the addresses in ARL table, or the incoming Source MAC address matches the address in ARL table, but the port number is mismatched.

- The third mode is Simplified Mode (when EAP Mode = 11'b). In this mode, the unknown Source MAC address packets would be forwarded to CPU rather than dropped. Otherwise, it is same as the Extended Mode operation.

NOTE: The BCM53134P checks only the destination addresses to qualify EAPOL frames. Ethernet type fields, packet type fields, or non-IEEE 802.1Q frames are not checked.

2.16 DoS Attack Prevention

The BCM53134P supports the detection of the following DoS (Denial of Service) attack types based on register setting, which can be programmed drop or not to drop each type of DoS packets respectively.

Table 4: DoS Attacks Detected by BCM53134P

DoS Attack Type	Description
IP_LAND	IPDA = IPSA in an IPv4/IPv6 datagram
TCP_BLAT	DPort = SPort in a TCP header carried in an unfragmented IP datagram or in the first fragment of a fragmented IP datagram
UDP_BLAT	DPort = SPort in a UDP header carried in an unfragmented IP datagram or in the first fragment of a fragmented IP datagram
TCP_NULLScan	Seq_Num = 0 and all TCP_FLAGS = 0 in a TCP header carried in an unfragmented IP datagram or in the first fragment of a fragmented IP datagram
TCP_XMASScan	Seq_Num = 0, FIN = 1, URG = 1, and PSH = 1 in a TCP header carried in an unfragmented IP datagram or in the first fragment of a fragmented IP datagram
TCP_SYNFINScan	SYN = 1 and FIN = 1 in a TCP header carried in an unfragmented IP datagram or in the first fragment of a fragmented IP datagram
TCP_SYNErrror	SYN = 1, ACK = 0, and SRC_Port < 1024 in a TCP header carried in an unfragmented IP datagram or in the first fragment of a fragmented IP datagram
TCP_ShortHDR	The length of a TCP header carried in an unfragmented IP datagram or the first fragment of a fragmented IP datagram is less than MIN_TCP_Header_Size
TCP_FragError	The Fragment_Offset = 1 in any fragment of a fragmented IP datagram carrying part of TCP data
ICMPv4_Fragment	The ICMPv4 protocol data unit carried in a fragmented IPv4 datagram
ICMPv6_Fragment	The ICMPv6 protocol data unit carried in a fragmented IPv6 datagram
ICMPv4_LongPing	The ICMPv4 ping (echo request) protocol data unit carried in an unfragmented IPv4 datagram with its Total Length indicating a value greater than the MAX_ICMPv4_Size + size of IPv4 header
ICMPv6_LongPing	The ICMPv6 ping (echo request) protocol data unit carried in an unfragmented IPv6 datagram with its payload length indicating a value greater than the MAX_ICMPv6_Size

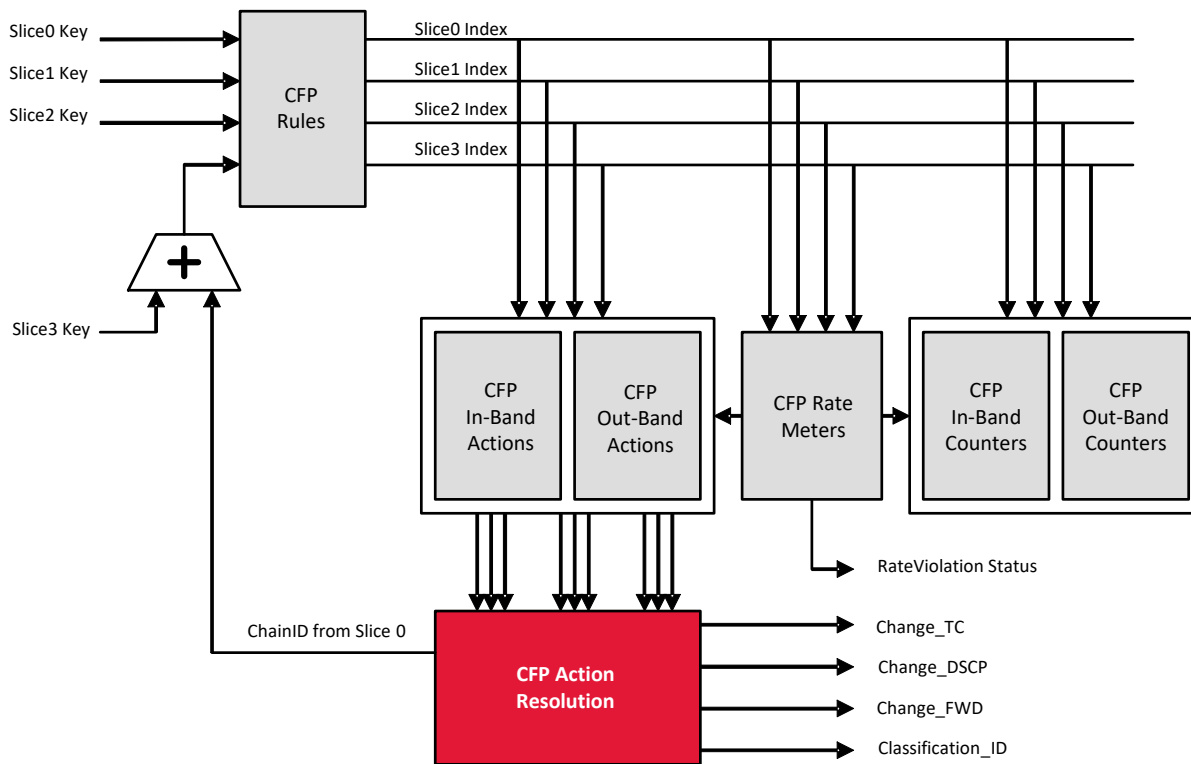
- MIN_TCP_Header_Size is programmable between 0 and 255 bytes, inclusive. The default value is set to 20 bytes (TCP header without options).
- MAX_ICMPv4_Size is programmable between 0 and 9.6 KB, inclusive. The default value is set to 512 bytes.
- MIN_TCP_Header_Size is programmable between 0 and 9.6 KB, inclusive. The default value is set to 512 bytes.
- The default control setting for all types of DoS attacks is not to drop the DoS attack packet.
- The smurf/fraggle type of DoS packet can be detected and dropped through CFP rules by matching configurable directed IP broadcast addresses and the appropriate ICMP type or UDP port.
- The ping flooding or SYN/ACK flooding types of DoS attack can be detected and rate-metered through CFP rules by matching the appropriate ICMP type or the appropriate TCP control flags.
- It is globally configurable whether to perform the SA learning operation with the received packets of the DoS attack type defined in the registers, regardless of the individual DoS attack types.
- Once a packet is detected as a DOS attack type that must be dropped, the packet is dropped regardless of ARL or CFP forwarding decisions, but its forwarding based on mirroring function is not affected.

2.17 Compact Field Processor

The BCM53134P is designed with a Compact Field Processor (CFP) to provide support for ACL (Access Control Logic) implementation and some proprietary protocols, which require nonstandard actions. CFP is a versatile filter and classification processor. The CFP implements ContentAware processing based on TCAM technology.

The overall CFP structure is illustrated in [Figure 11](#).

Figure 11: CFP Engine



There are three predefined slice formats, IPv4, IPv6, and non-IP formats. Formats are selected through the L3_Framing field. Each format has up to three user-defined slice variations: Slice0, Slice1, and Slice2. Slice variations are selected through the Slice ID field. In the case of IPv6, an extra slice (Slice3) is available which can be chained for an expanded search field. However, only Slice0 can be chained to Slice3.

There are a total of 128 rules (for all formats and slices combined) available to users.

Users can define the rules for search keys and set the corresponding mask bits. The contents of the rule searching keys are generated from the packet parser for each packet arriving at the BCM53134P. Although all the CFP rules are stored in a single physical TCAM entity, each slice of rules represents a logical portion of CFP rules and can be assigned with any number of rules for the slice (up to the physical size limitation).

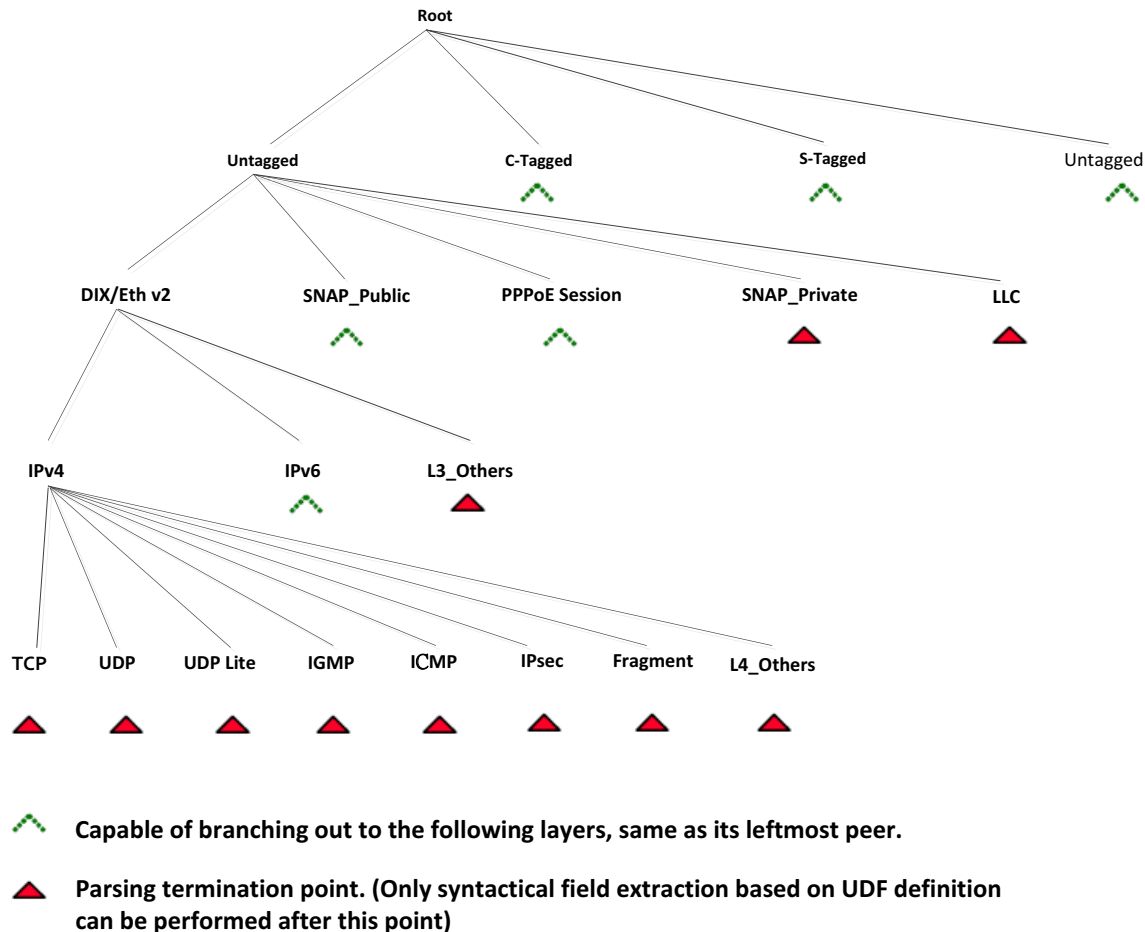
Each CFP rule is associated with a dedicated rate meter that is running at a granularity of 64 Kb/s, and a pair of packet counters. There is an in-band counter for statistics collection of packets arriving within the provisioned rate profile, and an out-band counter for statistics collection of packets arriving outside the provisioned rate profile. The counters are selectively activated based on the rate metering result.

Each arriving packet is allowed up to four slices of rule searching. The matched action sets are converged using the multi-action resolution procedure to generate the switching actions in terms of forwarding/filtering, priority assignment, remarking, classification, etc. Actions can be either metering-dependent or metering-independent.

2.17.1 Parser

In each port, every received packet would go through the following parsing process to compose the CFP rule search slice keys for the CFP based packet switching operations. The packet parsing flow is shown in [Figure 12](#).

Figure 12: Overall Parsing Flow for CFP



At layer 2, the parser is able to recognize up to 2 levels of VLAN tags, and the following basic Ethernet framing types: DIXv2/Ethernet II, LLC, SNAP (Public to carry EtherType protocol, Private to carry system vendor specific protocols, and PPPoE session). The parsing of L2 or L2+ tunneling protocol headers following the before mentioned Ethernet framing, such as MPLS, etc. is not supported.

At layer 3, the parser is able to recognize only IPv4 and IPv6 specified by the preceding EtherType. In case that there is another IP (v4 or v6) header after the initial IP header (in various flavors of IOverIP tunneling protocols), it is treated as upper layer protocol header by the parser; that is, parsing within an L3 tunnel is not supported.

At layer 4 and above, the parser is able to extract upper layer protocol fields starting from where L3 header parsing is finished, such as TCP/UDP header fields.

The detailed syntax and semantics parsing of each layer is described in the following sections.

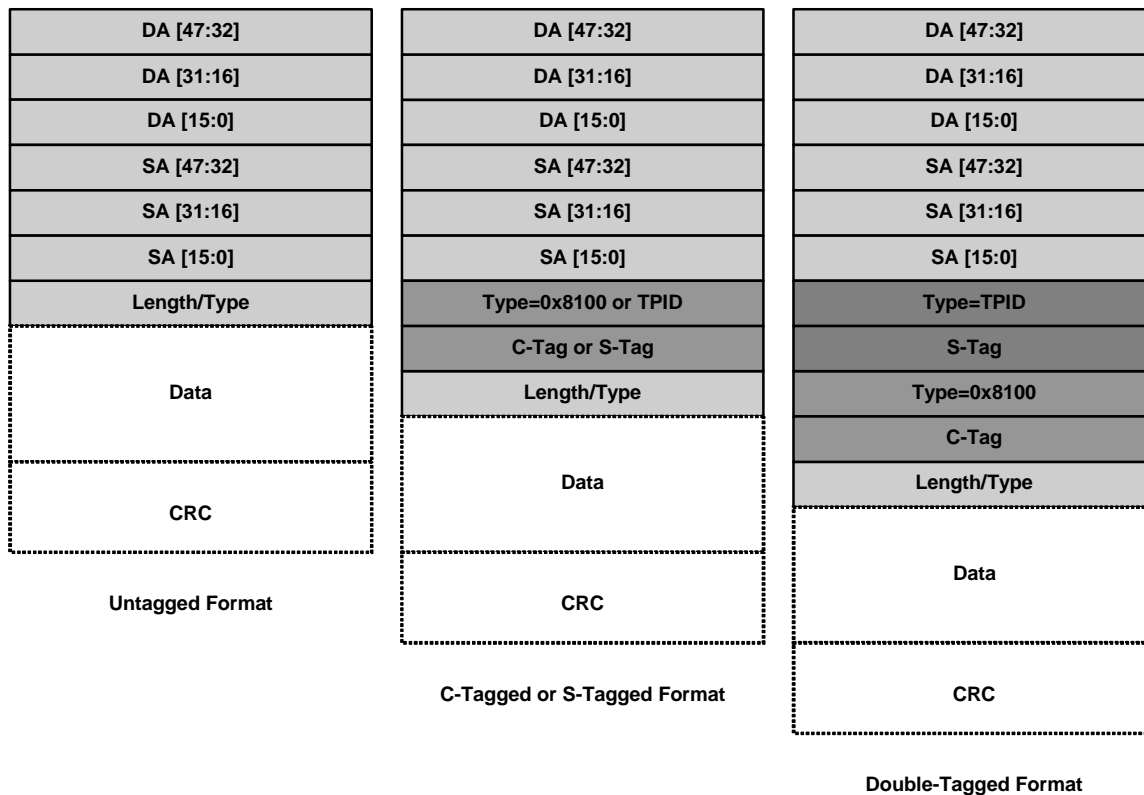
2.17.1.1 L2 Framing Structure Parsing

At Layer 2, the following framing structures are of concern to the parser design for the BCM53134P.

2.17.1.1.1 VLAN Tagging Structure Parsing

It is expected that the Ethernet packet received from a networking port can be of Untagged, Single-Tagged (C-Tagged or S-Tagged), or Double-Tagged format, as shown in [Figure 13](#).

Figure 13: VLAN Tagging Structures



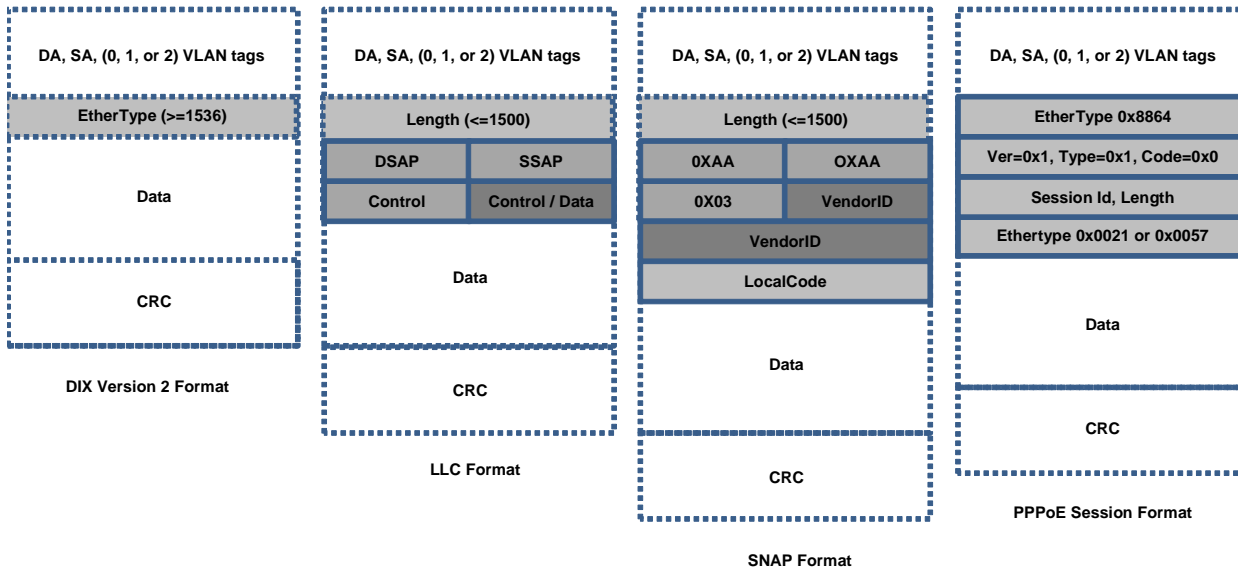
The TPID is globally programmable (defaults to 0x88A8) indicating the S-VLAN Tag. It is typically programmed differently from 0x8100 (which indicates C-Tag), but may be programmed as 0x8100 to indicate S-Tag.

When the TPID is programmed different from 0x8100, C-Tag and S-Tag can be explicitly differentiated. However, when the TPID is programmed as 0x8100, the differentiation between C-Tag and S-Tag relies on whether the port where the packet is received is programmed for user network interface (UNI) or service-provider network interface (NNI) operation: If the port is UNI, the first (outmost) tag received is considered as a C-Tag, otherwise it is considered as an S-Tag.

2.17.1.1.2 Ethernet Framing Structure Parsing

It is expected that packets received from a networking port can be of the following framing structures.

Figure 14: Ethernet Framing Structures



The DIXv2 (Ethernet II) format is the dominating Ethernet framing structure for IP traffic. The LLC format is the IEEE 802 standard framing structure mainly for OSI protocol stack support (such as SNA, NETBEUI). The SNAP format is the expanded version of LLC framing structure to support EtherType.

For the DIXv2 format, the parser extracts the EtherType field and continues the parsing operation for the higher layer protocol headers based on the EtherType field. In this format, the end of the EtherType field is considered the location for EndOfL2Header from which the upper layer parsing starts.

For the LLC format, the parser extracts the DSAP/SSAP field and continues the parsing operation for field extraction syntactically based on the corresponding UDF field definitions. In this format, the end of the SSAP field is considered the location for EndOfL2Header from which the upper layer parsing starts.

For the SNAP format, the parser extracts the VendorID and LocalCode continues the parsing operation for field extraction syntactically based on the corresponding UDF field definitions. However, if and only if the VendorID is zero, the LocalCode is treated as EtherType, and the parser continues the upper layer parsing as the parsing operation for the DIX v2 format. Therefore, the SNAP format is further divided as the following two categories: SNAP_Public format, which is used for encapsulation of standard EtherType based protocols, and SNAP_Private format, which is used for encapsulation of vendor specific protocols.

The parser can optionally parse PPPOE Session Stage packets (Ethertype = 0x8864) described in RFC2516 with:

- Version 1
- Type 1
- Code 0
- PPP Protocol 0x21 for IPv4, or 0x57 for IPv6

The parser can extract the IP header fields (including IP SA, DA, DSCP, IP Protocol, and so on) from IPv4/IPv6 packets encapsulated in the parsed PPPOE Session stage packets.

2.17.1.2 L3 Framing Structure Parsing


At layer 3, the following two framing structures are of concern to the parser design for the BCM53134P.

- IPv4 (qualified by preceding EtherType=0x0800)
- IPv6 (qualified by preceding EtherType=0x86DD)

2.17.1.2.1 IPv4 Header

IPv4 header is of variable size between 20 octets and 60 octets inclusive, in granularity of 4 octets, and has the following structures.

Figure 15: IPv4 Header Structure



Version: 4 for IPv4.

- IHL (Internet Header Length): Length of the IP header in 4-octet words. Minimum is 5 (indicating 20 octets); maximum is 15 (indicating 60 octets).
- TOS (Type of Service): Parameter fields specifying aggregate PHB. Its original TOS definitions have been superseded by the DSCP definitions.
- Total Length: Length of the IP datagram in octets, including IP header and data.
- Identification: An identifying value assigned by the sender to aid in assembling the fragments of a datagram. It must be unique for a source-destination pair and protocol for the time the datagram will be active in the Internet system.
- DF (Don't Fragment): Control for IP datagram fragmentation.
- MF (More Fragments): Indication of more fragments following the current one.
- Fragment Offset: Indication of where the fragment is located in the fragmented IP datagram, measured in units of 8-octets from the beginning of the original IP datagram.
- TTL (Time to Live): A timer value to track the lifetime of the datagram.
- Protocol: Indication of the next encapsulated protocol. If fragmentation occurs, this field will be same for all fragments.
- Header Checksum: A 16-bit one's complement checksum of the IP header.
- Source IP Address: 32-bit IPv4 address of the sender.
- Destination IP Address: 32-bit IPv4 address of the intended receiver(s).
- Options & Padding: A number of variable sized TLVs. Padding is used to keep 32-bit alignment for the IPv4 header.

2.17.1.2.2 IPv6 Header

The IPv6 (base) header is of 40 octets fixed length, and has the following structures.

Figure 16: IPv6 (Base) Header Structure

VER	Traffic Class	Flow Label	
Payload Length		Next Header	Hop Limit
Source IP Address [127:96]			
Source IP Address [95:64]			
Source IP Address [63:32]			
Source IP Address [31:0]			
Destination IP Address [127:96]			
Destination IP Address [95:64]			
Destination IP Address [63:32]			
Destination IP Address [31:0]			

Version: 6 for IPv6.

- Traffic Class: Parameter fields specifying aggregate PHB. It has the same definition as the TOS in IPv4.
- Flow Label: Experimental field to QoS treatment. Default value is zero, which means the packet does not belong to a specific flow.
- Payload Length: Length of the IPv6 payload in octets. All IPv6 extension headers following the IPv6 header are considered part of IPv6 payload.
- Next Header: Types of header immediately following the IPv6 header, its definition is same as Protocol field in IPv4 header. Some values indicate the type of IPv6 extension headers, some values indicate the type of higher layer protocols.
- Hop Limit: Same as TTL field in IPv4 header.
- Source IP Address: 128-bit IPv6 address of the sender.
- Destination IP Address: 128-bit IPv6 address of the intended receiver(s).

2.17.1.2.3 IPv6 Extension Header

Each IPv6 header may optionally be followed by one or multiple IPv6 extension header as part of its payload. Each type of IPv6 extension header is of integer number of 8 octets long, and has the following generic structures. Note that the AH extension header have different format in which the Hdr Ext Length is replaced by what the AH spec called Payload Length, and counting in 4-octet unit.

Figure 17: IP Extension Header Structure

Next Header	Hdr Ext Length	Ext Type Dependent Fields
Ext Type Dependent Fields && Padding		

- Next Header: Types of header immediately following the current IPv6 extension header,

- Hdr Ext Length: Length of IPv6 extension header in 8-octet unit, not including the first 8 octets. A zero value in this field indicates that the IPv6 extension header is of 8 octets long.
- Ext Type Dependent Fields: The syntax and semantics of these fields vary depending on the extension type indicated by the preceding Next Header field.

Following the IPv6 (base) header, the extension headers (if any) need to follow the following order:

- Hop-By-Hop Options (preceding Next Header = 0) header: it can appear only immediately after IPv6 header)
- Destination options (preceding Next Header = 60) header: it is used for options to be processed by all the destinations indicated by the IPv6 header and subsequent destinations listed in the Routing header.
- Routing (preceding Next Header = 43) header: it is used for source routing.
- Fragment (preceding Next Header = 44) header: It is used for end-to-end fragment transportation.
- Authentication (AH) (preceding Next Header = 51) header: It is used for IPsec.
- Encapsulating Security Payload (ESP) (preceding Next Header = 50) header: It is used for IPsec.
- Destination Options (preceding Next Header = 60) header: It is used for options to be processed only by the final destination of the packet.

Each extension header may occur at most once in an IPv6 packet, except for the Destination Options header which may occur at most twice: once before Routing header, once after Routing header. Not all extension headers listed above would appear in all IPv6 packets.

2.17.1.3 L4 Framing Structure Parsing

Above the IPv4/v6 layer, the parser supports limited semantical field parsing for the following higher layer protocols, without verifying the corresponding checksum.

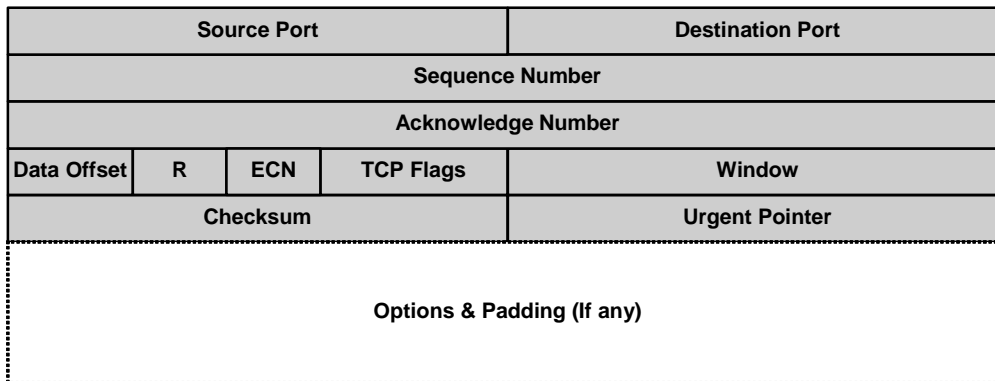
- TCP (preceding Next Header = 6)
- UDP (preceding Next Header = 17)
- UDPLite (preceding Next Header = 136)
- ICMPv4 (preceding Next Header = 1)
- IGMP (preceding Next Header = 2)
- ICMPv6 (preceding Next Header = 58)

The semantical parsing of these upper layer protocol headers are mainly for explicit switching behavior control. The slice key composition is based on syntactical field extraction based on UDF definitions.

2.17.1.3.1 TCP Header

The TCP header is of variable size between 20 octets and 60 octets inclusive, with granularity of 4 octets, and has the following structures.

Figure 18: TCP Header Structure



- Source Port: 16-bit source protocol interface.
- Destination Port: 16-bit destination protocol interface.
- Sequence Number: A parameter used for TCP protocol operations.
- Acknowledge Number: A parameter used for TCP protocol operations when ACK bit is set.
- Window: A parameter used for TCP protocol operations.
- Data Offset: The length the TCP header in 32-bit word unit. Minimum is 5 (no options), Maximum is 15.
- TCP Flags: 6 individual control bits for TCP protocol operations.
 - URG: Urgent Pointer field significant
 - ACK: Acknowledge field significant
 - PSH: Push Function
 - RST: Reset the connection
 - SYN: Synchronize sequence numbers
 - FIN: No more data from sender
- ECN: 3 individual bits used for Explicit Congestion Notification protocol operation (RFC 3168).
- Urgent Pointer: A parameter used for TCP protocol operations when URG bit is set.

- Checksum: 16-bit checksum calculated over pseudo header (including source/destination IP addresses, protocol, and length information), TCP header, and TCP data. Note that the length information used is different for IPv4 and IPv6 cases: IPv4 use Total Length, IPv6 use Upper Layer Packet Length, which excludes any extension headers.

2.17.1.3.2 UDP Header

The UDP header is of fixed size of 8 octets, and has the following structures.

Figure 19: UDP Header Structure

Source Port	Destination Port
Length	Checksum

- Source Port: 16-bit source protocol interface. It is set to 0 when not used.
- Destination Port: 16-bit destination protocol interface.
- Length: The length of UDP header and following UDP data. Minimum is 8 (UDP header only).
- Checksum: 16-bit checksum calculated over pseudo header (including source/destination IP addresses, protocol, and length information), UDP header, and UDP data. Note that the length information used is different for IPv4 and IPv6 cases: IPv4 uses Total Length, IPv6 uses Upper Layer Packet Length which excludes any extension headers.

2.17.1.3.3 UDPLite Header

The UDPLite header is of fixed size of 8 octets, and has the following structures.

Figure 20: UDPLite Header Structure

Source Port	Destination Port
Checksum Coverage	Checksum

- Source Port: 16-bit source protocol interface. It is set to 0 when not used.
- Destination Port: 16-bit destination protocol interface.
- Checksum Coverage: The number of bytes, counting from the first byte of the UDPLite header covered by the checksum. The UDPLite header **MUST** always be covered by the checksum. The value of 0 indicates that the entire UDPLite frame (header and data) is covered by the checksum (same as UDP).
- Checksum: 16-bit checksum calculated over pseudo header (including source/destination IP addresses, protocol, and length information), and the number of bytes specified by the Checksum Coverage. Note that the length information used is different for IPv4 and IPv6 cases: IPv4 uses Total Length, IPv6 uses Upper Layer Packet Length which excludes any extension headers.

2.17.1.3.4 ICMP/IGMP Headers

The IGMP, ICMPv4, and ICMPv6 headers all have the similar structures as follows.

Figure 21: ICMP/IGMP Message Structure

Type	Code	Checksum
Type/Code Dependent Data		

- Type: 8-bit field specifying type and format of the message.
- Code: 8-bit parameter for further qualification of the message.
- Checksum: 16-bit checksum to message protection. For IGMP/ICMPv4 message, the pseudo header is not included in the checksum calculation (different from that of TCP/UDP); for ICMPv6 message, the pseudo header is included in the checksum calculation (same as that of TCP/UDP).

2.17.1.4 User Defined Field (UDF) Extraction

Most packet fields that are used for the CFP rule searching key composition are extracted from each packet as UDFs, which are based on relative location specifications in order to provide flexibility of key composition for various networking applications.

2.17.1.4.1 UDF Offset Base Generation

The following four offset bases are used for UDF configurations. The value of each offset base varies depending on the parsing route the received packet goes through.

- StartOfPacket

This indicates the location where the MACDA starts.

- EndOfL2Header (following the VLAN tagging syntax checking)

If the frame is of DIXv2/Ethernet II or SNAP_Public framing, it indicates the location where the EtherType field ends.

If LLC framing or SNAP_Private, it indicates the location where the SSAP field ends. If the frame is a PPPoE Session, it indicates where PPP protocol field ends.

