



KSZ8841-PMQL

Single-Port Ethernet MAC Controller with PCI Interface

Features

- Fully Compliant with the IEEE802.3u Standard
- Supports 10/100BASE-T/TX
- Supports IEEE 802.3x Full-Duplex Flow Control and Half-Duplex Backpressure Collision Flow Control
- Supports Burst Data Transfers
- 8 KB Internal Memory for RX/TX FIFO Buffers
- Early TX/RX Functions to Minimize Latency Through the Device
- Serial EEPROM Configuration
- Single 25 MHz Reference Clock for Both PHY and MAC

Network Features

- Fully Integrated to Comply with IEEE 802.3u Standards
- 10BASE-T and 100BASE-TX Physical Layer Support
- Auto-Negotiation: 10/100 Mbps Full- and Half-Duplex
- Supports IEEE 802.1Q Multiple VLAN Tagging
- On-Chip Wave Shaping: No External Filters Required
- Adaptive Equalizer
- Baseline Wander Correction

Power Modes, Power Supplies, and Packaging

- Single Power Supply (3.3V) with 5V Tolerant I/O Buffers
- Enhanced Power Management Feature with Power-Down Feature to Ensure Low Power Dissipation During Device Idle Periods
- Comprehensive LED Indicator Support for Link, Activity, Full-/Half-Duplex, and 10/100 Speed (4 LEDs)
- Low-Power CMOS Design
- Commercial Temperature Range: 0°C to +70°C
- Industrial Temperature Range: -40°C to +85°C
- Available in 128-Pin PQFP

Additional Features

- Single Chip Ethernet Controller with IEEE 802.3u Support
- 32-bit/33 MHz PCI Bus for Different Host Processor Interfaces
- Dynamic Buffer Memory Scheme
 - Essential for Applications such as Video over IP where Image Jitter is Unacceptable
- Microchip LinkMD[®] Cable Diagnostic Capabilities to Determine Cable Length, Diagnose Faulty Cables, and Determine Distance to Fault
- Wake-on-LAN Functionality
 - Incorporates Magic Packet[™], Network Link State, and Wake-Up Frame Technology
- HP Auto MDI-X Crossover with Disable/Enable Option
- Enhanced Power Management Feature with Power-Down Feature

Applications

- Video Distribution Systems
- High-End Cable, Satellite, and IP Set-Top Boxes
- Video over IP
- Voice over IP (VoIP) and Analog Telephone Adapters (ATA)

Markets

- Fast Ethernet
- Embedded Ethernet
- Industrial Ethernet

TO OUR VALUED CUSTOMERS

It is our intention to provide our valued customers with the best documentation possible to ensure successful use of your Microchip products. To this end, we will continue to improve our publications to better suit your needs. Our publications will be refined and enhanced as new volumes and updates are introduced.

If you have any questions or comments regarding this publication, please contact the Marketing Communications Department via E-mail at docerrors@microchip.com. We welcome your feedback.

Most Current Data Sheet

To obtain the most up-to-date version of this data sheet, please register at our Worldwide Web site at:

<http://www.microchip.com>

You can determine the version of a data sheet by examining its literature number found on the bottom outside corner of any page. The last character of the literature number is the version number, (e.g., DS30000000A is version A of document DS30000000).

Errata

An errata sheet, describing minor operational differences from the data sheet and recommended workarounds, may exist for current devices. As device/documentation issues become known to us, we will publish an errata sheet. The errata will specify the revision of silicon and revision of document to which it applies.

To determine if an errata sheet exists for a particular device, please check with one of the following:

- Microchip's Worldwide Web site; <http://www.microchip.com>
- Your local Microchip sales office (see last page)

When contacting a sales office, please specify which device, revision of silicon and data sheet (include -literature number) you are using.

Customer Notification System

Register on our web site at www.microchip.com to receive the most current information on all of our products.

Table of Contents

1.0 Introduction	4
2.0 Pin Description and Configuration	5
3.0 Functional Description	11
4.0 Register Descriptions	21
5.0 Operational Characteristics	59
6.0 Electrical Characteristics	60
7.0 Timing Specifications	61
8.0 Selection of Isolation Transformers	64
9.0 Package Outline	65
Appendix A: Data Sheet Revision History	67
The Microchip Web Site	68
Customer Change Notification Service	68
Customer Support	68
Product Identification System	69

KSZ8841-PMQL

1.0 INTRODUCTION

1.1 General Description

The KSZ8841-series single-port chip includes PCI and non-PCI CPU interfaces. This data sheet describes the KSZ8841-PMQL with PCI CPU interface chips. For information on the KSZ8841 non-PCI CPU interface chips, refer to the KSZ8841-16M/-32M data sheet.