- EndOfL3Header

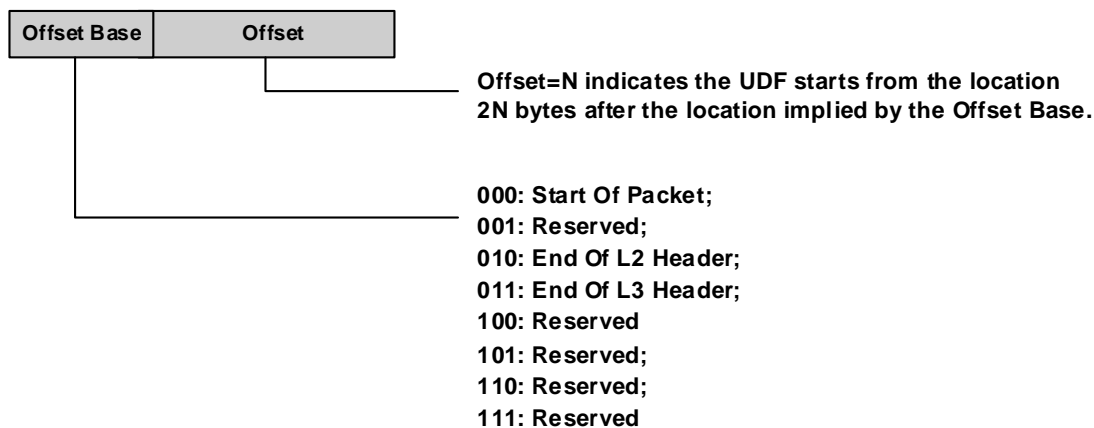
If the frame is of IPv4, it indicates the end of the IPv4 header (including options), except that when the IPv4 header is followed by an AH header, it is adjusted to indicate the end of the following AH header.

If the frame is of IPv6, it indicates the end of the IPv6 (base) header when there is no following extension header(s), or it indicates the end of last extension header the parser is able to parse when the IPv6 (base) header is followed by extension header(s).

2.17.1.4.2 UDF Specification

Each UDF is specified with 3-bit Offset Base value and the 5-bit Offset value relative to the specified Offset Base to indicate the start of the UDF fields. The granularity of the Offset value is 16-bit, and the length of each UDF is fixed as 16-bit.

Figure 22: UDF Configuration



The following show a few typical well-known fields expressed in UDF specifications.

- MACDA: Takes 3 UDFs, defined as {000, 00000}, {000, 00001}, and {000, 00010};
- MACSA: Takes 3 UDFs, defined as {000, 00011}, {000, 00100}, and {000, 00101};
- IPv4 SA: Takes 2 UDFs, defined as {010, 00110} and {010, 00111};
- IPv4 DA: Takes 2 UDFs, defined as {010, 01000} and {010, 01001};
- IPv6 SA: Takes 8 UDFs, defined as {010, 00100}, {010, 00101}, {010, 00110}, {010, 00111}, {010, 01000}, {010, 01001}, {010, 01010}, and {010, 01011};
- IPv6 DA: Takes 8 UDFs, defined as {010, 01100}, {010, 01101}, {010, 01110}, {010, 01111}, {010, 10000}, {010, 10001}, {010, 10010}, and {010, 10011};
- TCP/UDP SRC Port: Takes 1 UDF, defined as {011, 00000};
- TCP/UDP DST Port: Takes 1 UDF, defined as {011, 00001};
- ICMP Type & Code: Takes 1 UDF, defined as {011, 00000}.

2.17.2 Search Key Composition

The CFP rule searching engine supports up to four slices of searches; that is, a packet received is allowed to search the CFP rule set with up to four different keys. Slice0 has the option to be chained with the chain slice (Slice 3).

The key composition of the same slice can be programmed differently according to the received packet categories: IPv4 (unicast/multicast) packet, IPv6 (unicast/multicast) packet, or Non-IP packet. The UDFs used for different key composition are described as follows:

- UDF_n_A0, ... , UDF_n_A8: These UDFs are used by IPv4 packets for Slice n (n = 0, 1, or 2).
- UDF_n_B0, ... , UDF_n_B8: These UDFs are used by IPv6 packets for Slice n (n = 0, 1, or 2).
- UDF_n_C0, ... , UDF_n_C8: These UDFs are used by NonIP packets for Slice n (n = 0, 1, or 2).
- UDF_n_D0, ... , UDF_n_D11: These UDFs are used by IPv6 packet for the chain slice.

For the nonchain configuration of (Slice0, Slice1, Slice2, Slice3), Slice0 is of the lowest slice priority, and Slice3 is the highest slice priority. For the chain configuration of (Slice1, Slice2, {Slice0, Slice3}), Slice1 is of the lowest slice priority, and {Slice0, Slice3} as a chained slice is the highest slice priority.

The CFP rule searching key is 232 bits wide, and the formats of each search key composition are described as follows:

Table 5: Slice n Key for IPv4 Packet (n = 0, 1, 2)

Field	Width (bits)	Slice Bit Map	Description
Source_Port Map	8	231:224	The ingress port select to which the rule applies. To enable a key to a port or ports, the corresponding port mask bit should be set to 0, regardless of the setting state in the key field.
S_Tag_Status	2	223:222	00 = The packet was originally received without S Tag. 01 = The packet was originally received with SVID = 0. 10 = Reserved 11 = The packet was originally received with SVID = not 0.
C_Tag_Status	2	221:220	00 = The packet was originally received without C Tag. 01 = The packet was originally received with CVID = 0. 10 = Reserved 11 = The packet was originally received with CVID = not 0.
L2_Framing	2	219:218	00 = DIXv2 01 = SNAP Public
L3_Framing	2	217:216	00 = IPv4
IP_TOS	8	215:208	Type of Service field in IPv4 header
IP_Protocol	8	207:200	Protocol field in IPv4 header
IP_Fragmentation	1	199	0 = Not fragmented 1 = Fragmented
Non_First_Fragment	1	198	0 = first fragment 1 = Not first fragmented
IP_Authentication	1	197	0 = Not authenticated 1 = Authenticated

Table 5: Slice n Key for IPv4 Packet (n = 0, 1, 2) (Continued)

Field	Width (bits)	Slice Bit Map	Description
TTL_Range	2	196:195	00: TTL = 0 01: TTL = 1 10: TTL = Others 11: TTL = 255
PPPOE_Session	1	194	Identifies a PPPOE Session stage packet
Reserved	1	193	Default to zero
UDF_Valid[8]	1	192	Indicates whether UDF_n_A8 is valid.
UDF_Valid[7:0]	8	191: 184	Indicates whether UDF_n_A[7:0] is valid.
S-Tag	16	183: 168	The SVLAN tag is carried in the packet explicitly or is generated implicitly based on the default setting of the ingress port.
C-Tag	8	167: 160	The CVLAN tag is carried in the packet explicitly or is generated implicitly based on the default setting of the ingress port.
C-Tag	8 (LSB)	159: 152	
UDF_n_A8	16	151: 136	Must be validated by UDF_Valid[8].
UDF_n_A7	8	135: 128	Must be validated by UDF_Valid[7].
UDF_n_A7	8 (LSB)	127: 120	
UDF_n_A6	16	119: 104	Must be validated by UDF_Valid[6].
UDF_n_A5	8	103: 96	Must be validated by UDF_Valid[5].
UDF_n_A5	8 (LSB)	95: 88	
UDF_n_A4	16	87: 72	Must be validated by UDF_Valid[4].
UDF_n_A3	8	71: 64	Must be validated by UDF_Valid[3].
UDF_n_A3	8 (LSB)	63: 56	
UDF_n_A2	16	55: 40	Must be validated by UDF_Valid[2].
UDF_n_A1	8	39: 32	Must be validated by UDF_Valid[1].
UDF_n_A1	8 (LSB)	31:24	
UDF_n_A0	16	23: 8	Must be validated by UDF_Valid[0].
Reserved	4	7: 4	Default to zero.
Slice_ID	2	3: 2	Logical CFP rule used to search for slice identification: 00 = Slice 0 01 = Slice 1 10 = Slice 2 11 = Slice 3
Slice_Valid	2	1: 0	These bits are used to validate the corresponding slice and must be set to 2'b11.

Table 6: Slice n Key for IPv6 Packet (n = 0, 1, 2)

Field	Width (bits)	Slice Bit Map	Description
Source_Port Map	8	231: 224	The ingress port select to which the rule applies. To enable a key to a port or ports, the corresponding port mask bit should be set to 0, regardless of the setting state in the key field.
S_Tag_Status	2	223: 222	00 = The packet was originally received without S Tag. 01 = The packet was originally received with SVID = 0. 10 = Reserved 11 = The packet was originally received with SVID = not 0.
C_Tag_Status	2	221: 220	00 = The packet was originally received without CTag. 01 = The packet was originally received with CVID = 0. 10 = Reserved 11 = The packet was originally received with CVID = not 0.
L2_Framing	2	219: 218	00 = DIXv2 01 = SNAP Public
L3_Framing	2	217: 216	01 = IPv6
IP_Traffic Class	8	215: 208	IPv6 header Traffic Class field.
IP_Next Header	8	207: 200	Last parsed next header from the IPv6 header/extension header chain.
IP_Fragmentation	1	199	0 = Not fragmented 1 = Fragmented
Non_First_Fragment	1	198	0 = first fragment 1 = Not first fragmented
IP_Authentication	1	197	0 = Not authenticated 1 = Authenticated
Hop Limit Range	2	196: 195	00: Hop Limit = 0 01: Hop Limit TTL = 1 10: Hop Limit = Others 11: Hop Limit = 255
PPPOE_Session	1	194	Identifies a PPPOE Session stage packet received on a port that has the PPPOE_Session_Parsing_en configuration bit set.
Reserved	1	193	Defaults to zero
UDF_Valid[8]	1	192	Indicates whether UDF_n_B8 is valid.
UDF_Valid[7:0]	8	191: 184	Indicates whether UDF_n_B[7:0] is valid.
S-Tag	16	183: 168	The SVLAN tag is carried in the packet explicitly or is generated implicitly based on the default setting of the ingress port.
C-Tag	8	167: 160	The CVLAN tag is carried in the packet explicitly or is generated implicitly based on the default setting of the ingress port.
C-Tag	8 (LSB)	159: 152	
UDF_n_B8	16	151: 136	Must be validated by UDF_Valid[8].

Table 6: Slice n Key for IPv6 Packet (n = 0, 1, 2) (Continued)

Field	Width (bits)	Slice Bit Map	Description
UDF_n_B7	8	135: 128	Must be validated by UDF_Valid[7].
UDF_n_B7	8 (LSB)	127: 120	
UDF_n_B6	16	119: 104	Must be validated by UDF_Valid[6].
UDF_n_B5	8	103: 96	Must be validated by UDF_Valid[5].
UDF_n_B5	8 (LSB)	95: 88	
UDF_n_B4	16	87: 72	Must be validated by UDF_Valid[4].
UDF_n_B3	8	71: 64	Must be validated by UDF_Valid[3].
UDF_n_B3	8 (LSB)	63: 56	
UDF_n_B2	16	55: 40	Must be validated by UDF_Valid[2].
UDF_n_B1	8	39: 32	Must be validated by UDF_Valid[1].
UDF_n_B1	8 (LSB)	31: 24	
UDF_n_B0	16	23: 8	Must be validated by UDF_Valid[0].
Reserved	4	7: 4	Defaults to zero.
Slice_ID	2	3: 2	Logical CFP rule used to search for slice identification: 00 = Slice 0 01 = Slice 1 10 = Slice 2 11 = Slice 3
Slice_Valid	2	1: 0	These bits are used to validate the corresponding slice and must be set to 2'b11.

Table 7: Slice n Key for Non-IP Packet (n = 0, 1, 2)

Field	Width (bits)	Slice Bit Map	Description
Source_Port Map	8	231: 224	The ingress port select to which the rule applies. To enable a key to a port or ports, the corresponding port mask bit should be set to 0, regardless of the setting state in the key field.
S_Tag_Status	2	223: 222	00 = The packet is originally received without S Tag. 01 = The packet is originally received with SVID = 0. 10 = Reserved 11 = The packet is originally received with SVID = not 0.
C_Tag_Status	2	221: 220	00 = The packet is originally received without CTag. 01 = The packet is originally received with CVID = 0. 10 = Reserved 11 = The packet is originally received with CVID = not 0.

Table 7: Slice n Key for Non-IP Packet (n = 0, 1, 2) (Continued)

Field	Width (bits)	Slice Bit Map	Description
L2_Framing	2	219: 218	00 = DIXv2 01 = SNAP Public 10 = LLC 11 = SNAP Private
L3_Framing	2	217: 216	11 = Non IP
EtherType/SAP	16	215: 200	Ether type when L2 framing = DIXv2, or SNAP_Public {DSAP, SSAP} when L2 framing = LLD or SNAP_Private.
Reserved	5	199: 195	Defaults to zero.
PPPOE_Session	1	194	Identifies a PPPOE Session stage packet
Reserved	1	193	Defaults to zero.
UDF_Valid[8]	1	192	Indicates whether UDF_n_C8 is valid.
UDF_Valid[7:0]	8	191: 184	UDF_n_C[7:0] valid indication.
S-Tag	16	183: 168	The SVLAN tag is carried in the packet explicitly or is generated implicitly based on the default setting of the ingress port.
C-Tag	8	167: 160	The CVLAN tag is carried in the packet explicitly or is generated implicitly based on the default setting of the ingress port.
C-Tag	8 (LSB)	159: 152	
UDF_n_C8	16	151: 136	Must be validated by UDF_Valid[8].
UDF_n_C7	8	135: 128	Must be validated by UDF_Valid[7].
UDF_n_C7	8 (LSB)	127: 120	
UDF_n_C6	16	119: 104	Must be validated by UDF_Valid[6].
UDF_n_C5	8	103: 96	Must be validated by UDF_Valid[5].
UDF_n_C5	8 (LSB)	95: 88	
UDF_n_C4	16	87: 72	Must be validated by UDF_Valid[4].
UDF_n_C3	8	71: 64	Must be validated by UDF_Valid[3].
UDF_n_C3	8 (LSB)	63: 56	
UDF_n_C2	16	55: 40	Must be validated by UDF_Valid[2].
UDF_n_C1	8	39: 32	Must be validated by UDF_Valid[1].
UDF_n_C1	8 (LSB)	31: 24	
UDF_n_C0	16	23: 8	Must be validated by UDF_Valid[0].
Reserved	4	7: 4	Defaults to zero.
Slice_ID	2	3: 2	Logical CFP rule used to search for slice identification: 00 = Slice 0 01 = Slice 1 10 = Slice 2 11 = Slice 3 (Chain slice)

Table 7: Slice n Key for Non-IP Packet (n = 0, 1, 2) (Continued)

Field	Width (bits)	Slice Bit Map	Description
Slice_Valid	2	1: 0	These bits are used to validate the corresponding slice and must be set to 2'b11.

Table 8: Chain Slice Key for IPv6 Packet

Field	Width(bits)	Slice Bit Map	Description
Reserved	8	231: 224	Defaults to zero.
Chain ID	8	223: 216	The abstract ID found from slice 0 chained rule searching.
Reserved	4	215: 212	Defaults to zero.
UDF_Valid[11:8]	4	211: 208	Indicates whether UDF_n_D[11:8] is valid.
UDF_Valid[7:0]	8	207: 200	Indicates whether UDF_n_D[7:0] is valid.
UDF_n_D11	8	199: 192	Must be validated by UDF_Valid[11].
UDF_n_D11	8 (LSB)	191: 184	
UDF_n_D10	16	183: 168	Must be validated by UDF_Valid[10].
UDF_n_D9	8	167: 160	Must be validated by UDF_Valid[9].
UDF_n_D9	8 (LSB)	159: 152	
UDF_n_D8	16	151: 136	Must be validated by UDF_Valid[8].
UDF_n_D7	16	135: 128	Must be validated by UDF_Valid[7].
UDF_n_D7	16	127: 120	
UDF_n_D6	16	119: 104	Must be validated by UDF_Valid[6].
UDF_n_D5	8	103: 96	Must be validated by UDF_Valid[5].
UDF_n_D5	8 (LSB)	95: 88	
UDF_n_D4	16	87: 72	Must be validated by UDF_Valid[4].
UDF_n_D3	8	71: 64	Must be validated by UDF_Valid[3].
UDF_n_D3	8 (LSB)	63: 56	
UDF_n_D2	16	55: 40	Must be validated by UDF_Valid[2].
UDF_n_D1	8	39: 32	Must be validated by UDF_Valid[1].
UDF_n_D1	8 (LSB)	31: 24	
UDF_n_D0	16	23: 8	Must be validated by UDF_Valid[0].
Reserved	4	7: 4	Defaults to zero.
Slice_ID	2	3: 2	Logical CFP rule used to search for slice identification: 00 = Slice 0 01 = Slice 1 10 = Slice 2 11 = Slice 3 (Chain slice)
Slice_Valid	2	1: 0	These bits are used to validate the corresponding slice and must be set to 2'b11.

2.17.3 Action Resolution

2.17.3.1 Policy Action Definitions

CFP actions based on the resolved rule (rule that was selected by the resolution procedure) and are described in [Table 9](#). Suffix IB indicates action taken if the packet is in-band (within policing profile) and suffix OB indicates action if the packet is out-of-band (outside policing profile).

Table 9: CFP Actions

Field	Bit Width	Description
Change_DSCP_OB	1	It indicates whether to modify the IP DSCP field of the matched packet based on the New_DSCP value.
New_DSCP_OB	6	New_DSCP value. (In IPv4 header, this field is called TOS field, and the IP checksum field needs to be updated accordingly. In IPv6 header, this field is called TrafficClass field, and there is no IP checksum to be updated)
Change_DSCP_IB	1	It indicates whether to modify the IP DSCP field of the matched packet based on the New_DSCP value.
New_DSCP_IB	6	New_DSCP value. (In IPv4 header, this field is called TOS field, and the IP checksum field needs to be updated accordingly. In IPv6 header, this field is called TrafficClass field, and there is no IP checksum to be updated)
Change_FWD_IB	2	It indicates whether to enforce new egress direction for the matched packet. Changes ARL forwarding bitmap when meeting is in-band or conforming. 00: No change to the ARL FWD_MAP. 01: Remove ports in DST_MAP_IB (AND) from ARL FWD_MAP. 10: Replace ARL FWD_MAP with DST_MAP_IB 11: Add the DST_MAP_IB (OR) to ARL FWD_MAP
DST_MAP_IB	8	It indicates the port(s) to which the packet is forwarded. port 8 (IMP0), port 7 (CPU) port 5 (IMP1), port 4-0
CHANGE_FWD_OB	2	Changes ARL forwarding bitmap when meeting is out-of-band or not conforming. 00: No change to the ARL FWD_MAP. 01: Remove ports in DST_MAP_OB (AND) from ARL FWD_MAP. 10: Replace ARL FWD_MAP with DST_MAP_OB 11: Add the DST_MAP_OB (OR) to ARL FWD_MAP
DST_MAP_OB	8	It indicates the port(s) to which the packet is forwarded. port 8 (IMP0), port 7 (CPU) port 5 (IMP1), port 4-0
ChainID	8	If it is the result of Slice 0 chained search, it indicates the ChainID to be used as part of Chain slice key. 0x00 indicates no valid ChainID. Otherwise, it indicates the Classification ID if the packet needs to be forwarded to CPU. 0x00 indicates no valid Classification ID.
Change_TC	1	It indicates whether to enforce new traffic class for the matched packet to be queued with the corresponding COS at its egress Ethernet port(s) (excluding IMP port) before being transmitted.

Table 9: CFP Actions (Continued)

Field	Bit Width	Description
New_TC	3	If the Change_TC action is chosen for a packet matching the CFP rule, then this field indicates the new Traffic Class to be used for determining the queue in which the packet will be stored; the color of the packet (unless overridden by the Change_Color action of the CFP); and the RED/WRED profile of the packet.
Change_TC_O	1	It indicates whether the new traffic class (New_TC_O) for the matched packet to be saved in the queue at its egress Ethernet port(s) <i>instead of</i> the TC that was used for determining the packet's color, COS, and RED/WRED profile. The saved TC_O in the packet is used for optionally re-mark a packet's PCP and DEI before it is transmitted. If software configures a CFP rule to override TC for packets destined to an IMP port (port 5 or 8 in IMP mode), then it MUST configure (CFP.New_TC = CFP.New_TC_O) and (CFP.Change_TC = CFP.Change_TC_O).
New_TC_O	3	If the Change_TC_O action is chosen for a packet matching the CFP rule, then this field indicates the new Traffic Class to be used for determining the PCP and DEI of a packet after it is scheduled for transmission on an Egress Ethernet or an IMP port. If software configures a CFP rule to override TC for packets destined to an IMP port (port 5 or 8 in IMP mode), then it MUST configure (CFP.New_TC = CFP.New_TC_O) and (CFP.Change_TC = CFP.Change_TC_O).
SPCP_remark_disable CPCP_remark_disable DEI_remark_disable	3	It indicates whether the SPCP, CPCP, and DEI fields in the packet should be remarked at the egress port. if the per-port remarking is enabled. If set, this per-flow configuration disables remarking of SPCP, CPCP, and DEI fields in the packet even when the per-port PCP remark enable configuration bit enables the SPCP, CPCP, and DEI fields remarking function for the port. However, for backward compatibility, this CFP control bit is ignored for CFI or DEI remarking if the cfi_remark_enable[port] control is set by software.
LoopBack_Enable	1	It indicates whether the packet is allowed to be forwarded to the port it is originally received from.
Reason_Code	6	It indicates the reasons why the packet is forwarded to CPU, when the corresponding Change_FWD action indicates packet forwarding to CPU.
STP_Bypass	1	It indicates whether the CFP generated forwarding decision is subject to the STP port state based filtering.
EAP_Bypass	1	It indicates whether the CFP generated forwarding decision is subject to the 802.1x EAP port state based filtering.
VLAN_Bypass	1	It indicates whether the CFP generated forwarding decision is subject to the VLAN based filtering.
RED_DEFAULT	1	Use the RED default profile
Change_COLOR	1	Force the input color of the matched packet of the Flow Policer
New_Color	2	The input color of the matched packet Flow Policer
MAC_Limit_Bypass	1	If the MAC Address Limit feature is enabled on a port and the DROP_MAC_Limit is set, then the MAC_Limit_Bypass action will override the drop decision because of the MAC address limit.

2.17.3.2 Metering/Statistics Selection

Since the bandwidth of the CFP metering/statistics mechanism design can allow only one access by each packet, the selection of CFP meter and counter is subject to the following restrictions.

- In case of multiple slice matches, only one slice can have the meter actions, and only the counter associated with the selected meter will accumulate the byte counts. Among the matched slices, the meter with highest slice priority should be selected for rate metering.
- For a matched slice whose meter is selected for rate metering, the meter independent action selection is not affected by the metering results, but the meter dependent action selection from the slice is affected by the metering results (in-band versus out-band).
- For a matched slice whose meter is not selected for rate metering, the meter independent action selection is not affected by the metering results from the selected slice, and the meter dependent actions from the slice are selected from the in-profile action set.

The Flow Policer (meter) can be configured to function as one of the following:

- RFC2698 compliant two rate three color (trTCM) policer
- RFC4115 compliant two rate three color (trTCM) policer
- MEF (MEF 6.1, MEF 10.2) compliant rate meter
- Pass through with the metering function disabled

A *Flow* is identified by a set of parameters that can be matched against CFP rules. The index of the matching CFP rule (127..0) then indexes directly into one of the 128 Flow Policers.

The metering and marking function is used to measure the ingress byte rate of a traffic flow and compare it to one or two rate limits. The measured data rate compared to the rate limits is used to *mark* each packet with a *color*.

- Green marks packets belonging to conforming flows that don't exceed the lower limit.
- Yellow marks packets whose flow exceeds the lower limit but not the upper limit.
- Red marks packets whose flows exceed the upper limit.

A packet's color is used to determine whether or not it should be discarded prior to queuing.

- If the ingress packet is marked as Green:
 - The PCP and DCSP of the packet can be optionally *marked down* or changed.
 - A different RED/WRED profile can be selected for more aggressive packet drops during egress queue congestion.
 - Counted in a 32-bit switch statistics for Green packets.
- If the ingress packet is marked Yellow:
 - The PCP and DCSP of the packet can be optionally *marked down* or changed.
 - A different RED/WRED profile can be selected for more aggressive packet drops during egress queue congestion.
 - Counted in a 32-bit switch statistics for Yellow packets.
- If a packet is marked with Red color, it is dropped in the switch and counted separately.

By default, marking down of DSCP in IPv4/IPv6 packets encapsulated in PPPOE Session Stage version 1 and type 1 packet is disabled through a global configuration register. Software can enable marking down of the DSCP field by setting the bit.

2.18 Multiple Spanning Tree Protocol

The BCM53134P support up to eight multiple spanning trees. When the EN_RX_BPDU bit = 1, the BCM53134P forwards BPDU packets to the management port only.

2.19 Software Reset

The BCM53134P provides Software Resets. Software Resets can be triggered by setting the register (page 00h, address 79h).

NOTE: Software Reset sets all the register and the table (ARL, VLAN) contents to the default values. Software Reset will not latch in the strap pin values, but the previous latched strap pin values are retained.

2.20 BroadSync HD

BroadSync HD is the enhancement to IEEE 802.3 MAC and IEEE 802.1D bridges to support the kind of low-latency isochronous services and guaranteed QoS that is required for many consumer electronics applications.

BCM53134P provides the BroadSync HD feature. BCM53134P always forwards BPDU, MRP packets to CPU for BroadSync HD applications, and handle IEEE 802.1 Time Sync Protocol.

The BCM53134P can identify a packet as a BroadSync HD packet if MAC DA matches with the programmed (registered and configured based on MRP protocols) MAC address. A MAC address can be a multicast or unicast address. The PCP of the incoming BroadSync HD packet can be any value that is programmed. The default values are 4 and 5. There are two dedicated queues for BroadSync HD Class 5 and Class 4 traffic per egress port. The BCM53134P enhances shaping and scheduling for BroadSync HD operation.

2.20.1 Time Base and Slot Generation

For BroadSync HD applications, the BCM53134P maintains a time base (32-bit counter) running at a granularity of 1 ns, which can be adjusted by CPU for synchronization with the BroadSync HD time master unit (Switch or Host) through the IEEE 802.1 Time Synchronized (TS) protocol (to be standardized). The TS protocol is implemented by the CPU, which requires the BCM53134P to perform the following operations.

- A received TS protocol packet is timestamped at the ingress port when the first byte (of MACDA) arrives, and is transferred along with the receiving timestamp to the CPU.
- A TS protocol packet initiated by the CPU (to be transmitted at an egress port) is timestamped at the egress port when the first byte (of MACDA) is transmitted, and the transmit timestamp recorded at the egress port is reported back to CPU.

It is required that the time synchronization point peers over an Ethernet link is chosen such that the link delay is perceived as constant, and the protocol exchange occurs at least every 10 ms over every link.

The CPU may be required to speed up or slow down the timebase maintained in BCM53134P based on the TS protocol execution. The BCM53134P provides the time base adjustment mechanism for graceful time changes based on CPU instructions.

In addition, the BCM53134P maintains counter mechanism to generate time slot for BroadSync HD traffic scheduling.

- A Slot is defined as 125 μ s, it is used to pace the BroadSync HD Class 5 traffic which has tight jitter requirements;
- A MacroSlot is configurable as 1 ms, 2 ms, or 4 ms (binary number of Slots). It is used to pace the BroadSync HD Class 4 traffic, which has relaxed jitter requirements.

The CPU may be required to make the slot wider or narrower based on the TS protocol execution. The BCM53134P provides the slot adjustment mechanism for graceful Slot width changes based on CPU instructions.

2.20.2 Transmission Shaping and Scheduling

2.20.2.1 BroadSync HD Class5 Media Traffic

The CoS5 queue is dedicated for BroadSync HD Class 5 traffic only, and a CoS5 packet is always the highest priority to be scheduled for transmission. It is allowed by the Shaper A that operates as follows.

- The Shaper A is an emulation of fixed-bandwidth pipe for Class 5 BroadSync HD traffic with tight jitter-to-handle interference from non-BroadSync HD or Class 4 BroadSync HD traffic, adaptively. Note that the preamble and IPG transmission are not taken into account for the pipe operation.
- Tunable parameters for the Shaper A are listed as follows.
 - MaxAVPacketSize indicates the maximum packet size allowed on an AV-Enabled port. It is a global setting using BroadSync HD Max Packet Size register.
 - Class5_BW indicates the reserved bandwidth for Class 5 BroadSync HD traffic at granularity of Byte (per slot, 125 μ s). It is a per-port setting.
 - Class5_Window indicates the jitter control for Class5 BroadSync HD transmission. It is a per-port setting.
- At the start of each Slot:
 - Reset the credit in the shaping bucket to Class5_BW, if the queue is empty.
 - Reset the credit in the shaping bucket to Class5_BW, if the queue is not empty and Class5_Window is set to 0.
 - Reset the credit in the shaping bucket to Class5_BW, if the queue is not empty, Class5_Window is set to 1, and the credit remained in the shaping bucket is greater than MaxAVPacketSize.
 - Add Class5_BW to the credit in the shaping bucket, if the queue is not empty, Class5_Window is set to 1, and the credit remained in the shaping bucket is less than or equal to MaxAVPacketSize.
- The credit in the shaping bucket decrements for every byte transmitted for the Class 5 BroadSync HD traffic through the port.
 - If the credit reaches 0 before the end of the current slot while transmitting a Class 5 BroadSync HD packet, the ongoing packet transmission is not interrupted and the credit stays at 0 until being reset at the start of next slot.
 - The credit decrements resumes at the next Slot if the ongoing transmission continues.
- As long as the credits in the shaping bucket is greater than 0, a Class 5 BroadSync HD packet is allowed to be scheduled for transmission.

2.20.2.2 BroadSync HD Class4 Media Traffic

The CoS4 queue is dedicated for BroadSync HD Class 4 traffic only, and a CoS4 packet always yields to CoS5 traffic (if allowed to be scheduled), but takes precedence over the traffic from CoS0~CoS3 queues or follows the weight ratio between CoS4 and CoS0~CoS3 for transmission scheduling, if it is allowed by the Shaper B that operates as follows.

- The Shaper B is an emulation of fixed bandwidth pipe for Class 4 BroadSync HD traffic with relaxed jitter to handle interference from non-BroadSync HD or Class 5 BroadSync HD traffic adaptively. It also statistically levels the Class 4 BroadSync HD transmission bursts towards the next hop switch to reduce the buffering requirements by using slot (instead of MacroSlot) as the pacing mechanism. The preamble and IPG transmission are not accounted for in the pipe operation.
- Tunable parameters for the Shaper B are listed as follows:
 - MacroSlot_Period indicates the periodic cycle time to shape the Class 4 traffic. It is a global setting to indicate 1 ms, 2 ms, or 4 ms.
 - MaxAVPacketSize indicates the maximum packet size allowed on an AV-Enabled port. It is a global setting. (same as for BroadSync HD Class 5 setting)
 - Class4_BW indicates the evenly divided bandwidth share per Slot, which is derived from dividing the reserved bandwidth for Class 4 BroadSync HD traffic at granularity of Byte (per MacroSlot) by the number of slots within a MacroSlot. It is a per-port setting.
- At the start of each slot,

- If the slot is the first one for the current MacroSlot, reset the credit bucket to Class4_BW+MaxAVPacketSize; (MaxAVPacketSize is used as the deficit base)
- Otherwise, add Class4_BW to the credit in the shaping bucket.
- The shaping credit bucket decrements for every byte transmitted for the Class 4 BroadSync HD traffic.

As long as the credits in the shaping bucket is greater than or equal to MaxAVPacketSize, a Class 4 BroadSync HD packet is allowed to be scheduled for transmission

2.21 CableChecker

The BCM53134P provides cable diagnostic capabilities for unmanaged environments. The actual cable diagnostic feature lies in the PHY functional block. The BCM53134P lets the user monitor the cable diagnostic results through LED display by setting the appropriate bits in the LED refresh registers.

The BCM53134P uses the existing LED display (which is already assigned to various functions) to indicate the cable diagnostic results. Table 10 shows the cable diagnostic result output for each LED function where 1 and 0 represent the LED indication pin status; 1 indicates active and 0 indicates inactive.

NOTE:

- The best way for a user to visualize the cable diagnostic test result through LEDs is to bring out the LINK status bit to the LED display along with other functions to be displayed per port. In this way, the user can observe the cable diagnostic result from the flashing (or lit) LED of other functions while LINK LED is off. The switch turn off the LINK status LED during the cable diagnostic mode.
- The cable diagnostic is expected to be most effective when the user cannot establish the link with the partner.

Table 10: Cable Diagnostic Output

LED Function in LED Function Register	Cable Diagnostic Output
PHYLED4	1 = Cable diagnostic failed 0 = Cable diagnostic passed
LNK	No output during the cable diagnostic mode
DPX	1 = Passed 0 = Failed
ACT	1 = Passed 0 = Failed
COL	1 = Passed 0 = Failed
LNK/ACT	No output during the cable diagnostic mode
DPX/COL	1 = Passed 0 = Failed
SPD10M	1 = Failed 0 = Passed
SPD100M	In LED function0 map 1 = Cable diagnostic passed 0 = Failed In LED function1 map 1 = Cable diagnostic failed 0 = Passed

Table 10: Cable Diagnostic Output (Continued)

LED Function in LED Function Register	Cable Diagnostic Output
SPD1G	1 = Passed 0 = Failed
10M/ACT	1 = Failed 0 = Passed
100M/ACT	In LED function0 map 1 = Cable diagnostic passed 0 = Failed In LED function1 map 1 = Cable diagnostic failed 0 = Passed
10–100M/ACT	1 = Failed 0 = Passed
1G/ACT	1 = Passed 0 = Failed
PHYLED3	1 = Failed 0 = Passed

2.22 Egress PCP Remarking

The BCM53134P provides an egress PCP remarking feature of the outer tag at each egress port which includes the CFI and PCP field modification based on the internal ARL/CFP-generated TC and CFP-generated Rate Violation (RV) status. The Egress PCP remarking process applies to Ethernet ports only. Each Ethernet port can provide a 16-entry mapping table indexed by {RV, TC} to map to the {New CFI, New PCP} field for the outgoing packet.

NOTE: For the AV-enabled egress port, the egress PCP for the non-BroadSync HD class of traffic must never be programmed with values of 100 and 101.

2.23 Address Management

The BCM53134P Address Resolution Logic contains the following features:


- Four bins per bucket address table configuration.
- Hashing of the MAC/VID address to generate the address table point.

The address management unit of the BCM53134P provides wire speed learning and recognition functions. The address table supports 4K unicast/multicast addresses using on-chip memory.

2.23.1 Address Table Organization

The MAC addresses are stored in embedded SRAM. Each bucket contains four entries or bins. The address table has 1K buckets with four entries in each bucket. This allows up to four different MAC addresses with the same hashed index bits to be simultaneously mapped into the address table. In the ARL DA/SA lookup process, it hashes a 10-bit search index and read out bin0 and bin1 in the first cycle, and read out bin2 and bin3 in the second cycle. These four entries are used for ARL routing and learning.

Figure 23: Address Table Organization



The index to the address table is computed using a hash algorithm based on the MAC address and the VLAN ID (VID) if enabled.

NOTE: In the Enable IEEE 802.1Q and VLAN Learning Mode both the MAC address and the VLAN ID (VID) are used to compute the hashed index. See [“IEEE 802.1Q VLAN” on page 19](#) for additional information.

The hash algorithm uses the CRC-CCITT polynomial. The input to the hash is reduced to a 16-bit CRC hash value. Bits [9:0] of the hash are used as an index to the approximately 4K locations of the address table.

The CRC-CCITT polynomial is:

$$x^{16} + x^{12} + x^5 + 1$$

2.23.2 Address Learning

Information is gathered from received unicast packets and learned or stored for the future purpose of forwarding frames addressed to the receiving port. During the receive process, the frame information (such as the Source Address [SA] and VID) is saved until completion of the packet. An entry is created in the ARL table memory if the following conditions are met:

- The packet has been received without error.
- The packet is of legal length.
- The packet has a unicast SA.
- If using IEEE 802.1Q VLAN, the packet is from an SA that belongs to the indicated VLAN domain.
- The packet does not have a reserved multicast destination address. This condition can be disabled using register settings.
- There is free space available in memory to which the hashed index points.

When unicast packets are dynamically learned, the VALID bit is set, the AGE bit is set, and the STATIC bit is cleared in the entry. See [Table 12 on page 57](#) for a description of a unicast ARL entry.

Multicast addresses are not learned into the ARL table, but must be written using one of the [“Programming Interfaces” on page 85](#).

2.23.3 Address Resolution and Frame Forwarding

Received packets are forwarded based on the information learned or written into the ARL table. Address resolution is the process of locating this information and assigning a forwarding destination to the packet. The destination address (DA) and VID of the received packet are used to calculate a hashed index to the ARL table. The hashed index key is used by the address resolution function to locate a matching ARL entry. The frame is assigned a destination based on the forward field (PORTID or IPMC0) of the ARL entry. If the address resolution function fails to return a matching ARL entry, the packet is flooded to all appropriate ports. The following two sections describe the specifics of address resolution and frame forwarding for [“Unicast Addresses” on page 57](#) and [“Multicast Addresses” on page 57](#).

2.23.3.1 Unicast Addresses

Frames containing a unicast destination address are assigned a forwarding field corresponding to a single port. The unicast address-resolution algorithm is as follows:

- If the multiport addressing feature is enabled and the DA matches one of the programmed multiport addresses, then it is forwarded accordingly. See [“Power Savings Modes” on page 61](#).
- The lower 10 bits of the hashed index key are used as a pointer into the address table memory, and the entry is retrieved.
- If the valid indicator is set and the address stored at one of the locations matches the index key of the packet received, the forwarding field port ID is assigned to the destination port of the packet.
 - If the destination port matches the source port, the packet is not forwarded.
- If the address resolution function fails to return a matching valid ARL entry and the unicast DLF forward bit is set, the frame is forwarded according to the port map.
- Otherwise, the packet is flooded to all appropriate ports.

See [Table 11](#) for definitions of the unicast index key and the assigned forwarding field. The forwarding field for a unicast packet is the port ID contained in the matching ARL entry. See [Table 12](#) for a description of a unicast ARL entry.

Table 11: Unicast Forward Field Definitions

EN_1QVLAN	Index Key	Forwarding Field
1	DA and VID	Port ID
0	DA	Port ID

Table 12: Address Table Entry for Unicast Address

Field	Description
VID	VLAN ID associated with the MAC address
VALID	1 = Entry is valid. 0 = Entry is empty.
STATIC	1 = Entry is static—Should not be aged out and is written and updated by software. 0 = Entry is dynamically learned and aged.
AGE	1 = Entry has been accessed or learned since last aging process. 0 = Entry has not been accessed since last aging process.
TC	MACDA-based TC (valid only for static entries). See “Quality of Service and Scheduling” on page 11 for more information.
Reserved	—
Reserved	Only 00 is valid.
PortID	Port identifier. The port associated with the MAC address.
MAC ADDRESS	48-bit MAC address.

Multicast ARL table entries are described in [Table 13 on page 58](#).

2.23.3.2 Multicast Addresses

Frames containing a multicast destination address are assigned a forwarding field corresponding to multiple ports specified in a port map. If no matching ARL entry is found, the packet is flooded to all appropriate ports.

The multicast address resolution algorithm is as follows:

- If the DA matches one of the globally assigned reserved addresses between 01-80-C2-00-00-00 and 01-80-C2-00-00-2F, the packet is handled as described in [Table 14 on page 59](#).
- If the multiport addressing feature is enabled and the DA matches one of the programmed Multiport Addresses, then it is forwarded accordingly.
- Otherwise, the lower 10 bits of the hashed index key are used as a pointer into the ARL table memory, and the entry is retrieved.
- If the valid indicator is set, and the address stored at the entry locations matches the index key of the packet received, the forwarding field port map is assigned to the destination port of the packet.
- If the address resolution function fails to return a matching valid ARL entry and the multicast DLF forward bit is set, the frame is forwarded according to the port map. There are two types of Multicast address groups handled by two separate Lookup Fail Forward registers.
- Otherwise, all other multicast and broadcast packets are flooded to all appropriate ports.

See [Table 13](#) for a description of a multicast ARL entry. See [“Accessing the ARL Table Entries” on page 60](#) for more information.

Table 13: Address Table Entry for Multicast Address

Field	Description
VID	VLAN ID associated with the MAC address
VALID	1 = Entry is valid. 0 = Entry is empty.
STATIC	1 = Entry is static— This entry is not aged out and is written and updated by software. 0 = Not defined.
AGE	AGE bit is ignored for static ARL table entries.
TC	MACDA-based TC (valid only for static entries). See “Quality of Service and Scheduling” on page 11 for more information.
Reserved	—
IPMC0 [8:0]	Multicast forwarding mask 1 = Forwarding enable. 0 = Forwarding disable.
MAC ADDRESS	48-bit MAC address.

Unicast ARL table entries are described in [Table 12 on page 57](#).

2.23.3.3 Reserved Multicast Addresses

[Table 14](#) summarizes the actions taken for specific reserved multicast addresses. Packets identified with these destination addresses are handled uniquely since they are designed for special functions.

Table 14: Behavior for Reserved Multicast Addresses

MAC Address	Function	IEEE 802.1 Specified Action	Unmanaged Mode Action	Managed Mode Action
01-80-C2-00-00-00	Bridge group address	Drop frame	Flood frame	Forward frame to IMP only.
01-80-C2-00-00-01	IEEE 802.3x MAC control frame	Drop frame	Receive MAC determines if it is a valid pause frame and then acts accordingly	Receive MAC determines if valid pause frame and acts accordingly.
01-80-C2-00-00-02	Reserved	Drop frame	Drop frame	Forward to frame management port only.
01-80-C2-00-00-03	IEEE 802.1X port-based network access control	Drop frame	Drop frame	Forward frame to management port only.
01-80-C2-00-00-04– 01-80-C2-00-00-0F	Reserved	Drop frame	Drop frame	Forward frame to management port only.
01-80-C2-00-00-10	All LANs bridge management group address	Forward frame	Flood frame	Forward frame to all ports including management port.
01-80-C2-00-00-11– 01-80-C2-00-00-1F	Reserved	Forward frame	Flood frame	Forward frame to all ports excluding management port.
01-80-C2-00-00-20	GMRP address	Forward frame	Flood frame	Forward frame to all ports excluding management port, or forward frame to management port only (by setting bit 4 of page 34, offset 04h register).
01-80-C2-00-00-21	GVRP address	Forward frame	Flood frame	Forward frame to all ports excluding management port, or forward frame to management port only (by setting bit 5 of page 34, offset 04h register).
01-80-C2-00-00-22– 01-80-C2-00-00-2F	Reserved	Forward frame	Flood frame ^a	Forward frame to all ports excluding management port.

a. Frames flood to all ports. Certain exclusions apply, such as VLAN restrictions.

NOTE: The managed mode is set through the register Page 00h, Address 0Bh, bit 0, and a management port can be selected through the register Page 02h, Address 00h. The management (IMP) port often has to be forced to link through the IMP Port State Override register Page 00h, Address 0Eh.

2.23.4 Static Address Entries

The BCM53134P supports static ARL table entries that are created and updated using one of the [“Programming Interfaces” on page 85](#). These entries can contain either unicast or multicast destinations. Static entries do not automatically learn MAC addresses or port associations and are not aged out by the automatic internal aging process.

2.23.5 Accessing the ARL Table Entries

ARL table entries are accessed by one of two mechanisms. The first mechanism uses the ARL read/write control, which allows an address-entry location to be read, modified, or written based on the value of a known MAC address. The second mechanism searches the ARL table sequentially, returning all valid entries.

2.23.5.1 Searching the ARL Table

The second method to access the ARL table is through the ARL search control. The entire ARL table is searched sequentially, revealing each valid ARL entry.

The ARL search and ARL read/write operations execute in parallel with other register accesses. This allows the host processor to start a read, write, or search process and then read/write other registers, returning periodically to see if the operation has completed.

2.23.6 Address Aging

The BCM53134P offers two types of Address Aging processes: Normal Aging and Fast Aging. When the address entries are marked as *static* in the ARL table, the Normal Aging process does not apply. In this case, the address entries can only be removed by the user, by Fast-Aging, or by RESET.

2.23.6.1 Normal Aging

The Normal Aging process removes the dynamically learned address entries in the ARL table if the entries are not accesses (no packets are addressing the MAC address) for the specified period of time in the Aging Time Control Register (page 02h, address 06h). Users can specify the aging time. Otherwise, the entries will age out by the default time of 300 seconds.

Users can also accelerate the Aging time by enabling the AGE_Accelerate bit (Page 04h, address 00h, bit 2). When this bit is set, the Aging time is reduced to 1/128 of the current Aging time.

2.23.6.2 Fast Aging

The Fast Aging process removes the address entries in the ARL table selectively specified by the Fast Aging Control Register (page 00h, address 88h). The fast aging function can be enabled per port or per VLAN ID. Users can select and specify the type of entries to be removed in the Fast-Aging Control Register (page 00h, address 88h).

- Multicast entries: multicast cannot be selected at the same time with Port ID selection.
- Spanning Tree ID
- VLAN ID: specified in the Fast-Aging VID Control Register (page 00h, address 8Ah).
- Port ID: specified in the Fast-Aging Port Control Register (page 00h, address 89h).
- Dynamic entries
- Static entries

Users can trigger the Fast-Aging process by setting the Fast-Age-Start bit in the Fast Aging Control Register (page 00h, address 88h). Once the Fast-Aging process is triggered, the process goes through each entry in the ARL table and removes the selected type of entry. Entries stored in the lower address are removed first.

2.24 Power Savings Modes

The BCM53134P offers different power savings modes for different operating states. All the power saving scheme are implemented without any external CPU requirement.

The various power savings modes are:

- **Auto Power-Down mode:** This is a stand-alone PHY feature that is enabled by a register bit setting. The PHY shuts off the analog portion of the circuitry when cable is not connected or the link partner power is down.
- **Energy Efficient Ethernet (EEE) mode:** Energy Efficient Ethernet is IEEE 802.1az, an extension of the IEEE 802.3 standard. IEEE defines support for the PHY to operate in Low Power Idle (LPI) mode. When enabled, this mode supports QUIET times during low link utilization, allowing the both sides of link to disable portions of each PHY's operating circuitry and save power.
- **Deep Green mode:** This mode also requires the CPU to recognize the long period power-down time and shut off the PHY power and the PLL to the PHY core. The CPU wakes up the PHY when a signal is detected at the PHY input.

2.24.1 Auto Power-Down Mode

Auto Power-Down mode saves PHY power consumption while the link is down. When the user enables the Auto Power-Down mode through a PHY register bit setting, the PHY goes into the power savings mode automatically whenever it is in link-down state. During the Power-Down state, the PHY wakes up every 2.7 or 5.4 seconds, depending on the register settings, and checks for a link signal. If no link signal is detected, then the PHY goes back to Power-Down state, or the PHY wakes up and resumes the link process.

Automatic Power-Down mode applies to the following conditions:

1. Cable is plugged in, but the link partner is shut down (for example, when a PC is off), so the port is in link down state.
2. Cable is unplugged, so the port is in link down state.

2.24.2 Energy Efficient Ethernet Mode

Energy Efficient Ethernet (EEE) power savings mode saves PHY power consumption while the link is up but when extended idle periods may exist between packet traffic. In EEE power savings mode PHY power consumption is scalable to the actual bandwidth utilization. The PHY can go in to *Quiet* mode (low-power idle mode) when there is no data to be transmitted. This feature is based on the latest IEEE 802.3az standard. The EEE supporting capability of the link partner is a must for this feature to work, and the discovery of the capability is during auto-negotiation through Link Layer Discovery Protocol (LLDP). This EEE feature is an embedded PHY feature and no external CPU is required.

In this mode, the MAC determines when to enter low power mode by examining the state of the transmit queues associated with each MAC. Four simple adjustments (settings) are used to trigger (optimize) the behavior of EEE control policy. These adjustments are:

- Global Buffer occupancy threshold
- Two-part sleep delay timer
- Minimum low-power idle duration timer
- Wake transition timer

The two-way communication between the PHY and its link partner is required for the PHY to achieve the power savings on both sides. The transmit PHY sends a sleep symbol to the link partner, and the link partner enters low power state. When the transmit PHY sends a wake symbol, the regular packet transfer mode resumes.

For details on how the mode works and how to set up the conditions, refer to the application note, *Energy Efficient Ethernet in the BCM53125*.

2.24.3 Deep Green Mode

The Deep Green Power Savings mode is a step deeper than the Auto Power-Down Power Savings mode. The Deep Green Power Savings mode can be enabled through the internal 8051 microcontroller by setting the EN_8051 strap pin high (default state). The Auto Power-Down Power Savings mode is per port, but the Deep Green Power Savings mode is for all the ports with a common PLL.

When all the ports that are sharing a PLL are linked down, the Auto Power-Down mode is enabled (register page 10h–14h, address 38h, shadow 01010b), the DLL Auto Power-Down mode is enabled (register page 10h–14h, address 38h, shadow 00101b), and then the PHYs enter the Deep Green Power Savings mode. In this mode, all the PHY circuits are powered down except the energy detection circuit, and the energy detection circuit constantly monitors the energy on the line. Upon signal energy detection, the microcontroller turns on the PLL along with PHY power to start the auto-negotiation process. The Deep Green Power Savings mode is most effective when the user expects no activities on the line for a long period of time.

2.25 Interrupt

The BCM53134P has an interrupt feature that can be programmed to generate two types of interrupts. The first type of interrupt is generated whenever link status change occurs in any of the ports. The second type of interrupt is generated when there is a packet(s) in IMP transmit queue to the CPU using the IMP port. The new type of interrupt is called IMP Sleep interrupt. The two types of interrupts can be enabled through Interrupt_Enable_Register (page 03h, address 08h), and both types of interrupts can be enabled together per port. However, the Link Status Change interrupts are normally used on port [4:0], and IMP Sleep interrupts are available only on Port 5 and the IMP port.

The purpose of the IMP Sleep interrupt is to wake up the external CPU connected to the IMP (or WAN) port before the switch to transmit the queued-up packet to the CPU. This interrupt can be used in any system with the CPU/switch combination. When there is no traffic to be sent to CPU, the CPU can go into sleep mode to conserve power. However, the CPU must be able to wake up when there is a packet(s) to be received, without losing any packets. For this situation, the BCM53134P can send out an interrupt to wake up the CPU and hold the packet(s) until the CPU is ready to accept the packets and process. The time to hold the packets to the CPU is programmable by the specific CPU based on its wake-up time requirement. The packet hold time (IMP sleep time) can be programmed through the register (page 03h, address 10h (for IMP port), address 14h (for Port 5)).

Once the interrupt is generated to wake up the CPU, the switch holds the packet(s) in the transmit queue till the IMP-sleep timer expires. The IMP-sleep timer is also programmable depending on the wake-up time required by the CPU.

The CPU has an option to enable each different type of interrupt through the Interrupt-Enable register. Once an interrupt is generated, the CPU can find out the source of the interrupt by reading the Interrupt Status register.

Chapter 3: System Functional Blocks

3.1 Overview

The BCM53134P includes the following blocks:

- “Media Access Controller”
- “Integrated 10/100/1000 PHY” on page 64
- “Frame Management” on page 72
- “MIB Engine” on page 75
- “Integrated High-Performance Memory” on page 82
- “Switch Controller” on page 82

Each of these blocks is discussed in more detail in the following sections.

3.2 Media Access Controller

The BCM53134P contains five 10/100/1000 MACs and one 10/100/1000/2.5G MAC.

The MAC automatically selects the appropriate speed (CSMA/CD or full-duplex) based on the PHY auto-negotiation result. In full-duplex mode, IEEE 802.3x PAUSE frame-based flow control is also determined through auto-negotiation. The MAC is IEEE 802.3, IEEE 802.3u, and IEEE 802.3x-compliant.

3.2.1 Receive Function

The MAC initiates frame reception following the assertion of receive data valid indication from the physical layer. The MAC monitors the frame for the following error conditions:

- Receive error indication from the PHY
- Runt frame error if frame is fewer than 64 bytes
- CRC error
- Long frame error if frame is greater than standard max frame size or 9,720 bytes for jumbo-enabled ports.

NOTE: Frames longer than standard max frame size are considered oversized frames. When jumbo-frame mode is enabled, only the frames longer than 9,720 bytes are bad frames and dropped.

If no errors are detected, the frame is processed by the switch controller. Frames with errors are discarded. Receive functions can be disabled using register settings.

3.2.2 Transmit Function

Frame transmission begins with the switch controller queuing a frame to the MAC transmitter. The frame data is transmitted as received from the switch controller. The transmit controller is responsible for preamble insertion, carrier deferral, collision backoff, and inter-packet gap enforcement.

In 10/100 Mb/s half-duplex mode, when a frame is queued for transmission, the transmit controller behaves as specified by the IEEE 802.3 requirements for frame deferral. Following deferral, the transmitter adds 8 bytes of preamble and SFD to the frame data received from the switch controller. If, during frame transmission, a collision is observed and the collision window timer has not expired, the transmit controller asserts jam and then executes the backoff algorithm. The frame is retransmitted when appropriate. On the 16th consecutive collision, the backoff algorithm starts over at the initial state, the collision counter is reset, and it attempts to transmit the current frame continue. Following a late collision, the frame is aborted, and the switch controller is allowed to queue the next frame for transmission.

While in full-duplex mode, the transmit controller ignores carrier activity and collision indication. Transmission begins after the switch controller queues the frame, and the 96-bit times of IPG have been observed. Transmit functions can be disabled using register settings.

3.2.3 Flow Control

The BCM53134P implement an intelligent flow-control algorithm to minimize the system impact resulting from traffic congestion. Buffer memory allocation is adaptive to the status of each port's speed and duplex mode, providing an optimal balance between flow management and per-port memory depth. The BCM53134P initiate flow control in response to buffer memory conditions on a per-port basis.

The MACs are capable of flow control in both full-and half-duplex modes.

3.2.3.1 10/100 Mb/s Half-Duplex

In 10/100 half-duplex mode, the MAC backpressures a receiving port by transmitting a 96-bit time jam packet to the port. A single jam packet is asserted for each received packet for the duration of the time the port is in the flow-control state.

3.2.3.2 10/100/1000 Mb/s Full-Duplex

Flow control in full-duplex mode functions as specified by the IEEE 802.3x requirements. In the receiver, MAC flow-control frames are recognized and, when properly received, set the flow-control pause time for the transmit controller. The pause time is assigned from the 2-byte pause time field following the pause opcode. MAC control PAUSE frames are not forwarded from the receiver to the switch controller.

When the switch controller requests flow control, the transmit controller transmits a MAC control PAUSE frame with the pause time set to maximum. When the condition that caused the flow control state is no longer present, a second MAC control PAUSE frame is sent with the pause time field set to 0.

3.3 Integrated 10/100/1000 PHY

There are four integrated PHY blocks in the BCM53134P. For more information, see [“Copper Interface” on page 84](#). The following sections describe the operations of the internal PHY block.

3.3.1 Encoder

There are four integrated PHY blocks in the BCM53134P. The PHY is the Ethernet transceiver that appropriately processes data presented by the MAC into an analog data stream to be transmitted at the MDI interface, which performs the reverse process on data received at the MDI interface. The registers of the PHY are read using the [“Programming Interfaces” on page 85](#). The following sections describe the operations of the internal PHY block. For more information, see [“Copper Interface” on page 84](#).

In 10BASE-T mode, Manchester encoding is performed on the data stream that is transmitted on the twisted-pair cable. The multimode transmit digital-to-analog converter (DAC) performs preequalization for 100m of Category 3 cabling.

In 100BASE-TX mode, the BCM53134P transmits a continuous data stream over the twisted-pair cable. The transmit packet is encapsulated by replacing the first two nibbles of preamble with a start-of-stream delimiter (/J/K codes) and appending an end-of-stream delimiter (/T/R codes) to the end of the packet. The transmitter repeatedly sends the idle code group between packets. The encoded data stream is serialized and then scrambled by the stream cipher block, as described in [“Stream Cipher” on page 67](#). The scrambled data is then encoded into MLT3 signal levels.

In 1000BASE-T mode, the BCM53134P simultaneously transmits and receives a continuous data stream on all 4 pairs of the Category 5 cable. Byte-wide data from the transmit data pins is scrambled when the transmit enable is asserted, and the trellis (a PAM5 symbol on each of the four twisted-pairs) is encoded into a four-dimensional code group and then inserted into the transmit data stream. The transmit packet is encapsulated by replacing the first 2 bytes of the preamble with a start-of-stream delimiter, and appending an end-of-stream delimiter to the end of the packet. When the transmit error input is asserted during a packet transmission, a transmit error code group is sent in place of the corresponding data code group. The transmitter sends idle code groups or carrier extend code groups between packets. Carrier extension is used by the MAC to separate packets within a multiple-packet burst and is indicated by asserting the transmit error signal and placing 0Fh on the transmit data pins while the transmit enable is low. A carrier extend error is indicated by replacing the transmit data input with 1Fh during carrier extension.

The encoding complies with IEEE standard IEEE 802.3ab and is fully compatible with previous versions of the Broadcom 1000BASE-T PHYs.

3.3.2 Decoder

In 10BASE-T mode, Manchester decoding is performed on the data stream.

In 100BASE-TX mode, following equalization and clock recovery, the receive data stream is converted from MLT3 to serial nonreturn-to-zero (NRZ) data. The NRZ data is descrambled by the stream cipher block, as described later in this document. The descrambled data is then deserialized and aligned into 5-bit code groups. The 5-bit code groups are decoded into 4-bit data nibbles. The start-of-stream delimiter is replaced with preamble nibbles, and the end-of-stream delimiter and idle codes are replaced with 0h. The decoded data is driven onto the MII receive data pins. When an invalid code group is detected in the data stream, the BCM53134P asserts the MII receive error (RX_ER) signal. RX_ER is also asserted when the link fails, or when the descrambler loses lock during packet reception.

In 1000BASE-T mode, the receive data stream is:

- Passed through the Viterbi decoder
- Descrambled
- Translated back into byte-wide data

The start-of-stream delimiter is replaced with preamble bytes, and the end-of-stream delimiter and idle codes are replaced with 00h. Carrier extend codes are replaced with 0Fh or 1Fh. Decoding complies with IEEE standard IEEE 802.3ab and is fully compatible with previous versions of Broadcom 1000BASE-T PHYs.

3.3.3 Link Monitor

In 10BASE-T mode, a link-pulse detection circuit constantly monitors the TRD pins for the presence of valid link pulses.

In 100BASE-TX mode, receive signal energy is detected by monitoring the receive pair for transitions in the signal level. Signal levels are qualified using squelch detect circuits. When no signal is detected on the receive pair, the link monitor enters the Link Fail state and the transmission and reception of data packets is disabled. When a valid signal is detected on the receive pair for a minimum of 1 ms, the link monitor enters the Link Pass state, and the transmit and receive functions are enabled.

Following auto-negotiation in 1000BASE-T mode, the master transceiver begins sending data on the media. The slave transceiver also begins transmitting when it has recovered the master transceiver's timing. Each end of the link continuously monitors its local receiver status. When the local receiver status has been good for at least 1 microsecond, the link monitor enters the Link Pass state, and the transmission and reception of data packets are enabled. When the local receiver status is bad for more than 750 ms, the link monitor enters the Link Fail state and the transmission and reception of data packets are disabled.

3.3.4 Digital Adaptive Equalizer

The digital adaptive equalizer removes intersymbol interference (ISI) created by the transmission channel media. The equalizer accepts sampled unequalized data from the analog-to-digital converter (ADC) on each channel and produces equalized data. The BCM53134P achieves an optimum signal-to-noise ratio by using a combination of feed forward equalization (FFE) and decision feedback equalization (DFE) techniques. Under harsh noise environments, these powerful techniques achieve a bit error rate (BER) of less than 1×10^{-12} for transmissions up to 100m on Category 5 twisted-pair cabling (100m on Category 3 UTP cable for 10BASE-T mode). The all-digital nature of the design makes the performance very tolerant to noise. The filter coefficients are self-adapting to accommodate varying conditions of cable quality and cable length.

3.3.5 Echo Canceler

Because of the bidirectional nature of the channel in 1000BASE-T mode, an echo impairment is caused by each transmitter. The output of the echo filter is added to the FFE output to remove the transmitted signal impairment from the incoming receive signal. The echo canceler coefficients are self-adapting to manage the varying echo impulse responses caused by different channels, transmitters, and environmental conditions.

3.3.6 Crosstalk Canceler

The BCM53134P transmits and receives a continuous data stream on four channels. For a given channel, the signals sent by the other three local transmitters cause impairments on the received signal because of near-end crosstalk (NEXT) between the pairs. It is possible to cancel the effect because each receiver has access to the data for the other three pairs that cause this interference. The output of the adaptive NEXT canceling filters is added to the FFE output to cancel the NEXT impairment.

3.3.7 Analog-to-Digital Converter

Each receive channel has its own 125 MHz analog-to-digital converter (ADC) that samples the incoming data on the receive channel and feeds the output to the digital adaptive equalizer. Advanced analog circuit techniques achieve the following results:

- Low offset

- High power-supply noise rejection
- Fast settling time
- Low bit error rate

3.3.8 Clock Recovery/Generator

The clock recovery and generator block creates the transmit and receive clocks for 1000BASE-T, 100BASE-TX, and 10BASE-T operation.

In 10BASE-T or 100BASE-TX mode, the transmit clock is locked to the 25 MHz crystal input, and the receive clock is locked to the incoming data stream.

In 1000BASE-T mode, the two ends of the link perform loop timing. One end of the link is configured as the master, and the other is configured as the slave. The master transmit and receive clocks are locked to the 25 MHz crystal input. The slave transmit and receive clocks are locked to the incoming receive data stream. Loop timing allows for the cancellation of echo and NEXT impairments by ensuring that the transmitter and receiver at each end of the link are operating at the same frequency.

3.3.9 Baseline Wander Correction

1000BASE-T and 100BASE-TX data streams are not always DC-balanced. Because the receive signal must pass through a transformer, the DC offset of the differential receive input can vary with data content. This effect, which is known as baseline wander, can greatly reduce the noise immunity of the receiver. The BCM53134P automatically compensates for baseline wander by removing the DC offset from the input signal, thereby significantly reducing the probability of a receive symbol error.

In 10BASE-T mode, baseline wander correction is not performed because the Manchester coding provides a perfect DC balance.

3.3.10 Multimode TX Digital-to-Analog Converter

The multimode transmit digital-to-analog converter (DAC) transmits PAM5, MLT3, and Manchester coded symbols. The transmit DAC performs signal wave shaping that decreases the unwanted high frequency signal components, reducing electromagnetic interference (EMI). The transmit DAC uses a current voltage driven output with internal terminations and hence, does not require external components or a magnetic supply for operation which reduces system complexity for routing and bill of materials.

3.3.11 Stream Cipher

In 1000BASE-T and 100BASE-TX modes, the transmit data stream is scrambled to reduce radiated emissions and to ensure that there are adequate transitions within the data stream. The 1000BASE-T scrambler also ensures that there is no correlation among symbols on the four different wire pairs and in the transmit and receive data streams. The scrambler reduces peak emissions by randomly spreading the signal energy over the transmit frequency range and eliminating peaks at certain frequencies. The randomization of the data stream also assists the digital adaptive equalizers and echo/crosstalk cancelers. The algorithms in these circuits require there to be no sequential or cross-channel correlation among symbols in the various data streams.

In 100BASE-TX mode, the transmit data stream is scrambled by exclusive ORing the encoded serial data stream. This is done with the output of an 11-bit wide linear feedback shift register (LFSR), producing a 2047-bit nonrepeating sequence.

In 1000BASE-T mode, the transmit data stream is scrambled by exclusive ORing the input data byte with an 8-bit wide cipher text word. The cipher text word generates each symbol period from eight uncorrelated maximal length data sequences that are produced by linear remapping of the output of a 33-bit wide LFSR. After the scrambled data bytes are encoded, the sign of each transmitted symbol is again randomized by a 4-bit wide cipher text word that is generated in the same manner as the 8-bit word. The master and slave transmitters use different scrambler sequences to generate the cipher text words. For repeater or switch applications, where all ports can transmit the same data simultaneously, signal energy is randomized further by using a unique seed to initialize the scrambler sequence for each PHY.

The receiver descrambles the incoming data stream by exclusive ORing it with the same sequence generated at the transmitter. The descrambler detects the state of the transmit LFSR by looking for a sequence representing consecutive idle code groups. The descrambler locks to the scrambler state after detecting a sufficient number of consecutive idle codes. The BCM53134P enables transmission and reception of packet data only when the descrambler is locked. The receiver continually monitors the input data stream to ensure that it has not lost synchronization by checking that inter-packet gaps containing idles or frame extensions are received at expected intervals. When the BCM53134P detects loss of synchronization, it notifies the remote PHY of the inability to receive packets (1000BASE-T mode only) and attempts to resynchronize to the received data stream. If the descrambler is unable to resynchronize for a period of 750 ms, the BCM53134P is forced into the Link Fail state.

In 10BASE-T mode, scrambling is not required to reduce radiated emissions.

3.3.12 Wire Map and Pair Skew Correction

During 1000BASE-T operation, the BCM53134P has the ability to automatically detect and correct some UTP cable wiring errors. The symbol decoder detects and compensates for (internal to the BCM53134P) the following errors:

- Wiring errors caused by the swapping of pairs within the UTP cable.
- Polarity errors caused by the swapping of wires within a pair.

The BCM53134P also automatically compensates for differences in the arrival times of symbols on the four pairs of the UTP cable. The varying arrival times are caused by differing propagation delays (commonly referred to as delay skew) between the wire pairs. The BCM53134P can tolerate delay skews of up to 64 ns long. Auto-negotiation must be enabled to take advantage of the wire map correction.

During 10/100 Mb/s operation, pair swaps are corrected. Delay skew is not an issue though, because only one pair of wires is used in each direction.

3.3.13 Automatic MDI Crossover

During copper auto-negotiation, one end of the link must perform an MDI crossover so that each transceiver's transmitter is connected to the other receiver. The BCM53134P can perform an automatic media-dependent interface (MDI) crossover, eliminating the need for crossover cables or cross-wired (MDIX) ports. During auto-negotiation, the BCM53134P normally transmits and receives on the TRD pins.

When connecting to another device that does not perform MDI crossover, the BCM53134P automatically switches its TRD in pairs when necessary to communicate with the remote device. When connecting to another device that does have MDI crossover capability, an algorithm determines which end performs the crossover function.

During 1000BASE-T operation, the BCM53134P swaps the transmit symbols on pairs 0 and 1 and pairs 2 and 3 if auto-negotiation completes in the MDI crossover state. The 1000BASE-T receiver automatically detects pair swaps on the receive inputs and aligns the symbols properly within the decoder. The automatic MDI crossover function cannot be disabled when in 1000BASE-T mode. During 10BASE-TX and 100BASE-T operation, pair swaps automatically occur within the device and do not require user intervention. The automatic MDI crossover function by default works only when auto-negotiation is enabled. This function can be disabled during auto-negotiation using a register write.

NOTE: This function operates only when the copper auto-negotiation is enabled.

3.3.14 10/100BASE-TX Forced Mode Auto-MDIX

The automatic MDI crossover function can also be enabled when in forced 10BASE-T or forced 100BASE-TX mode. This feature allows the user to disable the copper auto-negotiation in either 10BASE-T or 100BASE-TX and still take advantage of the automatic MDI crossover function. Whenever the forced link is down for at least 4 seconds, then auto-negotiation is internally enabled with its automatic MDI crossover function until link pulses or 100TX idles are detected. Once detected, the PHY returns to forced mode operation.

The user should set the same speed in register 0 and the auto-negotiation advertisement register 4.

NOTE: This function operates only when the copper auto-negotiation is disabled.

3.3.15 Resetting the PHY

The BCM53134P provides a hardware reset pin, $\overline{\text{RESET}}$, which resets all internal nodes to a known state. Hardware reset is accomplished by holding the $\overline{\text{RESET}}$ pin low for at least 1 ms. Once $\overline{\text{RESET}}$ is brought high, the PHY will complete its reset sequence within 5 ms. All outputs will be inactive until the PHY has completed its reset sequence. The PHY will keep the inputs inactive for 5 ms after the deassertion of hardware reset. The hardware configuration pins and the PHY address pins will be read on the deassertion of hardware reset.

The BCM53134P also has a software reset capability. To enable the software reset, a 1 must be written to the bit. This bit is self-clearing, meaning that a second write operation is not necessary to end the reset. There is no effect if 0 is written to this bit. Mode pins that are labeled sample on reset (SOR) are latched during hardware reset. However, software resets do not latch new values for the SOR mode pins.

3.3.16 PHY Address

The BCM53134P has four unique PHY addresses for MII management of the internal PHYs. The PHY addresses for each port are as follows,

- PHY address for Port 0 is 0
- PHY address for Port 1 is 1
- PHY address for Port 2 is 2
- PHY address for Port 3 is 3
- SGMII PHY address is 4

3.3.17 Super Isolate Mode

When in Super Isolate mode, the transmit and receive functions on the Copper Media Dependent Interface are disabled (no link is established with the PHY's copper link partner). Any data received from the switch is ignored by the BCM53134P and no data is sent from the BCM53134P.

3.3.18 Standby Power-Down Mode

The BCM53134P can be placed into standby power-down mode using software commands. In this mode, all PHY functions except for the serial management interface are disabled. There are three ways to exit standby power-down mode:

- Clear MII Control register (address 00h), bit 11 = 0.
- Set the software RESET bit 15.
- Assert the hardware $\overline{\text{RESET}}$ pin.

Read or write operations to any MII register, other than MII Control register, while the device is in the standby power-down mode returns unpredictable results. Upon exiting standby power-down mode, the BCM53134P remains in an internal reset state for 40 μs and then resumes normal operation.

3.3.19 Auto Power-Down Mode

The BCM53134P can be placed into auto power-down mode. Auto power-down mode reduces device power when the signal from the copper link partner is not present. The auto power-down mode works whether the device is in Auto-negotiation Enabled or Forced mode. This mode is enabled by setting bit 5 = 1 of Auto Power-Down register. When auto power-down mode is enabled, the BCM53134P automatically enters the low-power mode when energy on the line is lost, and it resumes normal operation when energy is detected. The energy-detect circuit is always enabled even when a port is in low-power mode. When the BCM53134P is in auto power-down mode, it wakes up after 2.7s or 5.4s, which determined by bit 4 of Auto Power-Down register, and sends link pulses to the link partner. The BCM53134P enters normal operation and establishes a link if energy is detected.

NOTE: Auto power-down mode is a Broadcom proprietary feature and is based on IEEE standard.

3.3.20 External Loopback Mode

The External Loopback mode allows in-circuit testing of the BCM53134P as well as the transmit path through the magnetics and the RJ-45 connector. External loopback can be performed with and without a jumper block. External loopback with a jumper block tests the path through the magnetics and RJ-45 connector. External loopback without the jumper block tests only the BCM53134P's transmit and receive circuitry. In 1000BASE-T, 100BASE-TX, and 10BASE-T modes, a jumper block must be inserted into the RJ-45 connector to support external loopback. The jumper block should have the following RJ-45 pins connected together:

```
1-----3
2-----6
4-----7
5-----8
```

The following six tables describe how the external loopback is enabled for 1000BASE-T, 100BASE-TX, and 10BASE-T modes with and without a jumper block.

Table 15: 1000BASE-T External Loopback with External Loopback Plug

Register Writes	Comments
Write 1800h to 1000BASE-T Control register	Enable 1000BASE-T Master Mode
Write 0040h to MII Control register	Enable Force 1000BASE-T
Write 8400h to Auxiliary Control register	Enable External Loopback Mode with external loopback plug

Table 16: 1000BASE-T External Loopback Without External Loopback Plug

Register Writes	Comments
Write 1800h to 1000BASE-T Control register	Enable 1000BASE-T Master Mode
Write 0040h to MII Control register	Enable Force 1000BASE-T
Write 8400h to Auxiliary Control register	Enable External Loopback Mode
Write 0014h to Auxiliary Control register	Enable External Loopback Mode without external loopback plug

Table 17: 100BASE-TX External Loopback with External Loopback Plug

Register Writes	Comments
Write 2100h to MII Control register	Enable Force 100BASE-TX full-duplex mode

Table 18: 100BASE-TX External Loopback Without External Loopback Plug

Register Writes	Comment
Write 2100h to MII Control register	Enable Force 100BASE-TX full-duplex mode
Write 0014h to Auxiliary Control register	Enable external loopback mode without external loopback plug

Table 19: 10BASE-T External Loopback with External Loopback Plug

Register Writes	Comments
Write 0100h to MII Control register	Enable Force 10BASE-T full-duplex mode

Table 20: 10BASE-T External Loopback Without External Loopback Plug

Register Writes	Comments
Write 0100h to MII Control register	Enable Force 10BASE-T full-duplex mode
Write 0014h to Auxiliary Control register	Enable external loopback mode without external loopback plug

NOTE: To exit the External Loopback mode, a software or hardware reset is recommended.

3.3.21 Full-Duplex Mode

The BCM53134P supports full-duplex operation. While in full-duplex mode, a transceiver can simultaneously transmit and receive packets on the cable.

3.3.21.1 Copper Mode

When auto-negotiation is disabled, full-duplex operation can be enabled using register settings.

When auto-negotiation is enabled, the full-duplex capability is advertised for one of the following, depending on the register settings:

- 10BASE-T
- 100BASE-T
- 1000BASE-T

3.3.22 Master/Slave Configuration

In 1000BASE-T mode, the BCM53134P and its link partner perform loop timing. One end of the link must be configured as the timing master, and the other end as the slave. Master/slave configuration is performed by the auto-negotiation function. The auto-negotiation function first looks at the manual master/slave configuration bits advertised by the local PHY and the link partner. If neither PHY requests manual configuration, then the auto-negotiation function looks at the advertised repeater/DTE settings. If one PHY is advertised as a repeater port and the other is advertised as a DTE port, then the repeater port is configured as the master and the DTE port as the slave. Each end generates an 11-bit random seed if the two settings are equal, and the end with the higher seed is configured as the master. If the local PHY and the link partner generate the same random seed, then auto-negotiation is restarted.

If both ends of the link attempt to force the same manual configuration (both master or both slave), or the random seeds match seven consecutive times, then the BCM53134P sets the Master/Slave Configuration Fault bit in the 1000BASE-T Status register, and auto-negotiation is restarted. This is used to set the BCM53134P to manual master/slave configuration or to set the advertised repeater/DTE configuration.

3.3.23 Next Page Exchange

The 1000BASE-T configuration requires the exchange of three auto-negotiation next pages between the BCM53134P and its link partner. Exchange of 1000BASE-T Next Page information takes place automatically when the BCM53134P is configured to advertise 1000BASE-T capability.

The BCM53134P also supports software controlled Next Page exchanges. This includes the three 1000BASE-T Next Pages, which are always sent first. The BCM53134P automatically generates the appropriate message code field for the 1000BASE-T pages. When the BCM53134P is not configured to advertise 1000BASE-T capability, the 1000BASE-T Next Pages are not sent.

When the BCM53134P is not configured to advertise 1000BASE-T capability, the BCM53134P does not advertise Next Page ability.

3.4 Frame Management

The BCM53134P provides a Frame Management block that works in conjunction with one of the RGMII ports operating in IMP mode as the full-duplex packet-streaming interface to the external CPU, with in-band messaging mechanism for management purpose.

3.4.1 In-Band Management Port

The BCM53134P provides two PHY-less interface ports and supports a dual IMP ports (IMP port and Port 5) feature. One (IMP port) or both PHY-less ports (IMP port and Port 5) can be configured as the management port using register settings. In the dual IMP feature, all traffic to the CPU from LAN ports will be forwarded to IMP port, and all traffic to the CPU from WAN ports will be forwarded to Port 5. When the PHY-less interface port is defined as the Frame Management Port, it is referred to as the in-band management port (IMP).


The IMP can be used as a full-duplex 10/100/1000 Mb/s port, which can be used to forward management information to the external management agent, such as BPDUs, mirrored frames, or frames addressed to other static address entries that have been identified as a special interest to the management system. Two interfaces support only full-duplex mode.

NOTE: The signals of the IMP port are named as IMP_xxx in [“Signal Descriptions” on page 117](#). The signals of the WAN (port5) port are name as GMII_xxx.

As IMP is defined as the frame management port, normal frame data is forwarded to the port based on register settings. If these bits are cleared, no frame data will be forwarded to the Frame Management Port, with the exception that frames meeting the mirror ingress/egress rules criteria, will always be forwarded to the designated frame management port.

Packets transferred over the IMP port are tagged with the Broadcom proprietary header to carry the necessary information which is of interest to the management entity running on the CPU, as shown in [Figure 24](#), except for the PAUSE frame. The IMP port must support normal Ethernet pause based flow control mechanism.

Figure 24: IMP Packet Encapsulation Format



- The alignment of the regenerated FCS may vary depending on the original frame length.
- The regenerated FCS covers from the whole frame, including the BRCM tag.

The BRCM tag is designed for asymmetric operation across the IMP port. The information carried from the switching device to the CPU is different from the information carried from the CPU to the switching device.

Similarly, the host system must insert the BRCM tag fields into frames it wished to send into the management port, to be routed to specific egress ports. The OPCODE within the tag field determines how the frame is handled, and allows frames to be forwarded using the normal address lookup using a port ID designation within the Tag.

The BRCM tags are transmitted with the convention of highest significant octet first, followed by the next lowest significant octet, and so on, with the least significant bit of each octet transmitted out from the MAC first. Therefore, for the BRCM tag field in [Table 21](#), the most significant octet would be transmitted first (bits [24:31]), with bit 24 being the first bit transmitted.

3.4.2 Broadcom Tag Format for Egress Packet Transfer

When a packet is forwarded by the switching device to the external CPU for processing, the BRCM tag is formatted as shown in [Table 21](#).

Table 21: Egress Broadcom Tag Format (IMP to CPU)

31–29	28–24	23–16	15–8	7–5	4–0
OPCODE=000	Reserved	CLASSIFICATION_ID[7:0]	REASON_CODE[7:0]	TC[2:0]	SRC_PID[4:0]
63–61	60–38			37	36–32
OPCODE=001	Reserved			T/R	T/R_PID[4:0]
31–0					
TIME_STAMP[31:0]					

- **OPCODE 000**
This indicates the packet transfer with explicit reasons to help the external CPU to direct the packet for the appropriate packet processing entities.
- **CLASSIFICATION_ID [7:0]**
It indicates the packet flow ID classified based on CFP rules. 0x00 indicates no classification ID exists.
- **REASON_CODE [7:0]**
This indicates the reasons why the packet is forwarded to the external CPU so that the CPU can identify the appropriate software routines for packet processing.
 - Bit [0] indicates mirroring
 - Bit [1] indicates SA learning
 - Bit [2] indicates switching
 - Bit [3] indicates protocol termination
 - Bit [4] indicates protocol snooping
 - Bit [5] indicates flooding/exception processing
 - Bit [6] and bit [7] are reserved
- **TC [2:0]**
This indicates the traffic class classified by the switching device when forwarding the packet to the CPU.
- **SRC_PID [4:0]**
This indicates the ingress port of the switching device where the packet is received.
- **OPCODE 001**
This indicates a packet transfer with explicit timestamp recorded at the port where it was transmitted or received (indicated by the T/R_PID) for IEEE 802.1AS protocol implementation.
- **T/R**
This indicates the type of timestamp. 0 indicates the timestamp recorded when the packet was received through the port (indicated by the T/R_PID); 1 indicates the timestamp recorded when the packet was transmitted through the port (indicated by the T/R_PID).
- **T/R_PID [4:0]**
This indicates the port through which the packet was transmitted when T/R = 1, or the port through which the packet was received when T/R = 0.
- **TIME_STAMP [31:0]**

This carries the timestamp recorded at the port through which the packet was transmitted when T/R = 1, or the timestamp recorded at the port through which the packet was received when T/R = 0.

3.4.3 Broadcom Tag Format for Ingress Packet Transfer

For packet transfer from the external CPU to the switching device, the Broadcom tag is formatted as shown in [Table 22](#).

Table 22: Ingress BRCM Tag (CPU to IMP)

31–29	28–26	25–24	23–0	
OPCODE=000	TC[2:0]	TE[1:0]	Reserved	
31–29	28–26	25–24	23	22–0
OPCODE=001	TC[2:0]	TE[1:0]	TS	DST_MAP[22:0]

- **OPCODE 000**
It indicates that the external CPU is not dictating how the packet is forwarded, and the packet is forwarded by the switching device based on the original Ethernet packet information.
- **OPCODE 001**
This indicates the packet is forwarded to multiple (or single) egress ports by the switching device based on the explicit direction of the external CPU.
- **DST_MAP [22:0]**
This indicates the egress port bit map to which the external CPU intends to forward the packet. Bits [5:0] = Port [5:0], Bit 8 = IMP port.
- **TC [2:0]**
This indicates the traffic class with which the external CPU intends to forward the packet.
- **TS (timestamp request)**
This indicates whether the transmit timestamped at the egress port should be reported back to the external CPU.
- **TE (tag enforcement)**
This indicates the IEEE 802.1Q/P tagging/untagging encapsulation enforcement for the packet transmission.
00 = No enforcement (follow VLAN untag mask rules)
01 = Untag enforcement
10 = Tag enforcement
11 = Reserved

3.5 MIB Engine

The MIB Engine is responsible for processing status words received from each port. Based on whether it is a receive status or transmit status, appropriate MIB counters are updated. The BCM53134P implement 70-plus MIB counters on a per-port basis. MIB counters can be categorized into three groups: receive-only counters, transmit-only counters, and receive or transmit counters. This latter group can, as a group, be selectively steered to the receive or transmit process on a per-port basis. The section below describes each individual counter.

The BCM53134P offers the MIB snapshot feature per port. A snapshot of a selected port MIB registers can be captured and available to the users while MIB counters are continuing to count.

3.5.1 MIB Counters Per Port

Receive Only Counters (19) Description of Counter

RxDropPkts (32 bit) The number of good packets received by a port that were dropped due to a lack of resources (such as lack of input buffers) or were dropped due to a lack of resources before a determination of the validity of the packet was able to be made (such as receive FIFO overflow). The counter is incremented only if the receive error was not counted by the **RxExcessSizeDisc**, the **RxAlignmentErrors**, or the **RxFCSErrors** counters.

RxOctets (64 bit)—The number of data bytes received by a port (excluding preamble, but including FCS), including bad packets.

RxBroadcastPkts (32 bit)—The number of good packets received by a port that are directed to the broadcast address. This counter does not include errored broadcast packets or valid multicast packets. The maximum packet size can be programmed.

RxMulticastPkts (32 bit)—The number of good packets received by a port that are directed to a multicast address. This counter does not include errored multicast packets or valid broadcast packets. The maximum packet size can be programmed.

RxSACHanges (32 bit)—The number of times the SA of good receive packets has changed from the previous value. A count greater than 1 generally indicates the port is connected to a repeater-based network. The maximum packet size can be programmed.

RxUndersizePkts (32 bit)—The number of good packets received by a port that are less than 64 bytes long (excluding framing bits, but including the FCS).

RxOversizePkts (32 bit)—The number of good packets received by a port that are greater than standard max frame size. The maximum packet size can be programmed.

RxFragments (32 bit)—The number of packets received by a port that are less than 64 bytes (excluding framing bits) and have either an FCS error or an alignment error.

RxJabbers (32 bit)—The number of packets received by a port that are longer than standard max frame size and have either an FCS error or an alignment error.

RxUnicastPkts (32 bit)—The number of good packets received by a port that are addressed to a unicast address. The maximum packet size can be programmed.

RxAlignmentErrors (32 bit)—The number of packets received by a port that have a length (excluding framing bits, but including FCS) between 64 and standard max frame size, inclusive, and have a bad FCS with a nonintegral number of bytes.

RxFCSErrors (32 bit)—The number of packets received by a port that have a length (excluding framing bits, but including FCS) between 64 and standard max frame size inclusive, and have a bad FCS with an integral number of bytes.

RxGoodOctets (64 bit)—The total number of bytes in all good packets received by a port (excluding framing bits, but including FCS). The maximum packet size can be programmed.

JumboPktCount (32 bit)—The number of good packets received by a port that are greater than the standard maximum size and less than or equal to the jumbo packet size, regardless of CRC or alignment errors.

RxPausePkts (32 bit)—The number of PAUSE frames received by a port. The PAUSE frame must have a valid MAC Control Frame EtherType field (88–08h), have a destination MAC address of either the MAC Control frame reserved multicast address (01-80-C2-00-00-01) or the unique MAC address associated with the specific port, a valid PAUSE opcode (00–01), be a minimum of 64 bytes in length (excluding preamble but including FCS), and have a valid CRC. Although an IEEE 802.3-compliant MAC is permitted to transmit PAUSE frames only when in full-duplex mode with flow control enabled and with the transfer of PAUSE frames determined by the result of auto-negotiation, an IEEE 802.3 MAC receiver is required to count all received PAUSE frames, regardless of its half/full-duplex status. An indication that a MAC is in half-duplex with the RxPausePkts incrementing indicates a noncompliant transmitting device on the network.

RxSymbolErrors (32 bit)—The total number of times a valid-length packet was received at a port and at least one invalid data symbol was detected. The counter increments only once per carrier event and does not increment on detection of a collision during the carrier event.

RxDiscard (32 bit)—The number of good packets received by a port that were discarded by the Forwarding Process.

InRangeErrors (32 bit)—The number of packets received with good CRC and one of the following: (1) The value of length/type field is between 46 and 1500 inclusive, and does not match the number of (MAC client data + PAD) data octets received, OR (2) The value of length/type field is less than 46, and the number of data octets received is greater than 46 (which does not require padding).

OutOfRangeErrors (32 bit)—The number of packets received with good CRC and the value of length/type field is greater than 1500 and less than 1536.

Transmit Counters Only (19)—Description of Counter

TxDropPkts (32 bit)—This counter is incremented every time a transmit packet is dropped due to lack of resources (such as transmit FIFO underflow), or an internal MAC sublayer transmit error not counted by either the TxLateCollision or the TxExcessiveCollision counters.

TxOctets (64 bit)—The total number of good bytes of data transmitted by a port (excluding preamble but including FCS).

TxBroadcastPkts (32 bit)—The number of good packets transmitted by a port that are directed to a broadcast address. This counter does not include errored broadcast packets or valid multicast packets.

TxMulticastPkts (32 bit)—The number of good packets transmitted by a port that are directed to a multicast address. This counter does not include errored multicast packets or valid broadcast packets.

TxCollisions (32 bit)—The number of collisions experienced by a port during packet transmissions.

TxUnicastPkts (32 bit)—The number of good packets transmitted by a port that are addressed to a unicast address.

TxSingleCollision (32 bit)—The number of packets successfully transmitted by a port that have experienced exactly one collision.

TxMultipleCollision (32 bit)—The number of packets successfully transmitted by a port that have experienced more than one collision.

TxDeferredTransmit (32 bit)—The number of packets transmitted by a port for which the first transmission attempt is delayed because the medium is busy. This applies only to the half-duplex mode, while the Carrier Sensor Busy.

TxLateCollision (32 bit)—The number of times that a collision is detected later than 512 bit-times into the transmission of a packet.

TxExcessiveCollision (32 bit)—The number of packets that are not transmitted from a port because the packet experienced 16 transmission attempts.

TxPausePkts (32 bit)—The number of PAUSE events at each port.

TxFramelnDisc (32 bit)—The number of valid packets received which are discarded by the forwarding process due to lack of space on an output queue (not maintained or reported in the MIB counters). Located in the Congestion Management registers (page 0Ah). This attribute increments only if a network device is not acting in compliance with a flow control request, or the BCM53134P internal flow-control/buffering scheme has been configured incorrectly.

TxQ0PKT (32 bit)—The total number of good packets transmitted on CoS0, which is specified in MIB queue select register when QoS is enabled.

TxQ1PKT (32 bit)—The total number of good packets transmitted on CoS1, which is specified in MIB queue select register when QoS is enabled.

TxQ2PKT (32 bit)—The total number of good packets transmitted on CoS2, which is specified in MIB queue select register when QoS is enabled.

TxQ3PKT (32 bit)—The total number of good packets transmitted on CoS3, which is specified in MIB queue select register when QoS is enabled.

TxQ4PKT (32 bit)—The total number of good packets transmitted on CoS4, which is specified in MIB queue select register when QoS is enabled.

TxQ5PKT (32 bit)—The total number of good packets transmitted on CoS5, which is specified in MIB queue select register when QoS is enabled.

Transmit or Receive Counters (10) Description of Counter

Pkts64Octets (32 bit)—The number of packets (including error packets) that are 64 bytes long.

Pkts65to127Octets (32 bit)—The number of packets (including error packets) that are between 65 and 127 bytes long.

Pkts128to255Octets (32 bit)—The number of packets (including error packets) that are between 128 and 255 bytes long.

Pkts256to511Octets (32 bit)—The number of packets (including error packets) that are between 256 and 511 bytes long.

Pkts512to1023Octets (32 bit)—The number of packets (including error packets) that are between 512 and 1023 bytes long.

Pkts1024toMaxPktOctets (32 bit)—The number of packets that (include error packets) are between 1024 and the standard maximum packet size, inclusive.

EEE Counters/Description of Counters

LPI Idle count (32 bit)—EEE low-power idle event.

In asymmetric mode, this is simply a count of the number of times that the lowPowerAssert control signal has been asserted for each MAC. In symmetric mode, this is the count of the number of times both lowPowerAssert and the lowPowerIndicate (from the receive path) are asserted simultaneously.

LPI Duration Count (32 bit)—EEE low-power idle duration.

This counter accumulates the number of microseconds that the associated MAC/PHY is in the low-power idle state. The unit is 1 μ sec.

Total number of counters per port: 45

Table 23 identifies the mapping of the BCM53134P MIB counters and their generic mnemonics to the specific counters and mnemonics for each of the key IETF MIBs that are supported. Direct mappings are defined. However, there are several additional statistics counters, which are indirectly supported that make up the full complement of the counters required to fully support each MIB. These are shown in Table 24 on page 81.

Finally, Table 25 on page 81 identifies the additional counters supported by the BCM53134P and references the specific standard or reason for the inclusion of the counter.

Table 23: Directly Supported MIB Counters

BCM53134P MIB	Ethernet-Like MIB RFC 1643	Bridge MIB RFC 1493	MIB II Interface RFC 1213/1573	RMON MIB RFC 1757
RxDropPkts	dot3StatsInternalMACReceivesErrors	dot1dTpPortInDiscards	ifInDiscards	–
RxOctets	–	–	ifInOctets	etherStatsOctets
RxBroadcastPkts	–	–	ifInBroadcastPkts	etherStatsBroadcastPkts
RxMulticastPkts	–	–	ifInMulticastPkts	etherStatsMulticastPkts
RxSACChanges	Note 2	Note 2	Note 2	Note 2
RxUndersizePkts	–	–	–	etherStatsUndersizePkts
RxOversizePkts	dot3StatsFrameTooLongs	–	–	etherStatsOversizePkts
RxFragments	–	–	–	etherStatsFragments
RxJabbers	–	–	–	etherStatsJabbers
RxUnicastPkts	–	–	ifInUcastPkts	–
RxAlignmentErrors	dot3StatsAlignmentErrors	–	–	–
RxFCSErrors	dot3StatsFCSErrors	–	–	–
RxGoodOctets	–	–	–	–
RxExcessSizeDisc	Note 2	Note 2	Note 2	Note 2
RxPausePkts	Note 2	Note 2	Note 2	Note 2
RxSymbolErrors	Note 2	Note 2	Note 2	Note 2
Note 1	–	–	ifInErrors	–
Note 1	–	–	ifInUnknownProtos	–
Note 1	–	dot1dTpPortInFrames	–	–
TxDropPkts	dot3StatsInternalMACTransmitErrors	–	ifOutDiscards	–
TxOctets	–	–	ifOutOctets Note 3	–
Note 1	–	dot1dTpPortOutFrames	–	–
TxBroadcastPkts	–	–	ifOutBroadcastPkts	–
TxMulticastPkts	–	–	ifOutMulticastPkts	–
TxCollisions	–	–	–	etherStatsCollisions
TxUnicastPkts	–	–	ifOutUcastPkts	–
TxSingleCollision	dot3StatsSingleCollisionFrames	–	–	–

Table 23: Directly Supported MIB Counters (Continued)

BCM53134P MIB	Ethernet-Like MIB RFC 1643	Bridge MIB RFC 1493	MIB II Interface RFC 1213/1573	RMON MIB RFC 1757
TxMultipleCollision	dot3StatsMultipleCollisionFrames	–	–	–
TxDeferredTransmit	dot3StatsDeferredTransmissions	–	–	–
TxLateCollision	dot3StatsLateCollision	–	–	–
TxExcessiveCollision	dot3StatsExcessiveCollision	–	–	–
TxFramelnDisc	Note 2	Note 2	Note 2	Note 2
TxPausePkts	Note 2	Note 2	Note 2	Note 2
Note 4	dot3StatsCarrierSenseErrors	–	–	–
Note 1	–	–	ifOutErrors	–
Pkts64Octets	–	–	–	etherStatsPkt64Octets
Pkts65to127Octets	–	–	–	etherStatsPkt65to127Octets
Pkts128to255Octets	–	–	–	etherStatsPkt128to255Octets
Pkts256to511Octets	–	–	–	etherStatsPkt256to511Octets
Pkts512to1023Octets	–	–	–	etherStatsPkt512to1023Octets
Pkts1024toMaxPktOctets	–	–	–	etherStatsPkt1024toMaxPktOctets
Note 1	–	–	–	etherStatsDropEvents
Note 1	–	–	–	etherStatsPkts
Note 1	–	–	–	etherStatsCRCAAlignErrors
Note 4	dot3StatsSQETestErrors	–	–	–

Note 1: Derived by summing two or more of the supported counters. See [Table 24](#) for specific details.

Note 2: Extensions required by recent standards developments or BCM53134P operation specifics.

Note 3: The MIB II interfaces specification for if OutOctets includes preamble/SFD and errored bytes. Because IEEE 802.3-compliant MACs have no requirement to keep track of the number of transmit bytes in an errored frame, this count is impossible to maintain. The TxOctets counter maintained by the BCM53134P is consistent with good bytes transmitted, excluding preamble, but including FCS. The count can be adjusted to more closely match the OutOctets definition by adding the preamble for TxGoodPkts and possibly an estimate of the octets involved in TxCollisions and TxLateCollision.

Note 4: The attributes TxCarrierSenseErrors and TxSQETestErrors are not supported in the BCM53134P. These attributes were originally defined to support coax-based AUI transceivers. The BCM53134P integrated transceiver design means these error conditions are eliminated. MIBs intending to support such counters should return a value of 0 (not supported).

Table 24: Indirectly Supported MIB Counters

BCM53134P MIB	Ethernet-Like MIB RFC 1643	Bridge MIB RFC 1493	MIB II Interface RFC 1213/1573	RMON MIB RFC 1757
RxErrorPkts = RxAlignmentErrors + RxFCSErrors + RxFragments + RxOversizePkts + RxJabbers	—	—	ifInErrors	—
—	—	—	ifInUnknownProtos	—
RxGoodPkts = RxUnicastPkts + RxMulticastPkts + RxBroadcastPkts	—	dot1dTpPortIn Frames	—	—
DropEvents = RxDropPkts + TxDropPkts	—	—	—	etherStatsDrop Events
RxTotalPkts = RxGoodPkts + RxErrorPkts	—	—	—	etherStatsPkts
RxCRCAlignErrors = RxCRCSErrors + RxAlignmentErrors	—	—	—	etherStatsCRCAlign Errors
—	dot3StatsSQETest Errors	—	—	—
RxFramesTooLong = RxOversizePkts + RxJabber	dot3StatsFrameToo Longs	—	—	—
TxGoodPkts = TxUnicastPkts + TxMulticastPkts + TxBroadcastPkts	—	dot1dTpPortOut Frames	—	—
TxErrorPkts = TxExcessiveCollision + TxLateCollision Note 1	—	—	ifOutErrors	—

NOTE: The number of packets transmitted from a port that experienced a late collision or excessive collisions. While some media types operate in half-duplex mode, frames that experience carrier sense errors are also summed in this counter. The BCM53134P integrated design means this error condition is eliminated.

Table 25: BCM53134P Supported MIB Extensions

BCM53134P MIB	Appropriate Standards Reference
RxSACChanges	IEEE 802.3u Clause 30—Repeater Port Managed Object Class aSourceAddressChanges.
RxExcessSizeDisc	The BCM53134P cannot store packets in excess of 1536 bytes (excluding preamble/SFD, but inclusive of FCS). This counter indicates packets that were discarded by the BCM53134P due to excessive length.
RxPausePkts	IEEE 802.3x Clause 30—PAUSE Entity Managed Object Class aPAUSEMACCtrlFramesReceived.

Table 25: BCM53134P Supported MIB Extensions (Continued)

BCM53134P MIB	Appropriate Standards Reference
RxSymbolErrors	IEEE 802.3u Clause 30—Repeater Port Managed Object Class aSymbolErrorDuringPacket.
TxFramelnDisc	Internal diagnostic use for optimization of flow control and buffer allocation algorithm.
TxPausePkts	The number of PAUSE events at a given port.

3.6 Integrated High-Performance Memory

The BCM53134P embed a high-performance SRAM for storing the following:

- Packet data: 128 KB
- ARL table
- VLAN table
- TX queues
- Descriptors

This eliminates the need for external memory and allows for the implementation of extremely low-cost systems.

The internal RAM controller efficiently executes memory transfers and achieves nonblocking performance for stand-alone 5-port applications.

3.7 Switch Controller

The core of the BCM53134P devices is a cost-effective and high-performance switch controller. The controller manages packet forwarding between the MAC receive and transmit ports through the frame buffer memory with a store and forward architecture. The switch controller encompasses the functions of buffer management, memory arbitration, and transmit descriptor queuing.

3.7.1 Buffer Management

The frame buffer memory is divided into pages (units of data consisting of 256 bytes each). Each received packet may be allocated more than one page. For example, six pages are required to store a 1522-byte frame. Frame data is stored in the buffer memory as the packet is received. After reception, the frame is queued to the egress port(s) transmit queue. This list tracks the transmission of the packet. After successful packet transmission, the buffer memory is released to the free buffer pool.

3.7.2 Memory Arbitration

Processes requesting access to the internal memory include the receive and transmit frame data handlers, address resolution, the VLAN lookup, learning and aging functions, egress descriptor update, and output-port queue managers. These processes are arbitrated to provide fair access to the memory and minimize the latency of critical processes to provide a fully nonblocking solution.


3.7.3 Transmit Output Port Queues

Frames are maintained in the egress port using a linked list. Two levels of linked lists are used to maintain one output queue (see [Figure 25](#)). The first level is the TXQ linked list, and the second level is the buffer tag linked list. The TXQ linked list is used to maintain frame TC order for each port. For each frame, the buffer tag linked list is used to maintain the order of the buffer pages corresponding to each frame.

Each egress port supports up to six transmit queues for servicing Quality of Service (QoS). All six transmit queues share the 512 entries of the TXQ table. The TXQ table is maintained as a linked list, and each node in the TXQ uses one entry in the TXQ table. The TXQ size for each priority can be programmed to up to 512 entries.

When the QoS function has been turned off, the switch controller maintains one output queue for each egress port. The TXQ table is maintained in a per-port individual internal memory. Each node in the queue represents a pointer that points to a frame buffer tag. Each buffer tag includes frame information and a pointer to the next buffer tag. Each buffer tag has an associated page allocated in the frame buffer. For a packet with a frame size larger than 256 bytes, multiple buffer tags are required. For instance, a 9720-byte jumbo frame requires 38 buffer tags for handling the frame.

Figure 25: TXQ and Buffer Tag Structure



Chapter 4: System Interfaces

4.1 Overview

The BCM53134P include the following interfaces:

- [“Copper Interface”](#)
- [“Frame Management Port Interface” on page 84](#)
- [“WAN Interface” on page 85](#)
- [“Configuration Pins” on page 85](#)
- [“Programming Interfaces” on page 85](#)
- [“LED Interfaces” on page 110](#)
- [“Digital Voltage Regulator \(LDO\)” on page 115](#)

Each interface is discussed in detail in these sections.

4.2 Copper Interface

The internal PHYs transmit and receive data using the analog copper interface. This section discusses the following topics:

- [“Auto-Negotiation”](#)
- [“Line-side \(Remote\) Loopback Mode” on page 84](#)

4.2.1 Auto-Negotiation

The BCM53134P negotiate a mode of operation over the copper media using the auto-negotiation mechanism defined in the IEEE 802.3u and IEEE 802.3ab specifications. When the auto-negotiation function is enabled, the BCM53134P automatically choose the mode of operation by advertising its abilities and comparing them with those received from its link partner. The BCM53134P can be configured to advertise the following modes:

- 1000BASE-T full-duplex and/or half-duplex.
- 100BASE-TX full-duplex and/or half-duplex.
- 10BASE-T full-duplex and/or half-duplex.

The transceiver negotiates with its link partner and chooses the highest common operating speed and duplex mode, commonly referred to as highest common denominator (HCD). Auto-negotiation can be disabled by software control, but is required for 1000BASE-T operation.

4.2.2 Line-side (Remote) Loopback Mode

The line-side loopback mode allows the testing of the copper interface from the link partner. This mode is enabled by setting bit 15 of the Miscellaneous Test register. The MDI receive packet is passed through the PCS and sent back out as the MDI transmit packet. The PCS receive data appears on the internal MAC interface.

4.3 Frame Management Port Interface

The dedicated frame management port provides high-speed connection to transfer management packets to an external management agent. For more information about frame management, see [“Frame Management” on page 72](#). The port is configurable to RGMII using strap pins or software configuration.

NOTE: The Frame Management port interface supports only full-duplex mode.

NOTE: The BCM53134P supports EEE features for external PHYs connected on the IMP and GMII (port5) only through the GMII interface.

4.3.1 RGMII Interface

The Reduced Gigabit Media Independent Interface (RGMII) serves as a digital data interface between the BCM53134P and an external management entity or an external PHY to provide additional data port capacity. Transmit and receive data is clocked on the rising and falling edge of the clocks. This reduces the number of data signals crossing the MAC interface without affecting the data transmission rate. The RGMII transmits data synchronously using the TXD[3:0] and RXD[3:0] data signals. The BCM53134P offers either 3.3V/2.5V or 1.8V/1.5V RGMII interface with an external device.

4.4 WAN Interface

The BCM53134P provides one SGMII interface (port 5) for WAN port or integrated gateways application. Port 5 is also capable of being second IMP port only when Port 8 is configured as first IMP port. (for example, Port 5 cannot be an IMP port by itself.)

Port 5 can be configured as an IMP port in the BCM53134P if dual-IMP is enabled.

NOTE: The BCM53134P supports EEE features for external PHYs connected on the IMP through the GMII interface.

4.5 Configuration Pins

Initial configuration of the BCM53134P takes place during power-on/reset by loading internal control values from hardware strap pins. The value of the pin is loaded when the reset sequence completes, and the pin transitions to normal operation. Pull-up or pull-down resistors can be added to these pins to control the device configuration. If the pins are left floating, the default value is determined based on the internal pull-up or pull-down configuration. See [“Signal Descriptions” on page 117](#) for more information.

4.6 Programming Interfaces

The BCM53134P can be programmed using the SPI interface or the EEPROM interface. The interfaces share a common pin set that is configured using the CPU_EEPROM_SEL strap pin. The [“SPI-Compatible Programming Interface”](#) provides access for a general-purpose microcontroller, allowing read and write access to the internal BCM53134P register space. It is configured to be compatible with the Motorola Serial Peripheral Interface (SPI) protocol. Alternatively, the [“EEPROM Interface” on page 100](#) can be connected to an external EEPROM for writing register values upon power-up initialization.

The internal address space of the BCM53134P devices is broken into a number of pages. Each page groups a logical set of registers associated with a specific function. Each page provides a logical address space of 256 bytes, although, in general, only a small portion of the address space in each page is utilized.

An explanation follows for using the serial interface with an SPI-compatible CPU ([“SPI-Compatible Programming Interface”](#)) or an EEPROM ([“EEPROM Interface” on page 100](#)). Either mode can be selected with the CPU_EEPROM_SEL strap pin. Either mode has access to the same register space.

4.6.1 SPI-Compatible Programming Interface

One way to access the BCM53134P internal registers is to use the SPI-compatible interface. This four-pin interface is designed to support a fully functional, bi-directional Motorola serial peripheral interface (SPI) for register read/write accesses. The maximum speed of operation is 25 MHz. The SPI interface shares pins with the EEPROM interface. To select the SPI interface, pull up or float the CPU_EEPROM_SEL pin. (The internal pull-up resistor defaults SPI interface over EEPROM interface.)

The SPI is a four-pin interface comprises the following:

- Device select (\overline{SS} : slave select, input to BCM53134P)
- Device clock (SCK: which operates at speeds up to 25 MHz, input to BCM53134P)
- Data write line (MOSI: Master Out/Slave In, input to BCM53134P)
- Data read line (MISO: Master In/Slave Out, output from BCM53134P)

NOTE: All the RoboSwitch™ SPI interfaces are designed to operate in slave mode. Therefore, the SCK and SS signals are driven by the external master host device when accessing the BCM53134P registers. For more detailed descriptions, refer to the *Motorola SPI spec MC68HC08AS20-Rev. 4.0*.

NOTE: The internal SPI bus can be in a busy state in which it does not allow an external device to access the switch registers during the power-up. The SPI bus is held in the busy state while the internal 8051 controller is being initialized. This busy period can be as long as 350 ms from the de-asserting of the RESET line.

4.6.1.1 \overline{SS} : Slave Select

The \overline{SS} signal is used to select a slave device and to indicate the beginning of transmission. The BCM53134P SPI interface operates in the clock phase one (CPHA = 1) transmission format. In this format, the SS signal is driven active low while the SCK signal is high, and remains low throughout the transmission including multiple-byte transfers. The minimum time requirement between \overline{SS} operation is 200 ns.

4.6.1.2 SCK: Serial Clock

The serial clock SCK maximum operating frequency is 25 MHz for the BCM53134P family of devices. The SCK is used to clock data into and out of the Slave ROBO device. The SCK signal is expected to remain high when the interface is idle. This is because the BCM53134P SPI design is based on CPOL = 1 (Clock Polarity = 1). This is not programmable on BCM53134P. The BCM53134P is designed so that data is driving by the falling edge and sampling by the rising edge of the SCK clock. This clock is not a free-running clock; it is generated only during a data transaction, and remains high when the clock is idle.

4.6.1.3 MOSI: Master Output Slave Input

The MOSI signal is used by the master device to transmit the data to the slave device. The data is put on the bus and is expected to be clocked in by a rising edge of the SCK clock signal. This line is used to issue a command and to set the register page and address value of read/write operations.

4.6.1.4 MISO: Master Input Slave Output

The MISO signal is used by the Slave device to output the data to the master device. The data is put on the bus and is expected to be clocked out by a rising edge of the SCK clock signal. This line is used to transmit the status and the content of the register of read operation.

A layer of protocol is added to the basic SPI definition to facilitate transfers from the BCM53134P. This protocol establishes the definition of the first 2 bytes issued by the master to the BCM53134P slave during an SPI transfer. The first byte issued from the SPI master in any transaction is defined as a command byte, which is always followed by a register address byte, and any additional bytes are data bytes.

The SPI mode supports two different access mechanisms, normal SPI and fast SPI, determined by the content of the command byte. Table 26 shows the normal SPI command byte, and Table 27 shows the Fast SPI command byte. These two mechanisms should not be mixed in an implementation; the CPU should always initiate transfers consistently with only one of the two mechanisms.

Table 26: Normal SPI Command Byte

0	1	1	MODE = 0	CHIP ID 2 (MSB)	CHIP ID 1	CHIP ID 0 (LSB)	Read/Write (0/1)
---	---	---	----------	-----------------	-----------	-----------------	------------------

Table 27: Fast SPI Command Byte

Byte Offset (MSB)	Byte Offset	Byte Offset (LSB)	MODE = 1	CHIP ID 2 (MSB)	CHIP ID 1	CHIP ID 0 (LSB)	Read/Write (0/1)
-------------------	-------------	-------------------	----------	-----------------	-----------	-----------------	------------------

The MODE bit (bit 4) of the command byte determines the meaning of bits 7:5. If bit 4 is a 0, it is a normal SPI command byte, and bits 7:5 should be defined as 011b. If bit 4 is a 1, bits 7:5 indicate a fast SPI command byte, and bits 7:5 indicate the byte offset into the register that the BCM53134P starts to read from (byte offsets are not supported for write operations).

In command bytes, bits [3:1] indicate the CHIP ID to be accessed. Because the BCM53134P operates as a single-chip system, the CHIP ID is 000.

NOTE: The SS# signal must also be active for any BCM53134P device to recognize that it is being accessed.

Bit 0 of the command byte is the R/W signal (0 = Read, 1 = Write) and determines the transmission direction of the data.

The byte following the command byte is an 8-bit register address. Initially, this sets the page address, followed by another command byte that contains the register base address in that page, which is used as the location to store the next byte of data received in the case of a write operation, or the next address from which to retrieve data in the case of a read operation. This base address increments as each byte of data is transmitted/received, allowing a contiguous block data from a register to be stored/read in a single transmission. When the fast SPI command byte mode is used, the actual start location of a read operation can be modified by the offset contained in bits 7:5 of the command byte. Reading/writing data from/to separate registers, even if those registers are contiguous in the current page, must be performed by supplying a new command byte and register address for each register, with the address as defined in the appropriate page register map.

Noncontiguous blocks are also stored/read through the use of multiple transmissions, which allow a new command byte and register base address to be specified. The \overline{SS} signal must remain low for the entire read or write transaction, as shown in Figure 26 and Figure 27, with the transaction terminated by the deassertion of the \overline{SS} line by the master.

Figure 26: SPI Serial Interface Write Operation




Figure 27: SPI Serial Interface Read Operation






Figure 28 shows the typical connection block diagram for SPI interface with/without external PHY devices.

Figure 28: SPI Interface Without External PHY Device



4.6.1.5 External PHY Registers

The BCM53134P also uses the MDIO/MDC interface for polling registers of an external PHY. In this case, the MDIO/MDC interface polls the external PHY registers pulling the data internal to the BCM53134P. Then, the external PHYs and retrieved from the register data using the SPI interface. The MDIO/MDC interface is not used as a method to access internal PHY registers. This must be done using the SPI interface. In this case, the MDC/MDIO interface of the BCM53134P is acting as a master.

Figure 29: Accessing External PHY Registers

4.6.1.6 Reading and Writing BCM53134P Registers Using SPI

BCM53134P internal register read and write operations are executed by issuing a command followed by multiple accesses of the SPI registers in the BCM53134P. There are three SPI interface registers in the BCM53134P that are used by the master device to access the internal switch registers. The SPI interface registers are:

- SPI Page register (page: global, address: FFh): used to specify the value of the specific register pages.
- SPI Data I/O register (page: global, address: F0h): used to write and read the specific register's content.
- SPI Status Register (page: global, address: FEh): used to check for an operation completion.
 - Bit 7 = SPIF, SPI read/write complete flag
 - Bit 6 = Reserved
 - Bit 5 = RACK, SPI read data ready acknowledgment
 - Bit 4:3 = Reserved
 - Bit 2 = MDIO_Start, Start/Done MDC/MDIO operation
 - Bit 1 = Reserved
 - Bit 0 = Reserved

The BCM53134P SPI interface supports the following operating modes:

- Normal read mode
- Fast read mode
- Normal write mode


NOTE: The RoboSwitch family does not support fast-write mode.

4.6.1.7 Normal Read Operation

Normal Read operation consists of five transactions (five \overline{SS} operations):

1. Issue a Normal Read Command (opcode = 0x60) to poll the SPIF bit in the SPI Status register (0xFE) to determine the operation can start.
2. Issue a Normal Write command (opcode = 0x61) to write the register page value into the SPI Page register 0xFF.
3. Issue a Normal Read command (opcode = 0x60) to setup the required RoboSwitch register address.
4. Issue a Normal Read command (opcode = 0x60) to poll the RACK bit in the SPI status register (0xFE) to determine the completion of read (register content gets loaded in SPI Data I/O register).
5. Issue a Normal Read command (opcode = 0x60) to read the specific registers' content placed in the SPI Data I/O register (0xF0).


Figure 30: Normal Read Operation



Example: Read from 1000BASE-T Control register (page 10h, offset 12h).


1. Issue a Normal Read command (opcode = 0x60) to check the SPIF bit in the SPI Status register (0xFE).
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Read Command Byte: 0 1 1 0 0 0 0 (opcode = 0x60).
 - Clock in the SPI Status register address (0xFE).
 - Clock out the SPI Status register value: 0 0 0 0 0 0 0 (SPIF bit 7 = 0).
 - Deassert \overline{SS} while SCK is high idle state.

Figure 31: Normal Read Mode to Check the SPIF Bit of SPI Status Register




2. Issue a Normal Write command (opcode = 0x61) and write the accessed register page value of 0x10 into SPI Page Register (0xFF)—this step is required only if previous read/write was not to/from Page 10h.
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Write Command Byte: 0 1 1 0 0 0 1 (opcode = 0x61).
 - Clock in offset of Page register (0xFF).
 - Clock in the accessed register page value: 0 0 0 1 0 0 0 (Page register: 0x10).
 - Deassert \overline{SS} while SCK is high idle state.

Figure 32: Normal Read Mode to Setup the Accessed Register Page Value



3. Issue a Normal Read command (opcode = 0x60) and write the accessed register address value 0x12, and clock out 8 bits to complete the read cycle, but discard result (this is where the state machine triggers an internal data transfer from address 0x12 to the SPI Data I/O register).
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Read Command Byte: 0 1 1 0 0 0 0 (opcode = 0x60).
 - Clock in the address of accessed register address value (0x12).
 - Clock out eight clocks for the dummy read, and discard results on MISO.
 - Deassert \overline{SS} while SCK is high idle state.


Figure 33: Normal Read Mode to Setup the Accessed Register Address Value (Dummy Read)



NOTE: This dummy read is always eight clock cycles, whether or not it is an 8-bit register.

4. Issue a Normal Read command (opcode = 0x60) to read the SPI Status to check the RACK bit for completion of the register content transfer to the SPI Data I/O register. (This step may be repeated until the proper bit set is read.)
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Read Command Byte: 0 1 1 0 0 0 0 (opcode = 0x60).
 - Clock in offset for SPI Status Register (0xFE): 1 1 1 1 1 1 0.
 - Clock out the content of SPI Status bits.
 - Repeat the polling until the content of SPI Status Register value: 0 0 1 0 0 0 0 (RACK bit 5 = 1).
 - Deassert \overline{SS} while SCK is high idle state.


Figure 34: Normal Read Mode to Check the SPI Status for Completion of Read



5. Issue a Normal Read command (opcode = 0x60) to read the data from the SPI Data I/O register:

- Assert \overline{SS} while SCK is high idle state.
- Clock in Command Byte: 0 1 1 0 0 0 0 (opcode = 0x60).
- Clock in offset of SPI Data I/O Register (0xF0).
- Clock out first data byte on MISO line: 0 0 0 0 0 0 0 (byte 0 = bit 7 to bit 0 = MSB to LSB).
- Clock out next byte (in this case, last) on MISO line: 0 0 0 0 1 1 1 0 (byte 1 = bit 15 to bit 8).
- [Continue if more bytes].
- Deassert \overline{SS} while SCK is high idle state.

Figure 35: Normal Read Mode to Obtain the Register Content



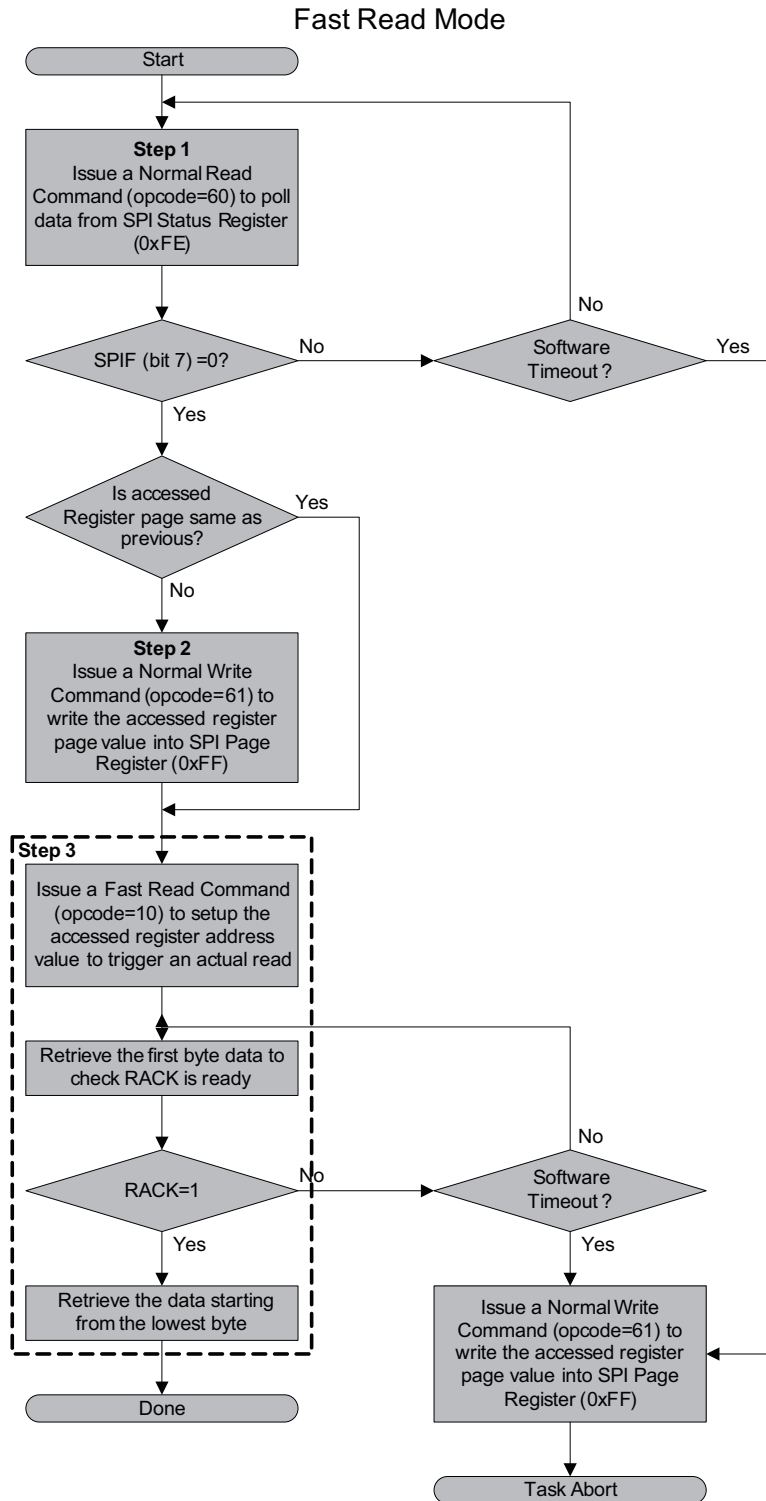
4.6.1.8 Fast Read Operation

Fast Read operation consists of three transactions (three \overline{SS} operations):

1. Issue a Normal Read Command (opcode = 0x60) to poll the SPIF bit in the SPI Status Register (0xFE) to determine the operation can start.
2. Issue a Fast Read command (opcode = 0x10) to setup the accessed Register Page value into the Page register (0xFF).
3. Issue a Fast Read command (opcode = 0x10) to setup the accessed register address value, to trigger an actual read, and retrieve the accessed register content till the completion.

Fast Read mode process is different from Normal Read mode, once the switch receives a fast read command followed by the register page and address information, the status and the data (register content) will be put on the MISO line without going through the SPI Status register or SPI Data I/O register. Once RACK bit of the bytes following the Fast Read command with Address information is recognized the register content will be put on MISO line immediately following the byte with RACK bit set. The Fast Read process is described in the following paragraphs with a flowchart followed by a step-by-step description.


Figure 36: Fast Read Operation



Example: Read from 1000BASE-T Control register (page 10h, offset 12h).


1. Issue a Normal Read command (opcode = 0x60) to check the SPIF bit in the SPI Status register (0xFE).
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Read Command Byte: 0 1 1 0 0 0 0 (opcode = 0x60).
 - Clock in the SPI Status register address (0xFE).
 - Clock in the accessed register page value: 0 0 0 0 0 0 0 (SPIF bit 7 = 0).
 - Deassert \overline{SS} while SCK is high idle state.

Figure 37: Normal Read Mode to Check the SPIF Bit of SPI Status Register




2. Issue a Normal Write command (opcode = 0x61) and write the accessed register page value of 0x10 in to SPI Page Register(0xFF)—this step is required only if previous read/write was not to/from Page 10h.
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Fast Read Command Byte: 0 11 0 0 0 0 0 1 (opcode = 0x61).
 - Clock in offset of Page register (0xFF).
 - Clock in the accessed register page value: 0 0 0 1 0 0 0 0 (Page register: 0x10).
 - Deassert \overline{SS} while SCK is high idle state.

Figure 38: Fast Read Mode to Setup New Page Value



3. Issue a Fast Read command (opcode = 0x10), followed by the Address of the accessed register (0x12), check for a read completion by checking the RACK bit in the SPI Status register, and finally clock out the read data.
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Fast Read Command Byte: 0 0 0 1 0 0 0 0 (opcode = 0x10).
 - Clock in the Address of accessed register (0x12).
 - Clock out Bytes Until Bit 0 or Bit 1 = 1: 0 0 0 0 0 0 0 1 (RACK bit 0 = 1).
 - Clock out first data byte: 0 0 0 0 0 0 0 0 (Byte 0 = Bit 7 to Bit 0).
 - Clock out next data (in this case, last) byte = 0 0 0 0 1 1 1 0 (Byte 1 = Bit 15 to Bit 8).
 - [Continue if more bytes].
 - Deassert \overline{SS} while SCK is high idle state.

Figure 39: Fast Read to Read the Register



NOTE: There is an erratum on the RACK output timing in Fast Read mode. The RACK (bit 0) must be sampled prior to toggling the clock to shift out the bit 0.

4.6.1.9 Normal Write Operation

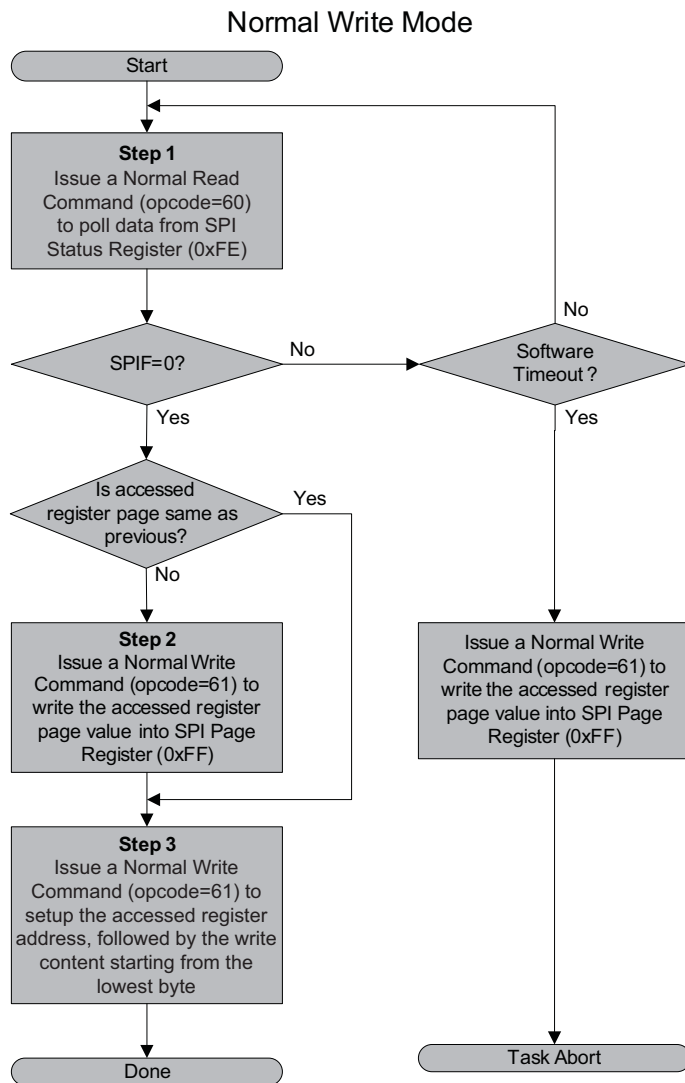
Normal Write operation consists of three transactions (three \overline{SS} operations):

1. Issue a Normal Read Command (opcode = 0x60) to poll the SPIF bit in the SPI Status register (0xFE) to determine the operation can start.
2. Issue a Normal Write command (opcode = 0x61) to setup the accessed register page value into the page register (0xFF).
3. Issue a Normal Write command (opcode = 0x61) to setup the accessed register address value, followed by the write content starting from a lower byte.

The Normal Write Mode process is described in the following paragraphs with a flowchart followed by a step-by-step description.

NOTE: The RoboSwitch does not support Fast Write Mode.


Figure 40: Normal Write Operation



Example: 0x1600h is written to 1000BASE-T Control register (page 0x10, offset 0x12).


1. Issue a Normal Read command (opcode = 0x60) to check the SPIF bit in the SPI Status register (0xFE).
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Read Command Byte = 0 1 1 0 0 0 0 0 (opcode = 0x60).
 - Clock in the SPI Status register address (0xFE).
 - Clock in the accessed register page value = 0 0 0 0 0 0 0 0 (SPIF bit 7=0).
 - Deassert \overline{SS} while SCK is high idle state.

Figure 41: Normal Read Mode to Check the SPIF Bit of SPI Status Register




2. Issue a Normal Write command (opcode = 0x61) and write the accessed register page value of 0x10 into SPI Page register (0xFF)—this step is required only if previous read/write was not from/to Page 0x10.
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Write Command Byte = 0 1 1 0 0 0 0 1 (opcode = 0x61).
 - Clock in offset of Page register (0xFF).
 - Clock in 1 byte of the accessed register page value (Page register 0x10).
 - Deassert \overline{SS} while SCK is high idle state.

Figure 42: Normal Write to Setup the Register Page Value



3. Issue a Normal Write command (opcode = 0x61) and write the Address of the accessed register followed by the write content starting from a lower byte.
 - Assert \overline{SS} while SCK is high idle state.
 - Clock in a Normal Write Command Byte = 0 1 1 0 0 0 1 (opcode = 0x61).
 - Clock in Offset of Address of accessed register (0x12).
 - Clock in lower data byte first = 0 0 0 0 0 0 0 (byte 0 = bit 7 to bit 0).
 - Clock in upper data byte next = 0 0 0 1 0 1 1 0 (byte 1 = bit 15 to bit 8).
 - [Continue if more bytes].
 - Deassert \overline{SS} while SCK is high idle state.

Figure 43: Normal Write to Write the Register Address Followed by Written Data




4.6.2 EEPROM Interface

The BCM53134P can be connected using the serial interface to a low-cost external serial EEPROM, enabling it to download register-programming instructions during power-on initialization. For each programming instruction fetched from the EEPROM, the instruction executes immediately and affects the register file.

During the chip-initialization phase, the switch identifies automatically the type of EEPROM it is connected to, then the data is sequentially read-in from the EEPROM after the internal memory has been cleared. The first data read-in is the HEADER and it matches a predefined magic code. In the case where the HEADER data does not match the instruction fetch, the process stops, and the EEPROM controller treats it as if no EEPROM exists. If the magic code matches, the fetch instruction process continues until it reaches the instruction length defined in the HEADER.

Figure 44: Serial EEPROM Connection



4.6.2.1 EEPROM Format

The EEPROM should be configured to x16 word format. The header contains key and length information as shown in [Table 28](#). The actual data stored in the EEPROM is byte-swapped as shown in [Table 29](#).

- Upper 5 bits are magic code 15h, which indicates that valid data follows.
- Bit 10 is for speed indication. A 0 means normal speed. A 1 indicates speedup. The default is 0.
- Lower 10 bits indicate the total length of all entries. For example:
 - 93C46 up to 64 words
 - 93C56 up to 128 words
 - 93C66 up to 256 words
 - 93C86 up to 1024 words

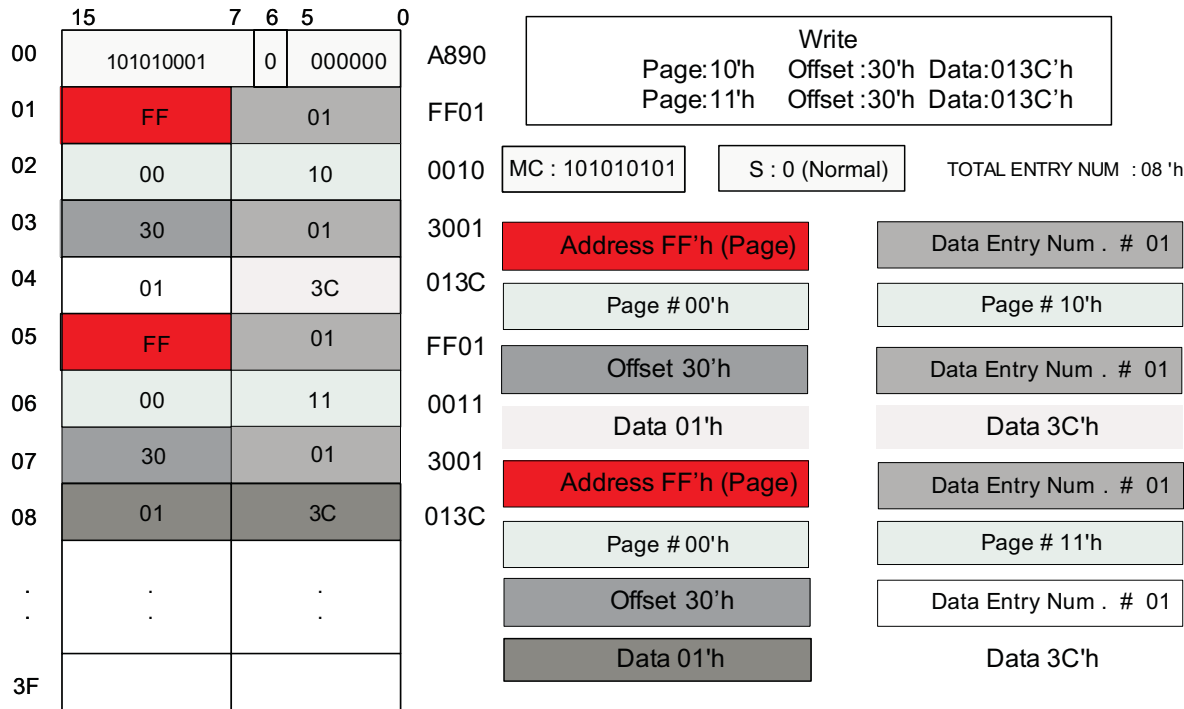
Table 28: EEPROM Header Format

Bits [15:11]	Bit 10	Bits [9:0]
Magic code, 15h	Speed	Total entry number 93C46: 0 to 63 93C56: 0 to 127 93C66: 0 to 255 93C86: 0 to 1023

Table 29: EEPROM Contents

Bits [7:0]	Bits [15:11]	Bit 10	Bits [9:8]
Total entry number	Magic code, 15h	Speed	Total entry number

Figure 45 on page 102 shows an EEPROM programming example.

Figure 45: EEPROM Programming Example

EEPROM Serial Stream: 90 A8 01 FF 10 00 01 30 3C 01 01 FF 11 00 01 30 3C 01

4.6.3 Serial Flash Interface

The BCM53134P offers a serial flash interface to store program code for the internal microcontroller (BCM8051 processor). The BCM53134P detects a flash memory device automatically and downloads the memory contents upon power-up. The main purpose of the stored code is to configure and run the power savings mode, such as any application that the user wishes to run that can fit in the internal BCM8051 memory. The embedded BCM8051 microcontroller has 32 KB of SRAM and 16 KB of ROM. The interface comprises four signal pins: chip select (FCS), Flash clock (FCLK), Flash Serial Out (FSO), and Flash Serial In (FSI). The maximum frequency of the interface is 25 MHz.

4.6.4 MDC/MDIO Interface

The BCM53134P offers an MDC/MDIO interface for accessing the switch registers as well as the PHY registers. An external management entity can access the switch registers through this interface when the SPI interface is not used. (when the SPI clock is in idle mode.) The switch registers are accessed through the Pseudo PHY interface, and the PHY registers are accessed directly by using direct PHY addresses from 0x00 ~ 0x04 (for example, the PHYs are mapped to Register page 0x10 ~ 0x14 when it is accessed through the SPI interface).

An external PHY can be connected to the GMII interface of the IMP port and Port 5. Through the SPI interface, by accessing the Page 88h and Page 85h, the external PHY MII registers can be accessed.

NOTE: The PHY registers are not accessible through the Pseudo PHY operation.

4.6.4.1 MDC/MDIO Interface Register Programming

The BCM53134P are designed to be fully compliant with the MII clause of the IEEE 802.3u Ethernet specification. The MDC pin of the BCM53134P sources a 2.5 MHz clock. Serial bidirectional data transmitted using the MDIO pin is synchronized with the MDC clock. Each MII read or write instruction is initiated by the BCM53134P and contains the following:

- **Preamble (PRE).** To signal the beginning of an MII instruction after reset, at least 32 consecutive 1-bits must be written to the MDIO pin. A preamble of 32 1-bits is required only for the first read or write following reset. A preamble of fewer than 32 1-bits causes the remainder of the instruction to be ignored.
- **Start of Frame (ST).** A 01 pattern indicates that the start of the instruction follows.
- **Operation Code (OP).** A read instruction is indicated by 10, while a write instruction is indicated by 01.
- **PHY Address (PHYAD).** A 5-bit PHY address follows, with the MSB transmitted first. The PHY address allows a single MDIO bus to access multiple PHY chips.
- **Register Address (REGAD).** A 5-bit register address follows, with the MSB transmitted first.
- **Turnaround (TA).** The next bit times are used to avoid contention on the MDIO pin when a read operation is performed. When a write operation is being performed, 10 must be sent by the BCM53134P chip during these two bit-times. When a read operation is being performed, the MDIO pin of the BCM53134P must be put in a high-impedance state during these bit times. The external PHY drives the MDIO pin to 0 during the second bit time.
- **Data.** The last 16 bits of the Instruction are the actual data bits. During a write operation, these bits are written to the MDIO pin with the most significant bit (MSB) transmitted first by the BCM53134P. During a read operation, the data bits are driven by the external PHY with the MSB transmitted first.

4.6.4.2 Pseudo-PHY

The MDC/MDIO can be used by an external management entity to read/write register values internal to the BCM53134P. This mode offers an alternative programming interface to the chip. The BCM53134P operate in slave mode with a PHY address of 30d. The following figures show the register setup flow chart for accessing the registers using the MDC/MDIO interface.

Figure 46: Pseudo-PHY MII Register Definitions

IEEE Reserved			Reg 0
			Reg 15
Page Number	Reserved	A	Reg 16
Register Address	Reserved	OP	Reg 17
Access Status			Reg 18
Reserved			
Access register bits [15:0]			Reg 24
Access register bits [31:16]			Reg 25
Access register bits [47:32]			Reg 26
Access register bits [63:48]			Reg 27
Reserved			Reg 31

Figure 47: Pseudo-PHY MII Register 16: Register Set Access Control Bit Definition

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	Bit #
Page Number								Reserved							A	Reg 16

bits [15:8] => Page number (RW)

bits [7:1] => Reserved

Bit 0 => Register set MDC/MDIO access enable (RW)

NOTE: The bit 0 (MDC/MDIO Access Enable) in register 16 should be released (set to 0) after a transaction is completed. This allows the SPI interface to access the switch register if required.

Figure 48: Pseudo-PHY MII Register 17: Register Set Read/Write Control Bit Definition

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	Bit #
Register Address								Reserved							OP	Reg 17

bits [15:8] => Register address (RW)

bits [7:2] => Reserved

bits [1:0] => OP code (RW/SC), 00 = No operation
 01 = Write operation
 10 = Read operation
 11 = Reserved

Figure 49: Pseudo-PHY MII Register 18: Register Access Status Bit Definition

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	Bit #
Reserved														E	P	Reg 18

bits [15:1] => Reserved

Bit 1 => Operation Error (RO/LH), when op_code = 2'b11, this bit is set to show operation error.

Bit 0 => Prohibit Access (RO/LH), for Page Number = 8'h1X, which are PHY MII registers.

Figure 50: Pseudo-PHY MII Register 24: Access Register Bit Definition

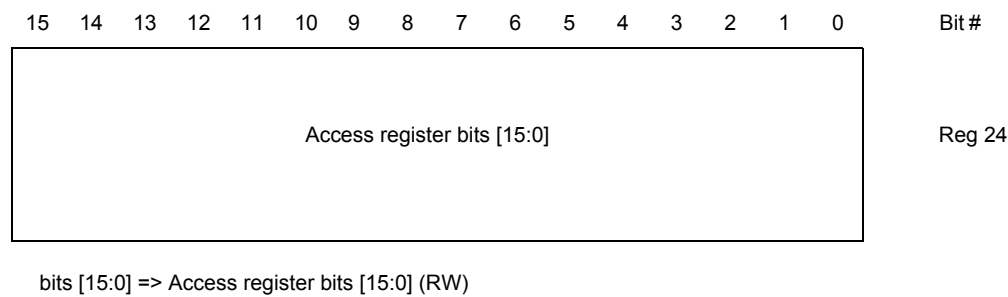


Figure 51: Pseudo-PHY MII Register 25: Access Register Bit Definition

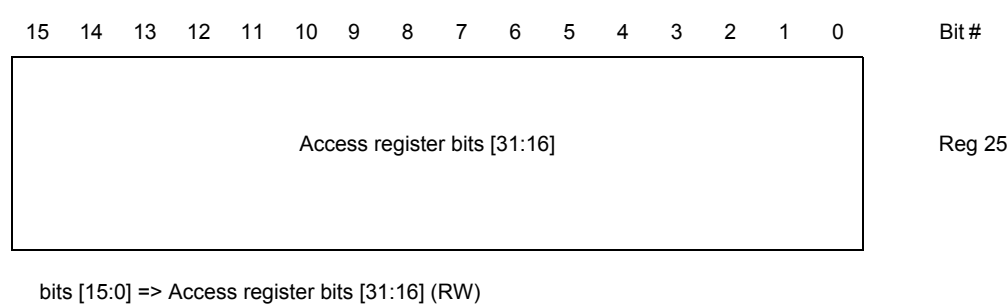


Figure 52: Pseudo-PHY MII Register 26: Access Register Bit Definition

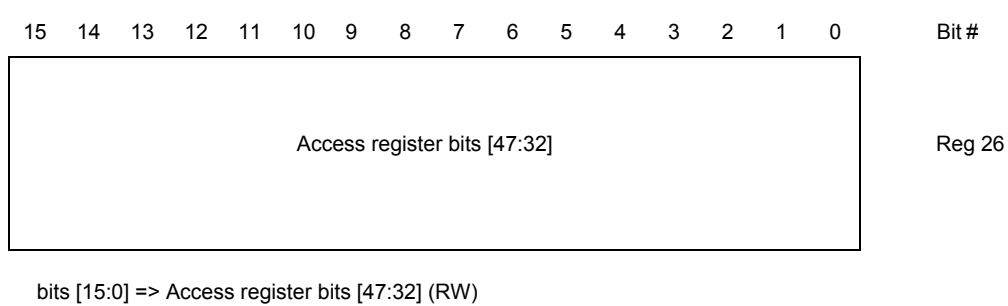


Figure 53: Pseudo-PHY MII Register 27: Access Register Bit Definition




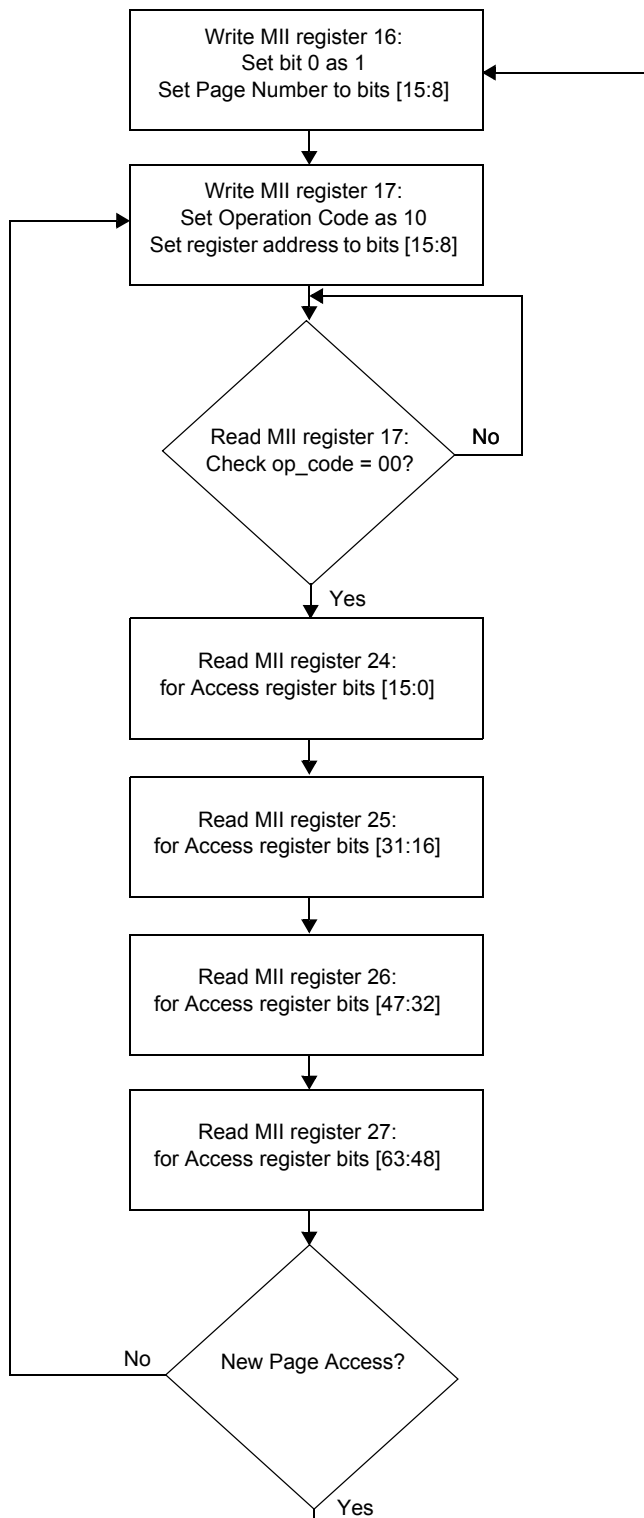
Figure 54: Read Access to the Register Set using the Pseudo-PHY (PHYAD = 11110) MDC/MDIO Path

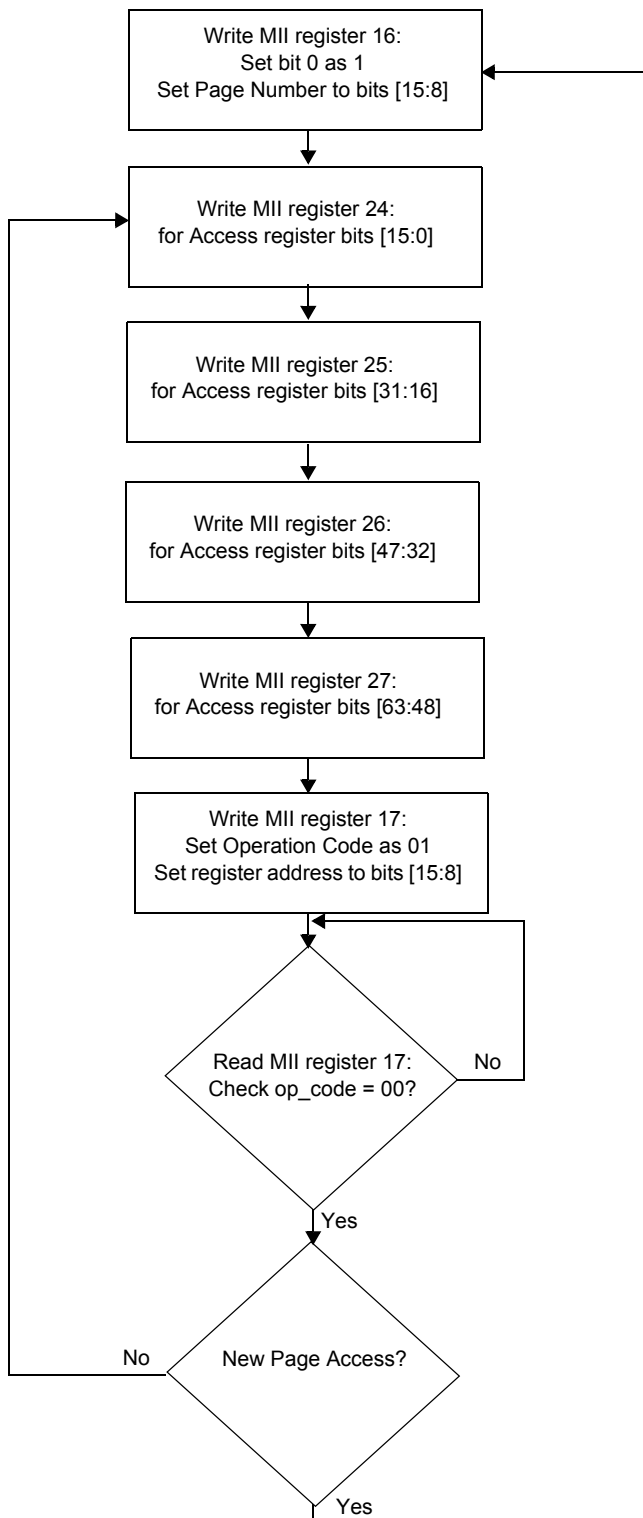
Figure 55: Write Access to the Register Set using the Pseudo-PHY (PHYAD = 11110) MDC/MDIO Path

Table 30 summarizes the complete management frame format.

Table 30: MII Management Frame Format

Operation	PRE	ST	OP	PHYAD	REGAD	TA	Data	Direction
Read	1 ... 1	01	10	AAAAA	RRRRR	ZZ Z0	Z ... Z D ... D	Driven by master Driven by slave
Write	1 ... 1	01	01	AAAAA	RRRRR	10	D ... D	Driven to master

See “MDC/MDIO Interface” on page 102 for more information regarding the timing requirements.

4.7 LED Interfaces

The BCM53134P provides flexible, programmable per-port status of various functions. The user can select predefined LED displays per port by setting the LED_MODE strap pins. Alternatively, the user can program different display functions and different sets of display functions for different ports by programming the LED Control registers. Normally the BCM53134P offers a total of 16 LED displays with four functions per port when the WAN (port 5) RGMII interface is **not** used, or a total of eight LED displays with two functions per port when the WAN (port 5) RGMII interface **is** used. The LED interface offers options to display different functions.

The options that are available for the LED display include the following:

- Number of displays per port
- Display functions per port
- LED lighting behavior

The total number of displays per port is based on the GMII_LED_DEL strap-pin option. When the GMII (port 5) is used, the 16-LED display option is not available. The LED[xx] signal allocation for each port is shown in Table 31, 8-LED Display Mode (WANLEDSEL=0).

The options for the number of LED displays are as follows:

- 16 LED displays total (WANLEDSEL = 1).
 - Four display functions per port.
 - LED[0:9] active state will be depends on the strap pin state. The LED active state will follow the same as the strap state. (for example, if the strap pin is pulled up, LED active state is high, and vice versa)
 - LED[10:15] active state is low. These LED pins will be always active-low, regardless of external/internal pull up or pull down.
- Eight LED displays total (WANLEDSEL = 0). This mode is forced when GMII (port 5) is used for data interface.
 - Two display functions per port
 - LED[0:7] active state depends on the strap pin state. The LED active state will follow the same as the strap state. (for example, if the strap pin is pulled up, the LED active state is high. If the strap pin is pulled low, the LED active state is low.)

The options for the different function displays per port are as follows:

- Using default display settings through the strap pins LED_MODE[1:0].
 - The different displays per LED_MODE settings are shown in Table 31. Each LED [xx] signal is assigned to a specific port. For example, if the port 3 and port 2 LED displays are disabled (register Page 00h, Address 16h = 0003), port 0 and port 1 LED display are still from LED pins LED[0~3] (port 0), LED[4~7] (port 1), just as if all four ports were used. If port 1 and port 0 LED displays are disabled (register Page 00h, Address 16h = 001C), port 2 and port 3 are still from LED pins LED[8~11] (port 2), and LED[12~15] (port 3), also just as if all four ports were used.
 - All the ports are displaying a same set of functions based on the strap pin settings. The display is fixed per port and per pin as shown in Table 31.
- Displaying one of two different display settings through programming the LED configuration registers.
 - Option to enable LED display per port (register Page 00h, Address 16h).
 - There are two sets of displays options the user can set through the LED Function Control Register 0 and 1 (register Page 00h, Address 10h, and 12h).
 - Each port can select one of two display functions (register Page 00h, Address 14h).

NOTE: For additional details about how to program LED-related registers, see the BCM53134 application note, *Layout and Design Guide* (53134-AN1xx-R).

The options for different LED lighting behavior (register Page 00h, Address 18h, and 1Ah) are as follows:

- Automatic mode: All the modes indication is in a steady state, except the activity state, which is blinking.
- Blink mode: All the status indication is blinking. The blinking rate can be set through the register page 00h, address 0Fh, bit [2:0].
- ON mode: Forces the LED to be on.

In serial LED pins, the status of the enabled ports is sent out with the highest port number first with the lowest port number last. Within a port, the highest (in terms of bit number) selected functionality in LED Function Control Register 0/1 is sent out first followed by lower (in terms of bit number) selected functionality in LED Function Control Register 0/1.

For parallel LED pins, lower ports are mapped to lower LED pins. Port0 is mapped to the lowest LED pins (LED[3:0] or LED[1:0] depending on 16-bit LED or 8-bit LED mode). Port1 is mapped to LED[7:4] or LED[3:2] (depending on 16-bit or 8-bit LED mode). Within a port, the lower (in terms of bit number) selected functionality is LED function Control Register 0/1 is mapped to lower LED pin.

Table 31: 8-LED Display Mode (WANLEDSEL=0)

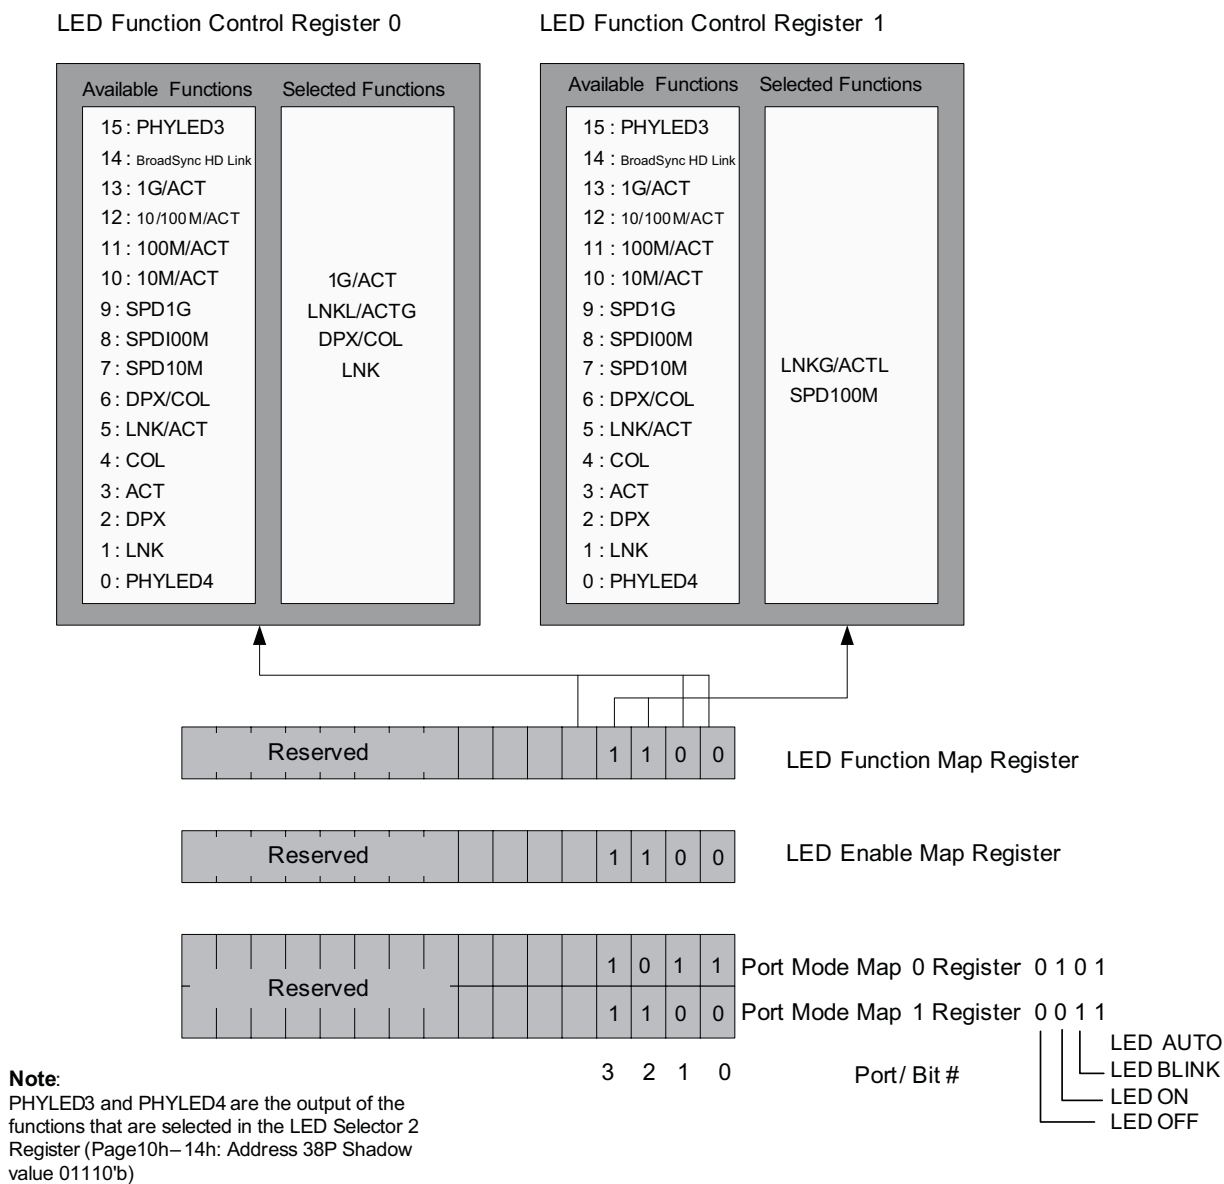
Port 0	Port 1	Port 2	Port 3	LED_MODE= 2'b11	2'b10	2'b01	2'b00
LED[1]	LED[3]	LED[5]	LED[7]	10M/ACT	LNK/ACT	DPX/COL	LNK/ACT
LED[0]	LED[2]	LED[4]	LED[6]	DPX	DPX	PHYLED4	PHYLED4

Table 32: 16-LED Display Mode (WANLEDSEL=1)

Port 0	Port 1	Port 2	Port 3	LED_MODE= 2'b11	2'b10	2'b01	2'b00
LED[3]	LED[7]	LED[11]	LED[15]	1G/ACT	SPD1G	1G/ACT	SPD1G
LED[2]	LED[6]	LED[10]	LED[14]	100M/ACT	SPD100M	10_100/ACT	SPD100M
LED[1]	LED[5]	LED[9]	LED[13]	10M/ACT	LNK/ACT	DPX/COL	LNK/ACT
LED[0]	LED[4]	LED[8]	LED[12]	DPX	DPX	PHYLED4	PHYLED4

Figure 56 shows the LED Interface register structure.

Figure 56: LED Interface Register Structure Diagram

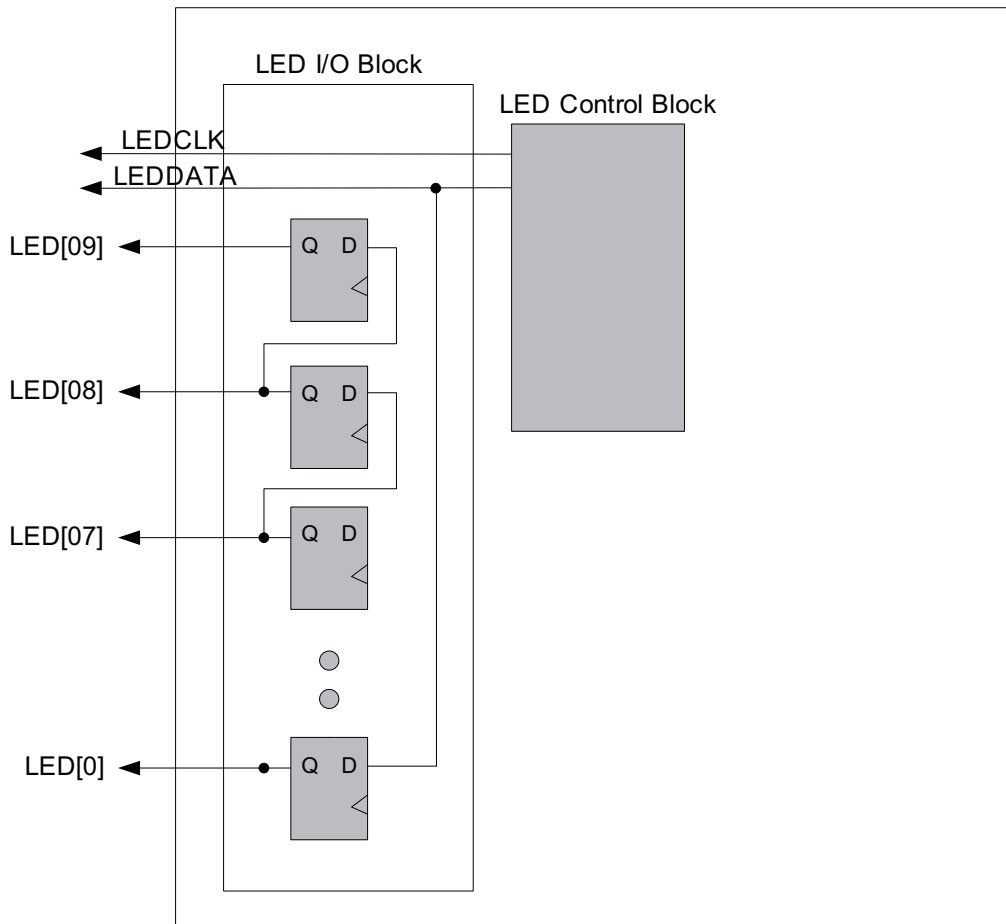


The BCM53134P offers two LED Interfaces: a parallel LED interface and a serial interface. As shown in [Figure 57 on page 113](#), the source of the LED status stream is the same for both interfaces. The status bit stream is based on the programmed register settings. The Parallel LED Interface provides all the shifting and storing of the status internally so that it does not require any external shift registers, but it requires more I/O pins to be connected on the part. The active level of the LED DATA signal can be low or high, depending on the strap pin configuration of each LED signal in the parallel LED interface output. The determination of the active state is shown in “[Dual Input Configuration/LED Output Function](#)” on [page 114](#).

The serial LED interface is output through two pins (LEDDATA and LEDCLK), saving the number of I/O pins but requiring the user to design in the external shift registers. The serial LED interface provides the LED display of port 0 to port 3. The active level of the LED DATA signal is low for each status in the serial LED interface output.

NOTE: In serial LED mode, all LEDs are active low. In parallel LED mode, the LED active state is described in the previous pages.


Figure 57: LED Interface Block Diagram



NOTE: There is a difference between the serial LED and the parallel LED display mode. In parallel LED mode, there is a fixed number of pins assigned to a specific port. In 16 LED display mode, there are four fixed LED signals assigned to each port and four LED signal pins are fixed to each port, even if the user decided to display 3, 2, or 1 LED functions per port. In serial LED mode, only the selected number of LED functions will be shifted out, so there is no gap between each port when the user displays 2, 3, or 4 functions per port. If the user displays/enables only three LED functions per port, then different port LED functions will be shifted out at every 4th function.

A dual LED is used for displaying more than one status using one LED cell. By packing two different-colored LEDs into one holder, a dual LED can display more than two states in one cell. [Figure 58](#) shows typical dual LED usage. The green LED displays LNKG/ACT status, while the yellow LED displays LNKF/ACT status.

Figure 58: Dual LED Usage Example



4.7.1 Dual Input Configuration/LED Output Function

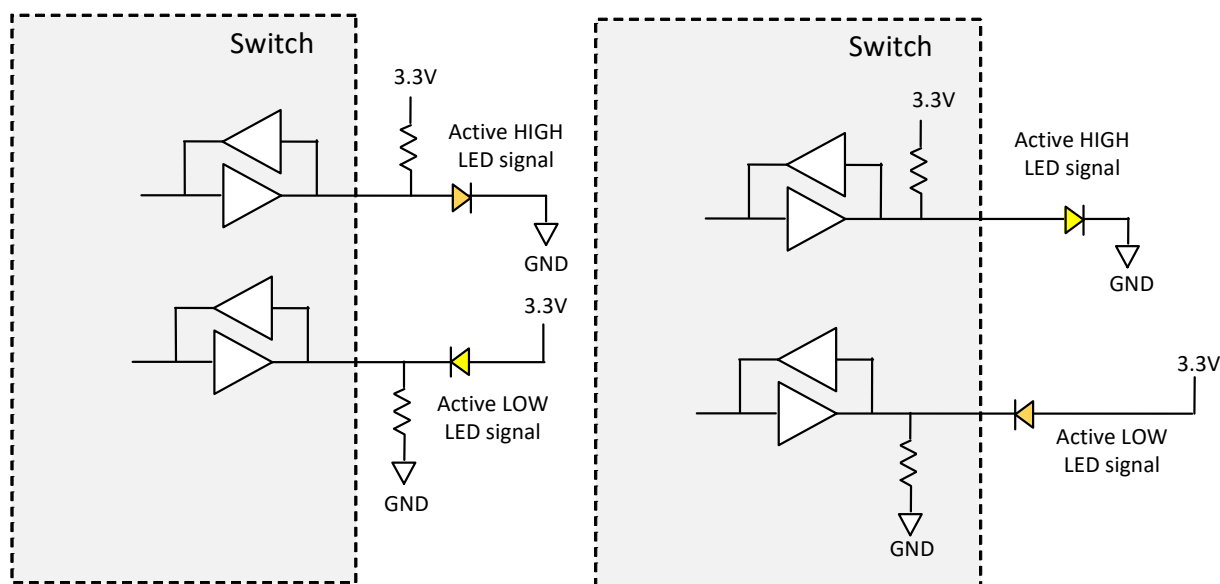
There are 10 LED pins that have secondary functions. These pins serve as input pins during the power-on/reset sequence. The logic level of the pin is sampled at reset and configures the secondary function. After the reset process is completed, the pin acts as an output LED during normal operation. The polarity of the output LED is determined based on the latched input value at reset. For example, if the value at the pin is high during reset, the LED output during normal operation is active-high. The user must first decide, based on the individual application, the values of the input configuration pin shown in [Table 33](#) to provide the correct device configuration. The LED circuit must then be configured to accommodate either an active-low or an active-high LED output.

Table 33: Input Configuration/LED Output Function

LED Output Pins	Input Configuration Pins	Internal Default Pull-ups and Pull-downs	Active State
LED[0]	–	Pull-down	Same as strap pin state
LED[1]	SKIp_RAM_BIST	Pull-up	Same as strap pin state
LED[2]	IMP_VOL_SEL	Pull-down	Same as strap pin state
LED[3]	CLKFREQ[1]	Pull-up	Same as strap pin state
LED[4]	WAN MODE	Pull-down	Same as strap pin state
LED[5]	–	Pull-down	Same as strap pin state
LED[6]	IMP MODE	Pull-down	Same as strap pin state
LED[7]	–	Pull-down	Same as strap pin state
LED[8]	CPU_EEPROM_SEL	Pull-up	Same as strap pin state
LED[9]	HW_FWDG_EN	Pull-up	Same as strap pin state

NOTE: For LEDs whose Active State is same as strap pin state only, if the signal is pulled up/down, the LED is active high/low.

NOTE: Refer to the LED Interface Design Guidelines application note for LED interface design consideration, and the LED function behavior difference between the BCM53134P A0 and B0 chips.

Figure 59: Dual Input Configuration/LED Output Function

Note: When LED signal pins are pulled up or down through an external or an internal termination due to the strap pin configuration, the active states of LED signals are as follows:

- If the signal is pulled up, the LED is active high.
- If the signal is pulled low, the LED is active low.

4.8 Digital Voltage Regulator (LDO)

The BCM53134P LDO generates a 1.8V power supply. The 1.8V is used internally as an intermediate voltage level in 28 -nm technology.

Chapter 5: Hardware Signal Definitions

5.1 I/O Signal Types

The following conventions are used to identify the I/O types. The I/O pin type is useful in referencing the DC pin characteristics.

Table 34: I/O Signal Type Definitions

Abbreviation	Description
XYZ	Active-low signal
3T	3.3V tolerant
A	Analog pin type
B	Bias pin type
CS	Continuously sampled
D	Digital pin type
DNC	Do not connect
GND	Ground
I	Input
I/O	Bidirectional
IPU	Input with internal pull-up
O _{3S}	Tristated signal
O _{DO}	Open-drain output
O	Output
PD	Internal pull-down
SOR	Sample on reset
PWR	Power pin supply
PU	Internal pull-up
XT	Crystal pin type

5.2 Signal Descriptions

Table 35: Signal Descriptions

Signal Name	Type and Default State	Description
PHY Interface		
TDP0_0	Bi	TDP[port#]_ch# , TDN[port#]_ch# are Transmit/Receive Pairs. In TRP/N [port number]_channel number for 1000BASE-T mode, differential data from the media is transmitted and received on all four signal pairs. In auto-negotiation and 10BASE-T and 100BASE-TX modes, the BCM53134P normally transmits on TRP/N[port #]_0 and receives on TRD[port#]_1.
TDN0_0	Bi	
TDP0_1	Bi	
TDN0_1	Bi	
TDP0_2	Bi	
TDN0_2	Bi	
TDP0_3	Bi	
TDN0_3	Bi	
TDP1_0	Bi	
TDN1_0	Bi	
TDP1_1	Bi	
TDN1_1	Bi	
TDP1_2	Bi	
TDN1_2	Bi	
TDP1_3	Bi	
TDN1_3	Bi	
TDP2_0	Bi	
TDN2_0	Bi	
TDP2_1	Bi	
TDN2_1	Bi	
TDP2_2	Bi	
TDN2_2	Bi	
TDP2_3	Bi	
TDN2_3	Bi	
TDP3_0	Bi	
TDN3_0	Bi	
TDP3_1	Bi	
TDN3_1	Bi	
TDP3_2	Bi	
TDN3_2	Bi	
TDP3_3	Bi	
TDN3_3	Bi	

Table 35: Signal Descriptions (Continued)

Signal Name	Type and Default State	Description
RESET/Clock		
RESET_L	I, Pu	Hardware Reset Input. Active low Schmitt-triggered input. Resets the BCM53134P. Active low.
XTAL_P	I	50 MHz Crystal Input/Output. A continuous 50 MHz reference clock must be supplied to the BCM53134P by connecting a 50 MHz crystal between these two pins or by driving XTALI with a clock. When using a crystal, connect a loading capacitor from each pin to GND. NOTE: BCM53134 (A0) supports 50 MHz crystal reference clock input only. NOTE: BCM53134 (B0) supports both 25 MHz and 50 MHz crystal reference clock inputs by using the “CLKREF_SEL” strap pin to select the state. CLKREF_SEL is shared with the LED5 pin. NOTE: In B0 rev, SGMII interface operation is supported only with the 50 MHz crystal reference clock.
XTAL_N	I	
IMP Interface		
IMP_RXCLK	In, Pd	IMP port RGMII Interface Receive Clock 125 MHz for 1000 Mb/s operation, 25 MHz for 100 Mb/s operation and 2.5 MHz for 10 Mb/s operation.
IMP_RXD_0	I, Pd	IMP port RGMII Receive Data Inputs. For 1000 Mb/s operation, data bits RXD[3:0] are clocked-out on the rising edge of RXCLK, and data bits RXD[7:4] are clocked on the falling edge of RXCLK. In 10 Mb/s and 100 Mb/s modes, data bits RXD[3:0] are clocked on the rising edge of RXCLK.
IMP_RXD_1	I, Pd	
IMP_RXD_2	I, Pd	
IMP_RXD_3	I, Pd	
IMP_RXDV	I, Pd	IMP port Receive Data Valid. Active high. Indicates the data on the RXD[3:0] pins are encoded and transmitted. Connects to the TXEN of the external MAC/Management entity.
IMP_TXCLK	O, Pd	IMP Port RGMII Transmit Clock. This clock is driven to synchronize the transmit data in RGMII mode (125 MHz for 1000 Mb/s operation, 25 MHz for 100 Mb/s operation, and 2.5 MHz for 10 Mb/s operation). In RGMII mode, both edges of the clock are used to align with TXD[3:0].
IMP_TXD_0	O, Pd	IMP Port RGMII Transmit Data Output. For 1000 Mb/s operation, data bits TXD[3:0] are clocked on the rising edge of TXCLK, and data bits TXD[7:4] are clocked on the falling edge of TXCLK. For 10 Mb/s and 100 Mb/s, data bits TXD[3:0] are clocked on the rising edge of TXCLK. These output pins have internal 25Ω series termination resistor.
IMP_TXD_1	O, Pd	
IMP_TXD_2	O, Pd	
IMP_TXD_3	O, Pd	
IMP_TXEN	O, Pd	
IMP_VOL_REF	pwr, in	Reference point
WAN Port Interface		
WAN_RXCLK	I, Pd	RGMII Receive Clock 125 MHz for 1000 Mb/s operation, 25 MHz for 100 Mb/s operation and 2.5 MHz for 10 Mb/s operation.
WAN_RXD0	I, Pd	WAN Receive Data Input.
WAN_RXD1	I, Pd	
WAN_RXD2	I, Pd	
WAN_RXD3	I, Pd	
WAN_RXDV	I, Pd	RGMII/MII Receive Data Valid

Table 35: Signal Descriptions (Continued)

Signal Name	Type and Default State	Description
LED15_WAN_TXCLK	O, Pd	When WANLEDSEL = 1b'0: LED15_WAN_TXCLK is used as WAN transmit clock. When WANLEDSEL = 1b'1: LED15_WAN_TXCLK is used as LED15 and the polarity is always active low.
LED10_WAN_TXD0	O, Pd	When WANLEDSEL = 1b'0: LED10_WAN_TXD0 is used as WAN transmit data output 0. When WANLEDSEL = 1b'1: LED10_WAN_TXD0 is used as LED10 and the polarity is always active low.
LED11_WAN_TXD1	O, Pd	When WANLEDSEL = 1b'0: LED11_WAN_TXD1 is used as WAN transmit data output 1. When WANLEDSEL = 1b'1: LED11_WAN_TXD1 is used as LED11 and the polarity is always active low.
LED12_WAN_TXD2	O, Pd	When WANLEDSEL = 1b'0: LED12_WAN_TXD2 is used as WAN transmit data output 2. When WANLEDSEL = 1b'1: LED12_WAN_TXD2 is used as LED12 and the polarity is always active low.
LED13_WAN_TXD3	O, Pd	When WANLEDSEL = 1b'0: LED13_WAN_TXD3 is used as RGMII transmit data output 3. When WANLEDSEL = 1b'1: LED13_WAN_TXD3 is used as LED13 and the polarity is always active low.
LED14_WAN_TXEN	O, Pd	When WANLEDSEL = 1b'0: LED14_WAN_TXEN is used as RGMII transmit enable. When WANLEDSEL = 1b'1: LED14_WAN_TXEN is used as LED14 and the polarity is always active low.
WAN_VOL_REF	I	–
SGMII Interface		
SGMII_RDN0	I	SGMII_Receive Pair
SGMII_RDP0	I	
SGMII_REFCLKN	I	SGMII Reference Clock Differential pair
SGMII_REFCLKP	I	–
SGMII_TDN0	O	SGMII Transmit Pair
SGMII_TDP0	O	
WANVOLSEL	I	Use to set the WAN interface operating voltage level. 0 = 3.3V/2.5V, 1 = 1.8V/1.5V.
Interrupt		
INT_L_LEDMODE1	O, Pu, SOR	Interrupt. This interrupt pin generates an interrupt based on the configuration in the Interrupt Enable register. It can be programmed to generate based on link status change of any port, or to generate an interrupt to a CPU entity when there is a packet(s) queued in the IMP transmit queue. This signal is active low. INT_L is a strap pin for LEDMODE1.
MDC/MDIO Interface		
MDC	Bi, Pd	Management Data I/O. In Master mode, this serial input/output data signal is used to read from and write to the MII registers of the external transceivers. In slave mode, it is used by an external entity to read/write to the switch registers using the Pseudo-PHY. See the MDC/MDIO interface for more information.
MDIO	Bi, Pd	Management Data Clock. In master mode, this 2.5 MHz clock sourced by BCM53134P to the external PHY device. In Slave mode, it is sourced by an external entity.
SPI/EEPROM Interface		

Table 35: Signal Descriptions (Continued)

Signal Name	Type and Default State	Description
SCK/SK	I,O, Pd	SPI Serial Clock. The clock input to the BCM53134P SPI interface is supplied by the SPI master, which supports up to 25 MHz, and is enabled if CPU_EEPROM_SEL is high during power-on reset. EEPROM Serial Clock. The clock output to an external EEPROM device and is enabled if CPU_EEPROM_SEL is low during power-on reset.
SS/CS	I,O, Pu	SPI Slave Select. Active-low signal that enables an SPI interface read or write operation. Enable if CPU_EEPROM_SEL is high during power-on reset. EEPROM Chip Select. Active-high control signal that enables a read operation from an external EEPROM device. Enable if CPU_EEPROM_SEL is low during power-on reset.
MISO/DO_ENEEE	O, Pu, SOR	SPI Master-In/Slave-Out. Output signal which transmits serial data during an SPI interface read operations. Enabled if CPU_EEPROM_SEL is high during Power-On Reset. EEPROM Data Out. Serial data output to an external EEPROM device. Enable if CPU_EEPROM_SEL is low during power-on reset. MISO is used for the strap pin for EN_EEE (Energy Efficient Ethernet). Enable EEE feature for switch MAC: 0 = disable 1 = enable (default)
MOSI/DI	I,O, Pu	SPI Master-Out/Slave-In. Input signal which receives control and address information for the SPI interface, as well as serial data during write operations. Enabled if CPU_EEPROM_SEL is high during power-on reset. EEPROM Data In. Serial data input to an external EEPROM device. Enabled if CPU_EEPROM_SEL is low during power-on reset. NOTE: This signal is tristated during RESET.
RS-232 Interface		
RS232_RXD	I, Pu	RS-232 Input
RS232_TXD_WANLEDSEL	O, Pd, SOR	RS-232 Output. This pin is used for the strap pin for WAN port LED Select 0 = WAN interface, 1 = LED interface. When set to 0, WAN interface is used for WAN interface. When set to 1, WAN interface is used for LED[15:8].
Flash Memory Interface		
FCK	O, Pd	Flash Memory Serial Clock. The clock output for serial Flash memory.
FCS_L_ENLOOPDET	O, Pd, SOR	Flash Memory Chip Select. Active low signal. Chip select to serial Flash memory device. FCS_L is used for the strap pin for ENLOOPDET to enable the Loop Detect feature. 1 = enable.
FSI	I, Pd	Serial Data Input. Serial data input from serial Flash memory.
FSO	O, Pd	Serial Data Output. Serial data output to drive serial Flash memory.
JTAG Interface		
TMS	I	JTAG Mode Select Input.
TRST_L	I	JTAG Test Reset. Active low. Resets the JTAG controller. This signal must be pulled low during normal operation.
TCK	I	JTAG Test Clock Input. Clock Input used to synchronize JTAG control and data transfers. If unused, may be left unconnected.
TDI	I	JTAG Test Data Input. Serial data input to the JTAG TAP Controller. Sampled on the rising edge of TCK. If unused, may be left unconnected.
TDO_EN8051	O, Pu	JTAG Test Data Output. TDO is used for the strap pin for EN_8051 . Enable the embedded BCM8051 microcontroller. The embedded BCM8051 microcontroller is enabled by default. 0 = disabled 1= enabled (default)

Table 35: Signal Descriptions (Continued)

Signal Name	Type and Default State	Description
JTCE	I, Pd	JTAG Capability Select.
LED Interface		
LED0	O, Pd	LED0. NOTE: This LED0 also shares strap pin Xtal_Bypass. The default is internal pull down to select the crystal clock input. <ul style="list-style-type: none"> ■ Xtal_Bypass = 0, selects Crystal mode for crystal input. ■ Xtal_Bypass = 1, selects single-ended CMOS clock input.
LED1	O, Pu	LED1.
LED2_IMPVOLSEL	O, Pd, SOR	LED2 is used for the strap pin for IMPVOLSEL, use to set the IMP interface operating voltage level. 0 = 3.3V/2.5V, 1 = 1.8V/1.5V.
LED3_CLKFREQ1	O, Pu, SOR	LED3 is used for the strap pin for CLK Frequency 1 bit.
LED4_WANMODE	O, Pd, SOR	LED4 is used for the strap pin for WAN Mode. When SGMII is support on Port 8 (IMP Port) and use the Port 5 (WAN port) as RGMII, this pin must be low.
LED5	O, Pd	LED5. NOTE: For BCM53134 (A0), this pin is LED5 only and not a strap pin. NOTE: For BCM53134 (B0), this is strap pin "CLKREF_SEL" and it is shared with LED5 to select the 50 MHz or 25 MHz clock input: <ul style="list-style-type: none"> – CLKREF_SEL = 1, selects 50 MHz crystal reference clock input. – CLKREF_SEL = 0, selects 25 MHz crystal reference clock input. NOTE: For SGMII applications, the reference clock must only be 50 MHz.
LED6_SGMII_P8_SEL	O, Pd, SOR	LED6 is used for the strap pin for SGMII_P8_SEL. NOTE: For BCM53134 (A0), this pin is LED6 only and not a strap pin. To use the IMP port as RGMII, this pin must be low, for example, never have external pull-up for this pin on A0 chip. NOTE: For BCM53134 (B0), this is strap pin "SGMII_P8_SEL" and is shared with LED6 to select the SGMII interface supported on P5 or P8. <ul style="list-style-type: none"> – SGMII_P8_SEL = 1, SGMII is supported on P8. – SGMII_P8_SEL = 0, SGMII is supported on P5 and P8 as RGMII.
LED7	O, Pd	LED7.
LED8_CPUEEPROMSEL	O, Pu, SOR	LED8 is used for the strap pin for CPU_EEPROM_SEL. CPU or EEPROM interface selection. CPU_EEPROM_SEL = 0: Enable EEPROM interface CPU_EEPROM_SEL = 1: Enable SPI Interface (default) The SPI interface must be selected (CPU_EEPROM_SEL=1) for Pseudo-PHY accesses through the MDC/MDIO Interface.
LED9_HWFWDGEN	O, Pu, SOR	LED9 is used for the strap pin for HW_FWDG_EN. Forwarding enable. When this pin is pulled high (default) at power-up, traffic is forwarded without any register settings, based on default register settings. When this pin is pulled low at power-up, frame forwarding is disabled. Traffic forwarding is enabled through Software Register bit set.
LEDCLK_LEDMODE0	O, Pd, SOR	LEDCLK is the LED Shift Clock. This clock is periodically active to enable LEDDATA to shift into external registers. LEDCLK are used for the strap pin for LED MODE[1:0] . Users can select predefined functions to be displayed for each port by setting the bits accordingly.
LEDDATA_CLKFREQ0	O, Pd, SOR	LEDDATA is Serial LED Data Output. Serial LED data for all ports is shifted out when LEDCLK is active. LEDDATA signal is used for the strap pin for Clock Frequency set bit 0. CLKFREQ[1:0] 00 = 125 MHz, 01 = 142 MHz, 10 = 200 MHz, 11 = 167 MHz.

Table 35: Signal Descriptions (Continued)

Signal Name	Type and Default State	Description
LED10	O, Pd	LED10. The polarity is always active low.
LED11	O, Pd	LED11. The polarity is always active low.
LED12	O, Pd	LED12. The polarity is always active low.
LED13	O, Pd	LED13. The polarity is always active low.
LED14	O, Pd	LED14. The polarity is always active low.
LED15	O, Pd	LED15. The polarity is always active low.
Power Interface		
XTAL_AVDD	pwr, in	1.0V for XTAL
AVDDL	pwr, in	1.0V PHY core
AVDDH ^a	pwr, in	3.3V/2.5V for analog I/O
DVDD	pwr, in	1.0V for core
VDDO	pwr, in	3.3V for digital I/O
VDD_1_8	pwr, in	1.8V input for intermediate state for internal use. (The output is from LDO_VOUT.)
PHY_BVDD ^a	pwr, in	3.3V/2.5V
VDD_PLL	pwr, in	1.0V for PLL
VDD_SGMII	pwr, in	1.0V for SGMII
OVDD_IMP	pwr, in	Power for IMP port I/O, 3.3V, 2.5V, 1.8V, or 1.5V
OVDD_WAN	pwr, in	Power for WAN port I/O, 3.3V, 2.5V, 1.8V, or 1.5V
AVSS	–	GND
VSS	–	GND
PHYPLL_GND	–	GND
IMP_VDDP	pwr, in	1.8/1.5V power for IMP port core
WAN_VDDP	pwr, in	1.8/1.5V power for WAN port core
PHY_VDD_PLL	pwr, in	1.0V for PHY PLL
LDO Interface		
LDO_AVDD	pwr, in	3.3V for LDO power input
LDO_VOUT	pwr, out	1.8V output from LDO
LDO_VSENSE	pwr, in	1.8V LDO sense input
Miscellaneous		
DNP	–	No physical ball (no solder mask)
NC	–	No connect
PHY_RDAC	–	6.04 K Ω resistor to GND is required.
ACT_LOOP_DETECT	I, Pd	Active loop detect function.

a. The BCM53134P supports 3.3V or 2.5V GPHY power rail.

Chapter 6: Pin Assignment

6.1 Pin List by Pin Number

Table 36: Pin List by Pin Number

Ball	Ball Name	Type
A01	TDN0_0	Bi-Dir
A02	TDN0_1	Bi-Dir
A03	TDN0_2	Bi-Dir
A04	TDN0_3	Bi-Dir
A05	TDP1_3	Bi-Dir
A06	TDP1_2	Bi-Dir
A07	TDP1_1	Bi-Dir
A08	TDP1_0	Bi-Dir
A09	TDP2_0	Bi-Dir
A10	TDP2_1	Bi-Dir
A11	TDP2_2	Bi-Dir
A12	TDP2_3	Bi-Dir
A13	TDN3_3	Bi-Dir
A14	TDN3_2	Bi-Dir
A15	TDN3_1	Bi-Dir
A16	TDN3_0	Bi-Dir
B01	TDP0_0	Bi-Dir
B02	TDP0_1	Bi-Dir
B03	TDP0_2	Bi-Dir
B04	TDP0_3	Bi-Dir
B05	TDN1_3	Bi-Dir
B06	TDN1_2	Bi-Dir
B07	TDN1_1	Bi-Dir
B08	TDN1_0	Bi-Dir
B09	TDN2_0	Bi-Dir
B10	TDN2_1	Bi-Dir
B11	TDN2_2	Bi-Dir
B12	TDN2_3	Bi-Dir
B13	TDP3_3	Bi-Dir
B14	TDP3_2	Bi-Dir
B15	TDP3_1	Bi-Dir
B16	TDP3_0	Bi-Dir
C01	PHY_RDAC	Input
C02	NC	No connect
C03	PHY_VDD_PLL	Power
C04	VSS	GND

Table 36: Pin List by Pin Number (Continued)

Ball	Ball Name	Type
C05	AVDDH	Power
C06	AVDDL	Power
C07	AVDDH	Power
C08	AVDDL	Power
C09	AVDDH	Power
C10	VSS	GND
C11	AVDDL	Power
C12	AVDDL	Power
C13	PHY_BVDD	Power
C14	AVDDH	Power
C15	VSS	GND
C16	VSS	GND
D01	VSS	GND
D02	VSS	GND
D03	VSS	GND
D04	VSS	GND
D05	DNP	No PHYSICAL BALL (no solder mask)
D06	DNP	No PHYSICAL BALL (no solder mask)
D07	DNP	No PHYSICAL BALL (no solder mask)
D08	DNP	No PHYSICAL BALL (no solder mask)
D09	DNP	No PHYSICAL BALL (no solder mask)
D10	DNP	No PHYSICAL BALL (no solder mask)
D11	DNP	No PHYSICAL BALL (no solder mask)
D12	DNP	No PHYSICAL BALL (no solder mask)
D13	VSS	GND
D14	VSS	GND
D15	MDC	–
D16	MDIO	–
E01	DNP	No PHYSICAL BALL (no solder mask)
E02	LED10_WAN_TXD0	–
E03	LED11_WAN_TXD1	–
E04	DNP	No PHYSICAL BALL (no solder mask)
E05	DNP	No PHYSICAL BALL (no solder mask)
E06	VSS	GND
E07	VSS	GND
E08	VSS	GND
E09	VSS	GND
E10	VSS	GND
E11	VSS	GND
E12	DNP	No PHYSICAL BALL (no solder mask)
E13	DNP	No PHYSICAL BALL (no solder mask)

Table 36: Pin List by Pin Number (Continued)

Ball	Ball Name	Type
E14	LED_0	–
E15	LED_1	–
E16	LED_2_IMPVOSEL	–
F01	LED12_WAN_TXD2	–
F02	LED13_WAN_TXD3	–
F03	LED14_WAN_TXEN	–
F04	DNP	No PHYSICAL BALL (no solder mask)
F05	VSS	GND
F06	VSS	GND
F07	VSS	GND
F08	VSS	GND
F09	VSS	GND
F10	VSS	GND
F11	VSS	GND
F12	VSS	GND
F13	DNP	No PHYSICAL BALL (no solder mask)
F14	LED_3_CLKFREQ1	–
F15	LED_4_WANMODE	–
F16	DNP	No PHYSICAL BALL (no solder mask)
G01	LED15_WAN_TXCLK	–
G02	WAN_VOL_REF	–
G03	VSS	GND
G04	DNP	No PHYSICAL BALL (no solder mask)
G05	VSS	GND
G06	VSS	GND
G07	VSS	GND
G08	VSS	GND
G09	VSS	GND
G10	VSS	GND
G11	VSS	GND
G12	VSS	GND
G13	DNP	No PHYSICAL BALL (no solder mask)
G14	LED_5	–
G15	LED6_SGMII_P8_SEL	–
G16	LED_7	–
H01	WAN_RXCLK	–
H02	WAN_RXD_0	–
H03	OVDD_WAN	–
H04	DNP	No PHYSICAL BALL (no solder mask)
H05	VSS	GND
H06	VSS	GND

Table 36: Pin List by Pin Number (Continued)

Ball	Ball Name	Type
H07	VSS	GND
H08	VSS	GND
H09	VSS	GND
H10	VSS	GND
H11	VSS	GND
H12	VSS	GND
H13	DNP	No PHYSICAL BALL (no solder mask)
H14	LED_DATA_CLKFREQ0	–
H15	LED_8_CPUEEPROMSEL	–
H16	LED_9_HWFWDGEN	–
J01	WAN_RXD1	–
J02	WAN_RXDV	–
J03	OVDD_WAN	Power
J04	DNP	No PHYSICAL BALL (no solder mask)
J05	VSS	GND
J06	VSS	GND
J07	VSS	GND
J08	VSS	GND
J09	VSS	GND
J10	VSS	GND
J11	VSS	GND
J12	VDDO	Power
J13	DNP	No PHYSICAL BALL (no solder mask)
J14	FSI	–
J15	LED_CLK_LEDMODE0	–
J16	DNP	No PHYSICAL BALL (no solder mask)
K01	WAN_RXD2	–
K02	WAN_RXD3	–
K03	WAN_VDDP	–
K04	DNP	No PHYSICAL BALL (no solder mask)
K05	VSS	GND
K06	DVDD	Power
K07	VSS	GND
K08	DVDD	Power
K09	VSS	GND
K10	DVDD	Power
K11	VSS	GND
K12	VDDO	Power
K13	DNP	No PHYSICAL BALL (no solder mask)
K14	FSO	–
K15	FCSL_ENLOOPDET	–

Table 36: Pin List by Pin Number (Continued)

Ball	Ball Name	Type
K16	FCK	–
L01	SGMII_TDN0	–
L02	VSS	GND
L03	VSS	GND
L04	DNP	No PHYSICAL BALL (no solder mask)
L05	VSS	GND
L06	DVDD	Power
L07	VSS	GND
L08	DVDD	Power
L09	VSS	GND
L10	DVDD	Power
L11	VSS	GND
L12	VDD_1P8	Power
L13	DNP	No PHYSICAL BALL (no solder mask)
L14	VDDO	Power
L15	NC	No connect
L16	ACT_LOOP_DETECT	–
M01	SGMII_TDP0	–
M02	SGMII_RDP0	–
M03	VDD_SGMII	Power
M04	DNP	No PHYSICAL BALL (no solder mask)
M05	DNP	No PHYSICAL BALL (no solder mask)
M06	DVDD	Power
M07	VSS	GND
M08	DVDD	Power
M09	VSS	GND
M10	DVDD	Power
M11	VSS	GND
M12	DNP	No PHYSICAL BALL (no solder mask)
M13	DNP	No PHYSICAL BALL (no solder mask)
M14	WANVOLSEL	Input
M15	NC	No connect
M16	DNP	No PHYSICAL BALL (no solder mask)
N01	DNP	No PHYSICAL BALL (no solder mask)
N02	SGMII_RDN0	–
N03	VDD_SGMII	Power
N04	LDO_VOUT	Power
N05	DNP	No PHYSICAL BALL (no solder mask)
N06	DNP	No PHYSICAL BALL (no solder mask)
N07	DNP	No PHYSICAL BALL (no solder mask)
N08	DNP	No PHYSICAL BALL (no solder mask)

Table 36: Pin List by Pin Number (Continued)

Ball	Ball Name	Type
N09	DNP	No PHYSICAL BALL (no solder mask)
N10	DNP	No PHYSICAL BALL (no solder mask)
N11	DNP	No PHYSICAL BALL (no solder mask)
N12	DNP	No PHYSICAL BALL (no solder mask)
N13	VDD_1P8	Power
N14	JTCE	Input
N15	TDO_EN8051	Output
N16	TCK	Input
P01	SGMII_REFCLKP	–
P02	SGMII_REFCLKN	–
P03	LDO_VSENSE	Power
P04	LDO_AVDD	Power
P05	NC	No connect
P06	VDD_PLL	Power
P07	VSS	GND
P08	OVDD_IMP	Power
P09	OVDD_IMP	Power
P10	IMP_VDDP	Power
P11	VSS	GND
P12	SS	Input
P13	SCK	Input
P14	NC	No connect
P15	TRST_L	Input
P16	TDI	Input
R01	SGMII_TESTN	Output
R02	SGMII_TESTP	Output
R03	VSS	GND
R04	XTAL_VDD	Power
R05	IMP_TXEN	–
R06	IMP_TXD1	–
R07	IMP_TXD2	–
R08	IMP_VOL_REF	Power/GND
R09	IMP_RXD3	–
R10	IMP_RXD1	–
R11	IMP_RXDV	–
R12	MOSI_DI	–
R13	INTR_L_LEDMODE1	–
R14	RESET_L	–
R15	TMS	Input
R16	DNP	No PHYSICAL BALL (no solder mask)
T01	VSS	GND

Table 36: Pin List by Pin Number (Continued)

Ball	Ball Name	Type
T02	VSS	GND
T03	XTAL_N	–
T04	XTAL_P	–
T05	DNP	No PHYSICAL BALL (no solder mask)
T06	IMP_TXD0	–
T07	IMP_TXD3	–
T08	IMP_TXCLK	–
T09	IMP_RXD0	–
T10	IMP_RXD2	–
T11	IMP_RXCLK	–
T12	MISO_DO_ENEEEE	–
T13	DNP	No PHYSICAL BALL (no solder mask)
T14	RS232_TXD_WANLEDSEL	–
T15	RS232_RXD	–
T16	VSS	GND

6.2 Pin List by Pin Name

Table 37: Pin List by Pin Name

Ball	Ball Name	Type
C05	AVDDH	Power
C07	AVDDH	Power
C09	AVDDH	Power
C14	AVDDH	Power
C06	AVDDL	Power
C08	AVDDL	Power
C11	AVDDL	Power
C12	AVDDL	Power
D05	DNP	No PHYSICAL BALL (no solder mask)
D06	DNP	No PHYSICAL BALL (no solder mask)
D07	DNP	No PHYSICAL BALL (no solder mask)
D08	DNP	No PHYSICAL BALL (no solder mask)
D09	DNP	No PHYSICAL BALL (no solder mask)
D10	DNP	No PHYSICAL BALL (no solder mask)
D11	DNP	No PHYSICAL BALL (no solder mask)
D12	DNP	No PHYSICAL BALL (no solder mask)
E01	DNP	No PHYSICAL BALL (no solder mask)
E04	DNP	No PHYSICAL BALL (no solder mask)
E05	DNP	No PHYSICAL BALL (no solder mask)
E12	DNP	No PHYSICAL BALL (no solder mask)
E13	DNP	No PHYSICAL BALL (no solder mask)
F04	DNP	No PHYSICAL BALL (no solder mask)
F13	DNP	No PHYSICAL BALL (no solder mask)
F16	DNP	No PHYSICAL BALL (no solder mask)
G04	DNP	No PHYSICAL BALL (no solder mask)
G13	DNP	No PHYSICAL BALL (no solder mask)
H04	DNP	No PHYSICAL BALL (no solder mask)
H13	DNP	No PHYSICAL BALL (no solder mask)
J04	DNP	No PHYSICAL BALL (no solder mask)
J13	DNP	No PHYSICAL BALL (no solder mask)
J16	DNP	No PHYSICAL BALL (no solder mask)
K04	DNP	No PHYSICAL BALL (no solder mask)
K13	DNP	No PHYSICAL BALL (no solder mask)
L04	DNP	No PHYSICAL BALL (no solder mask)
L13	DNP	No PHYSICAL BALL (no solder mask)
M04	DNP	No PHYSICAL BALL (no solder mask)
M05	DNP	No PHYSICAL BALL (no solder mask)
M12	DNP	No PHYSICAL BALL (no solder mask)
M13	DNP	No PHYSICAL BALL (no solder mask)

Table 37: Pin List by Pin Name (Continued)

Ball	Ball Name	Type
M16	DNP	No PHYSICAL BALL (no solder mask)
N01	DNP	No PHYSICAL BALL (no solder mask)
N05	DNP	No PHYSICAL BALL (no solder mask)
N06	DNP	No PHYSICAL BALL (no solder mask)
N07	DNP	No PHYSICAL BALL (no solder mask)
N08	DNP	No PHYSICAL BALL (no solder mask)
N09	DNP	No PHYSICAL BALL (no solder mask)
N10	DNP	No PHYSICAL BALL (no solder mask)
N11	DNP	No PHYSICAL BALL (no solder mask)
N12	DNP	No PHYSICAL BALL (no solder mask)
R16	DNP	No PHYSICAL BALL (no solder mask)
T05	DNP	No PHYSICAL BALL (no solder mask)
T13	DNP	No PHYSICAL BALL (no solder mask)
K06	DVDD	Power
K08	DVDD	Power
K10	DVDD	Power
L06	DVDD	Power
L08	DVDD	Power
L10	DVDD	Power
M06	DVDD	Power
M08	DVDD	Power
M10	DVDD	Power
K16	FCK	—
K15	FCSL_ENLOOPDET	—
J14	FSI	—
K14	FSO	—
T11	IMP_RXCLK	—
T09	IMP_RXD0	—
R10	IMP_RXD1	—
T10	IMP_RXD2	—
R09	IMP_RXD3	—
R11	IMP_RXDV	—
T08	IMP_TXCLK	—
T06	IMP_TXD0	—
R06	IMP_TXD1	—
R07	IMP_TXD2	—
T07	IMP_TXD3	—
R05	IMP_TXEN	—
P10	IMP_VDDP	Power
R08	IMP_VOL_REF	Power/GND
R13	INTR_L_LEDMODE1	—

Table 37: Pin List by Pin Name (Continued)

Ball	Ball Name	Type
N14	JTCE	Input
P04	LDO_AVDD	Power
N04	LDO_VOUT	Power
P03	LDO_VSENSE	Power
E14	LED_0	–
E15	LED_1	–
E16	LED_2_IMPVSSEL	–
F14	LED_3_CLKFREQ1	–
F15	LED_4_WANMODE	–
G14	LED_5	–
G15	LED6_SGMII_P8_SEL	–
G16	LED_7	–
H15	LED_8_CPUEEPROMSEL	–
H16	LED_9_HWFWDGEN	–
J15	LED_CLK_LEDMODE0	–
H14	LED_DATA_CLKFREQ0	–
E02	LED10_WAN_TXD0	–
E03	LED11_WAN_TXD1	–
F01	LED12_WAN_TXD2	–
F02	LED13_WAN_TXD3	–
F03	LED14_WAN_TXEN	–
G01	LED15_WAN_TXCLK	–
D15	MDC	–
D16	MDIO	–
T12	MISO_DO_ENEEE	–
R12	MOSI_DI	–
G02	WAN_VOL_REF	–
P05	NC	No connect
P14	NC	No connect
R01	SGMII_TESTN	–
R02	SGMII_TESTP	–
C02	NC	No connect
L15	NC	No connect
L16	ACT_LOOP_DETECT	–
M15	NC	No connect
P08	OVDD_IMP	Power
P09	OVDD_IMP	Power
H03	OVDD_WAN	Power
J03	OVDD_WAN	Power
C13	PHY_BVDD	Power
C01	PHY_RDAC	Input

Table 37: Pin List by Pin Name (Continued)

Ball	Ball Name	Type
C03	PHY_VDD_PLL	Power
R14	RESET_L	–
T15	RS232_RXD	–
T14	RS232_TXD_WANLEDSEL	–
P13	SCK	Input
N02	SGMII_RDN0	–
M02	SGMII_RDP0	–
P02	SGMII_REFCLKN	–
P01	SGMII_REFCLKP	–
L01	SGMII_TDN0	–
M01	SGMII_TDP0	–
P12	SS	Input
N16	TCK	Input
P16	TDI	Input
A01	TDN0_0	Bi-Dir
A02	TDN0_1	Bi-Dir
A03	TDN0_2	Bi-Dir
A04	TDN0_3	Bi-Dir
B08	TDN1_0	Bi-Dir
B07	TDN1_1	Bi-Dir
B06	TDN1_2	Bi-Dir
B05	TDN1_3	Bi-Dir
B09	TDN2_0	Bi-Dir
B10	TDN2_1	Bi-Dir
B11	TDN2_2	Bi-Dir
B12	TDN2_3	Bi-Dir
A16	TDN3_0	Bi-Dir
A15	TDN3_1	Bi-Dir
A14	TDN3_2	Bi-Dir
A13	TDN3_3	Bi-Dir
N15	TDO_EN8051	Output
B01	TDP0_0	Bi-Dir
B02	TDP0_1	Bi-Dir
B03	TDP0_2	Bi-Dir
B04	TDP0_3	Bi-Dir
A08	TDP1_0	Bi-Dir
A07	TDP1_1	Bi-Dir
A06	TDP1_2	Bi-Dir
A05	TDP1_3	Bi-Dir
A09	TDP2_0	Bi-Dir
A10	TDP2_1	Bi-Dir

Table 37: Pin List by Pin Name (Continued)

Ball	Ball Name	Type
A11	TDP2_2	Bi-Dir
A12	TDP2_3	Bi-Dir
B16	TDP3_0	Bi-Dir
B15	TDP3_1	Bi-Dir
B14	TDP3_2	Bi-Dir
B13	TDP3_3	Bi-Dir
R15	TMS	Input
P15	TRST_L	Input
L12	VDD_1P8	Power
N13	VDD_1P8	Power
P06	VDD_PLL	Power
M03	VDD_SGMII	Power
N03	VDD_SGMII	Power
J12	VDDO	Power
K12	VDDO	Power
L14	VDDO	Power
C04	VSS	GND
C10	VSS	GND
C15	VSS	GND
C16	VSS	GND
D01	VSS	GND
D02	VSS	GND
D03	VSS	GND
D04	VSS	GND
D13	VSS	GND
D14	VSS	GND
E06	VSS	GND
E07	VSS	GND
E08	VSS	GND
E09	VSS	GND
E10	VSS	GND
E11	VSS	GND
F05	VSS	GND
F06	VSS	GND
F07	VSS	GND
F08	VSS	GND
F09	VSS	GND
F10	VSS	GND
F11	VSS	GND
F12	VSS	GND
G03	VSS	GND

Table 37: Pin List by Pin Name (Continued)

Ball	Ball Name	Type
G05	VSS	GND
G06	VSS	GND
G07	VSS	GND
G08	VSS	GND
G09	VSS	GND
G10	VSS	GND
G11	VSS	GND
G12	VSS	GND
H05	VSS	GND
H06	VSS	GND
H07	VSS	GND
H08	VSS	GND
H09	VSS	GND
H10	VSS	GND
H11	VSS	GND
H12	VSS	GND
J05	VSS	GND
J06	VSS	GND
J07	VSS	GND
J08	VSS	GND
J09	VSS	GND
J10	VSS	GND
J11	VSS	GND
K05	VSS	GND
K07	VSS	GND
K09	VSS	GND
K11	VSS	GND
L02	VSS	GND
L03	VSS	GND
L05	VSS	GND
L07	VSS	GND
L09	VSS	GND
L11	VSS	GND
M07	VSS	GND
M09	VSS	GND
M11	VSS	GND
P07	VSS	GND
P11	VSS	GND
R03	VSS	GND
T01	VSS	GND
T02	VSS	GND

Table 37: Pin List by Pin Name (Continued)

Ball	Ball Name	Type
T16	VSS	GND
K03	WAN_VDDP	Power
H01	WAN_RXCLK	–
H02	WAN_RXD_0	–
J01	WAN_RXD1	–
J02	WAN_RXDV	–
K01	WAN_RXD2	–
K02	WAN_RXD3	–
M14	WANVOLSEL	Input
T03	XTAL_N	–
T04	XTAL_P	–
R04	XTAL_VDD	Power

Chapter 7: Electrical Characteristics

7.1 Absolute Maximum Ratings

Table 38: Absolute Maximum Ratings

Symbol	Parameter and Pins	Minimum	Maximum	Unit
AVDDL, DVDD, EPHY_PLLVDD, QGPHY1_PLLVDD, QGPHY2_PLLVDD	Supply voltage	GND – 0.3	1.1	V
OVDD, OTP_VDD, EPHY_BVDD, QGPHY1_BVDD, QGPHY2_BVDD, SWREG_VDDO, XTAL_AVDD	Supply voltage	GND – 0.3	3.63	V
I_I	Input current	–	–	mA
T_{STG}	Storage temperature	–40	125	°C
V_{ESD}	Electrostatic discharge	–	1800V	V
–	Input voltage: Digital input pins	–	–	V

NOTE: These specifications indicate levels where permanent damage to the device may occur. Functional operation is not guaranteed under these conditions. Operation at absolute maximum conditions for extended periods may adversely affect long-term reliability of the device.

7.2 Recommended Operating Conditions

Table 39: Recommended Operating Conditions

Symbol	Parameter	Pins	Minimum	Maximum	Unit
VDD	Supply voltage	–	0.95	1.1	V
		–	3.14	3.47	V
		–	–	–	V
		–	–	–	V
V_{IH}	High-level input voltage	All digital inputs	1.7	–	V
V_{IL}	Low-level input voltage for 2.5V	All digital inputs	–	0.7	V
V_{IL}	Low-level input voltage for 3.3V	All digital inputs	–	0.9	V
$-T_A$	Ambient operating temperature	–	0	70	°C

NOTE: The recommended minimum/maximum operating voltages are not final. The final numbers will be updated after the characterization.

7.3 Electrical Characteristics

Table 40: Electrical Characteristics

Symbol	Parameter	Pins	Conditions	Min.	Typical	Max.	Unit
I _{DD}	Supply current (for 1x RGMII and 1x SGMII operation)	1.0V power rail (analog)	Measured	–	183.3	–	mA
		1.0V power rail (digital)	Measured	–	272.5	–	mA
		3.3V power rail	Measured	–	197.5	–	mA
		OVDD_IMP and OVDD_WAN (2.5V for RGMII)	Measured	–	33.6	–	mA
V _{OH}	High-level output voltage	Digital output pins at 3.3V	I _{OH} = –8 mA	2.4	–	–	V
		Digital output pins at 2.5V	I _{OH} = –8 mA	2.0	–	–	V
		Digital output pins at 1.5V/1.8V	I _{OH} = –8 mA	1.0	–	–	V
V _{OL}	Low-level output voltage	Digital output pins at 3.3V	I _{OL} = 8 mA	–	–	0.4	V
		Digital output pins at 2.5V	I _{OL} = 8 mA	–	–	0.4	V
		Digital output pins at 1.5V/1.8V	I _{OL} = 8 mA	–	–	0.4	V
V _{IH}	High-level input voltage	Digital input pins at 3.3V and 2.5V	–	1.7	–	–	V
		Digital input pins at 1.5V/1.8V	–	0.9	–	–	V
V _{IL}	Low-level input voltage	Digital input pins for 3.3V	–	–	–	0.9	V
		Digital input pins for 2.5V	–	–	–	0.7	V
		Digital input pins at 1.5V/1.8V	–	–	–	0.5	V
I _I	Input current	Digital Inputs w/pull-up resistors	–	–	–	–	μA
		Digital Inputs w/pull-up resistors	–	–	–	–	μA
		Digital Inputs w/pull-down resistors	–	–	–	–	μA
		Digital Inputs w/pull-down resistors	–	–	–	–	μA
		All other digital inputs	–	–	–	–	μA
1. The measured result is based on 3.3V GPHY and 120m cable length. The 3m result will be 50 mW higher.							
2. The measured result of 2.5V GPHY power rail will be 80 mW lower than 3.3V GPHY.							

Table 41: Internal Voltage Regulator Electrical Characteristics

Parameter	Min.	Typ.	Max.	Unit
1.8V LDO output range	1.44	1.8	1.98	V
1.8V LDO output ripple	–	15	30	mV
1.8V LDO output accuracy	–	–	4.5	%
1.8V LDO output current	–	–	100	mA
1.8V LDO output current limit	200	300	400	mA
1.8V LDO power-up time	–	120	160	μ

Table 42: SGMII/SerDes DC Characteristics

Parameter	Symbol	Min.	Typ.	Max.	Unit
Receiver					
Input Voltage (Differential pk-pk), AC-coupled	V_{ID}	100	–	2000	mVp-p
Input Impedance (Differential), integrated on-chip	R_{IN}	–	100	–	Ω
Transmitter					
Output Voltage (Differential pk-pk), programmable	V_{OD}	–	1000	–	mVp-p
Output Impedance (Differential)	R_{OUT}	–	100	–	Ω


Table 43: 2.5 GbE SerDes DC Characteristics

Parameter	Symbol	Min.	Typ.	Max.	Unit
Differential Input Voltage	V_{IN}	–	–	1600	mVppd
Differential Input Impedance	VR_{IN}	–	100	–	Ω
Differential Output Voltage	V_{OD}	–	1000	–	mVppd
Differential Output Impedance	VR_{OUT}	–	100	–	Ω

Chapter 8: Timing Characteristics

8.1 Reset and Clock Timing

Figure 60: Reset and Clock Timing



NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 44: Reset and Clock Timing


Description	Parameter	Minimum	Typical	Maximum
XTALI period	t101	19.999 ns	20 ns	20.001 ns
XTALI high time	t102	9 ns	–	11 ns
XTALI low time	t103	9 ns	–	11 ns
RESET low pulse duration	t104	80 ms	100 ms	–
RESET rise time	t105	–	–	25 ns
Configuration valid setup to RESET rising	t107	100 ns	–	–
Configuration valid hold from RESET rising	t108	–	–	100 ns
Hardware initialization is complete. All the strap pin values are clocked in, and the internal registers can be accessed.	t109	5 ms before the registers can be accessed		

8.2 RGMII Interface Timing

The following specifies timing information regarding the IMP interface pins when configured in RGMII mode.

8.2.1 RGMII Output Timing (Normal Mode)

Figure 61: RGMII Output Timing (Normal Mode)



NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.


Table 45: RGMII Output Timing (Normal Mode)

Description	Parameter	Minimum	Typical	Maximum	Unit
MII1_TXC clock period (1000M mode)	—	7.2	8	8.8	ns
MII1_TXC clock period (100M mode)	—	36	40	44	ns
MII1_TXC clock period (10M mode)	—	360	400	440	ns
TskewT: Data to clock output skew	t201	–500 (1000M)	0	+500 (1000M)	ps
TskewT: Data to Clock at 1.5V/1.8V mode	t201	–750 (1000M)	0	+500 (1000M)	ps
Duty cycle for 1000M (GbE)	—	45	50	55	%
Duty cycle for 10/100M (FE)	—	40	50	60	%

NOTE: The output timing in 10/100M operation is always as specified in the delayed mode.

8.2.2 RGMII Output Timing (Delayed Mode)

Figure 62: RGMII Output Timing (Delayed Mode)




NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 46: RGMII Output Timing (Delayed Mode)

Description	Parameter	Minimum	Typical	Maximum	Unit
MII1_TXC clock period (1000M mode)	–	7.2	8	8.8	ns
MII1_TXC clock period (100M mode)	–	36	40	44	ns
MII1_TXC clock period (10M mode)	–	360	400	440	ns
TsetupT Data valid to clock transition: Available setup time at the output source (delayed mode)	t201D	1.2 (all speeds)	2.0	–	ns
TholdT Clock transition to data valid: Available hold time at the output source (delayed mode)	t202D	1.2 (all speeds)	2.0	–	ns
TsetupT Data valid to clock transition: Available setup time at the output source (1.5V/1.8V mode)	t201D	1.0 (all speeds)	2.0	–	ns
TholdT Clock transition to data valid: Available hold time at the output source (1.5V/1.8V mode)	t202D	1.0 (all speeds)	2.0	–	ns
Duty cycle for 1000M (GbE)	–	45	50	55	%
Duty cycle for 10/100M (FE)	–	40	50	60	%

8.2.3 RGMII Input Timing (Normal Mode)

Figure 63: RGMII Input Timing (Normal Mode)




NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 47: RGMII Input Timing (Normal Mode)

Description	Parameter	Minimum	Typical	Maximum	Unit
MII1_RXC clock period (1000M mode)	–	7.2	8	8.8	ns
MII1_RXC clock period (100M mode)	–	36	40	44	ns
MII1_RXC clock period (10M mode)	–	360	400	440	ns
TsetupR Input setup time: Valid data to clock	t301	1.0	2.0	–	ns
TholdR Input hold time: Clock to valid data	t302	1.0	2.0	–	ns
Duty cycle for 1000M (GbE)	–	45	50	55	%
Duty cycle for 10/100M (FE)	–	40	50	60	%

8.2.4 RGMII Input Timing (Delayed Mode)

Figure 64: RGMII Input Timing (Delayed Mode)



NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.


Table 48: RGMII Input Timing (Delayed Mode)

Description	Parameter	Minimum	Typical	Maximum	Unit
TsetupR	t301D	–1.0 (1000M)	–	–	ns
Input setup time (delayed mode)		–1.0 (10/100M)	–	–	ns
TholdR	t302D	3.0 (1000M)	–	–	ns
Input hold time (delayed mode)		9.0 (10/100M)	–	–	ns

8.3 MDC/MDIO Timing

The following specifies timing information regarding the MDC/MDIO interface pins.

Figure 65: MDC/MDIO Timing (Slave Mode)



NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 49: MDC/MDIO Timing (Slave Mode)

Description	Parameter	Minimum	Typical	Maximum	Unit
MDC cycle time	t401	80	–	–	ns
MDC high/low	–	30	–	–	ns
MDC rise/fall time	t402	–	–	10	ns
MDIO input setup time to MDC rising	t403	7.5	–	–	ns
MDIO input hold time from MDC rising	t404	7.5	–	–	ns
MDIO output delay from MDC rising	t405	0	–	45	ns

Figure 66: MDC/MDIO Timing (Master Mode)





Table 50: MDC/MDIO Timing (Master Mode)

Description	Parameter	Minimum	Typical	Maximum	Unit
MDC cycle time	t401	400	–	–	ns
MDC high/low	–	160	–	240	ns
MDC rise/fall time	t402	–	–	10	ns
MDIO input setup time to MDC rising	t403	20	–	–	ns
MDIO input hold time from MDC rising	t404	0	–	–	ns
MDIO output delay from MDC falling	t405m	–5	–	20	ns

8.4 Serial LED Interface Timing

The following specifies timing information regarding the LED interface pins.

Figure 67: Serial LED Interface Timing




NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 51: Serial LED Interface Timing

Description	Parameter	Minimum	Typical	Maximum	Unit
LED update cycle period	t301	–	42	–	ms
LEDCLK period	t302	–	320	–	ns
LEDCLK high-pulse width	t303	150	–	170	ns
LEDCLK low-pulse width	t304	150	–	170	ns
LEDCLK to LEDDATA output time	t305	140	–	180	ns

8.5 SPI Timings

Figure 68: SPI Timings, SS Asserted During SCK High



NOTE: $\overline{\text{SS}}$ should be asserted only while SCK is high.

NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 52: SPI Timings

Description	Parameter	Minimum	Typical	Maximum
SCK clock period	t601	40 ns	500 ns	—
SCK high/low time	t602	20 ns	250 ns	—
MOSI to SCK setup time	t603	5 ns	—	—
MOSI to SCK hold time	t604	5 ns	—	—
SCK to MISO valid	t605	—	—	10 ns
SCK to MISO valid	t605	—	—	10 ns
Time interval from assert of \overline{SS} to start of SCK	t607	10 ns	—	—
Time interval from end of SCK to deassert of \overline{SS}	t608	20 ns	—	—

8.6 JTAG Interface

JTAG timing is synchronous to the JTAG_TCK clock.

Figure 69: JTAG Interface

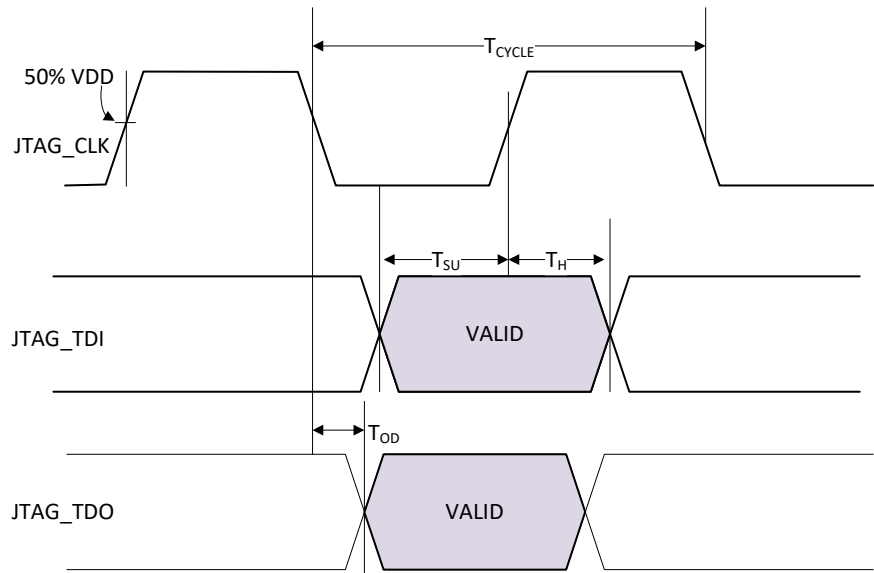



Table 53: JTAG Interface

Parameter	Description	Min.	Typ.	Max.	Unit
T_{CYCLE}	JTAG Cycle Time	50	–	–	ns
T_{SU}	Input Setup Time	12.5	–	–	ns
T_{TH}	Input Hold Time	12.5	–	–	ns
T_{OD}	Output Delay Time Measured from Falling Edge of JTAG_TCK	–	–	22	ns

8.7 EEPROM Timing

Figure 70: EEPROM Timing



NOTE: Advanced timing extracted from statistical data. It may or may not reflect the true timing of the device.

Table 54: EEPROM Timing

Description	Parameter	Minimum	Typical	Maximum
SCK clock frequency	t701	–	100 kHz	–
SCK high/low time	t702	–	5 μ s	–
SCK low to CS, DI valid	t703	–	–	500 ns
SCK low to CS, DI invalid	t704	500 ns	–	–
DO to SCK falling setup time	t705	200 ns	–	–
DO to SCK falling hold time	t706	200 ns	–	–

8.8 SGMII/SerDes Timing

Figure 71: SGMII/SerDes Interface Output Timing




Table 55: SGMII/SerDes Interface Output Timing

Parameter	Description	Min.	Typ.	Max.	Unit
t801	Transmit Data Signaling Speed	–	1.25	–	Gbaud
t802	Transmit Data Rise Time (20% to 80%)	60	–	320	ps
t803	Transmit Data Fall Time (20% to 80%)	60	–	320	ps
t805	Transmit Data Total Jitter	–	–	0.25	UI

Figure 72: SGMII/SerDes Interface Input Timing




Table 56: SGMII/SerDes Interface Input Timing

Parameter	Description	Min.	Typ.	Max.	Unit
t806	Receive Data Signaling Speed	–	1.25	–	Gbaud
v809	Receive Data Differential Input (pk-pk)	0.1	–	2.0	V

8.9 2.5GbE/SerDes Timing

Figure 73: 2.5GbE/SerDes Interface Output Timing




Table 57: 2.5GbE/SerDes Interface Output Timing

Parameter	Description	Min.	Typ.	Max.	Unit
t801	Transmit Data Signaling Speed	–	3.125	–	Gbaud
t802	Transmit Data Rise Time (20% to 80%)	30	–	130	ps
t803	Transmit Data Fall Time (20% to 80%)	30	–	130	ps
t805	Transmit Data Total Jitter	–	–	0.35	UI
T _{SKEW}	Transmit Differential Skew	–	–	0.15	ps

Figure 74: 2.5GbE/SerDes Interface Input Timing




Table 58: 2.5GbE/SerDes Interface Input Timing

Parameter	Description	Min.	Typ.	Max.	Unit
t806	Receive Data Signaling Speed	–	1.25	–	Gbaud
v809	Receive Data Differential Input (pk-pk)	–	–	1.6	V

Chapter 9: Thermal Characteristics

NOTE: The maximum allowed junction temperature is 110°C.

9.1 Package Only

Table 59: Package Only, 2s2p PCB, $T_A = 70^\circ\text{C}$, $P = 1.723\text{W}$

Device power dissipation, P (W)		1.72				
Ambient air temperature T _A (°C)		70				
θ _{JA} in still air (°C/W)		38.17				
θ _{JB} (°C/W)		17.32				
θ _{JC} (°C/W)		23.74				
2s2p board, pkg only						
Package Thermal Performance Curve						
Air Velocity		T _J	T _T	θ _{JA}	Ψ _{JT}	Ψ _{JB}
m/s	ft/min	(°C)	(°C)	(°C/W)	(°C/W)	(°C/W)
0	0	135.8	133.5	38.17	1.30	25.74
0.508	100	132.4	130.1	36.19	1.32	25.77
1.016	200	130.7	128.3	35.21	1.38	25.74
2.032	400	128.8	126.1	34.11	1.55	25.62
3.048	600	127.6	124.6	33.43	1.71	25.50

9.2 Package Only with Heat Sink (50 x 50 x 35 mm³)

Table 60: Package with External Heat Sink 50 x 50 x 35 mm³, 2s2p PCB, $T_A = 70^\circ\text{C}$, $P = 1.723\text{W}$

Device power dissipation, P (W)				1.72		
Ambient air temperature T _A (°C)				70		
θ _{JA} in still air (°C/W)				23.75		
θ _{JB} (°C/W)				17.32		
θ _{JC} (°C/W)				23.74		
2s2p board, 50 x 50 x 35 mm ³ estHS						
Package Thermal Performance Curve						
Air Velocity		T _J	T _T	θ _{JA}	Ψ _{JT}	Ψ _{JB}
m/s	ft/min	(°C)	(°C)	(°C/W)	(°C/W)	(°C/W)
0	0	110.9	82.0	23.75	16.78	17.47
0.508	100	105.9	76.4	20.82	17.10	17.19
1.016	200	104.8	75.3	20.20	17.15	17.15
2.032	400	104.2	74.7	19.87	17.15	17.14
3.048	600	104.2	74.5	19.73	17.14	17.14

9.3 Package Only

Table 61: Package Only, 2s2p PCB, $T_A = 55^\circ\text{C}$, $P = 1.723\text{W}$

Device power dissipation, P (W)		1.72				
Ambient air temperature T _A (°C)		55				
θ _{JA} in still air (°C/W)		38.65				
θ _{JB} (°C/W)		17.32				
θ _{JC} (°C/W)		23.74				
2s2p board, pkg only						
Package Thermal Performance Curve						
Air Velocity		T _J	T _T	θ _{JA}	Ψ _{JT}	Ψ _{JB}
m/s	ft/min	(°C)	(°C)	(°C/W)	(°C/W)	(°C/W)
0	0	121.6	119.4	38.65	1.29	25.74
0.508	100	117.8	115.5	36.43	1.31	25.77
1.016	200	116.0	113.6	35.38	1.36	25.75
2.032	400	114.0	111.3	34.23	1.54	25.63
3.048	600	112.8	109.8	33.52	1.70	25.51


9.4 Package Only with Heat Sink (19 x 19 x 5 mm³)

Table 62: Package with External Heat Sink 19.x 19 x 5 mm³, 2s2p PCB, $T_A = 55^\circ\text{C}$, $P = 1.723\text{W}$

Device power dissipation, P (W)		1.72				
Ambient air temperature T _A (°C)		55				
θ _{JA} in still air (°C/W)		31.70				
θ _{JB} (°C/W)		17.32				
θ _{JC} (°C/W)		23.74				
2s2p board, pkg only						
Package Thermal Performance Curve						
Air Velocity		T _J	T _T	θ _{JA}	Ψ _{JT}	Ψ _{JB}
m/s	ft/min	(°C)	(°C)	(°C/W)	(°C/W)	(°C/W)
0	0	109.6	84.5	31.70	14.58	19.66
0.508	100	105.3	79.8	29.20	14.82	19.46
1.016	200	102.1	76.0	27.31	15.13	19.17
2.032	400	98.7	72.0	25.34	15.49	18.80
3.048	600	96.9	69.9	24.33	15.68	18.59

Chapter 10: Mechanical Information

Figure 75: BCM53134P Mechanical Information



Chapter 11: Ordering Information

Table 63: Ordering Information

Part Number	Package	Ambient Temperature
BCM53134PKFBG	212-pin FBGA package	0°C to 70°C
BCM53134PIFBG	212-pin FBGA package	–45°C to 85°C

Revision History

53134P-DS105; November 15, 2018

Updated:

- [RGMII Interface](#)
- [Table 35, Signal Descriptions](#)
- [Table 40, Electrical Characteristics](#)
- [Table 45, RGMII Output Timing \(Normal Mode\)](#)
- [Table 46, RGMII Output Timing \(Delayed Mode\)](#)

53134P-DS104; July 23, 2018

Updated:

- Table 35, Signal Descriptions
- Table 36, Pin List by Pin Number 53134O/53134P/53134U
- Table 38, Pin List by Pin Name 53134O/53134P/53134U

53134P-DS103; March 15, 2018

Updated:

- Table 42, Electrical Characteristics

53134P-DS102; August 24, 2016

Updated:

- Figure 70: "MDC/MDIO Timing (Master Mode)," on page 188
- Table 52: "MDC/MDIO Timing (Master Mode)," on page 188
- Section 9: "Thermal Characteristics," on page 195
- Table 36: "Pin List by Pin Number 53134O/53134P/53134U," on page 150
- Table 38: "Pin List by Pin Name 53134O/53134P/53134U," on page 164

Added:

- "Egress VID Modification" on page 40

53134P-DS101; May 11, 2016

Updated:

- Figure 5: "Functional Block Diagram," on page 6
- "Multimode TX Digital-to-Analog Converter" on page 89
- Table 34: "Signal Descriptions," on page 142
- Table 35: "Pin List by Pin Number 53134O/53134P/53134U," on page 149
- Table 37: "Pin List by Pin Name 53134O/53134P/53134U," on page 163
- Table 60: "Package only, 2s2p PCB, TA = 70° C, P = 1.723W," on page 194
- Table 61: "Package with External Heat Sink 50 x 50 x 35 mm, 2s2p PCB, TA = 70° C, P = 1.723W," on page 194

- Table 62: “Package only, 2s2p PCB, TA = 55 °C, P = 1.723W,” on page 195
- Table 63: “Package with External Heat Sink 19.x 19 x 5 mm, 2s2p PCB, TA = 55° C, P = 1.723W,” on page 195

Added:

- “JTAG Interface” on page 190

53134P-DS100-DS100; December 18, 2015

Initial release.

Broadcom, the pulse logo, Connecting everything, ContentAware, RoboSwitch, BroadSync, CableChecker, Avago Technologies, Avago, and the A logo are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2018 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.



Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

Broadcom Limited:

[BCM53134PIFBG](#)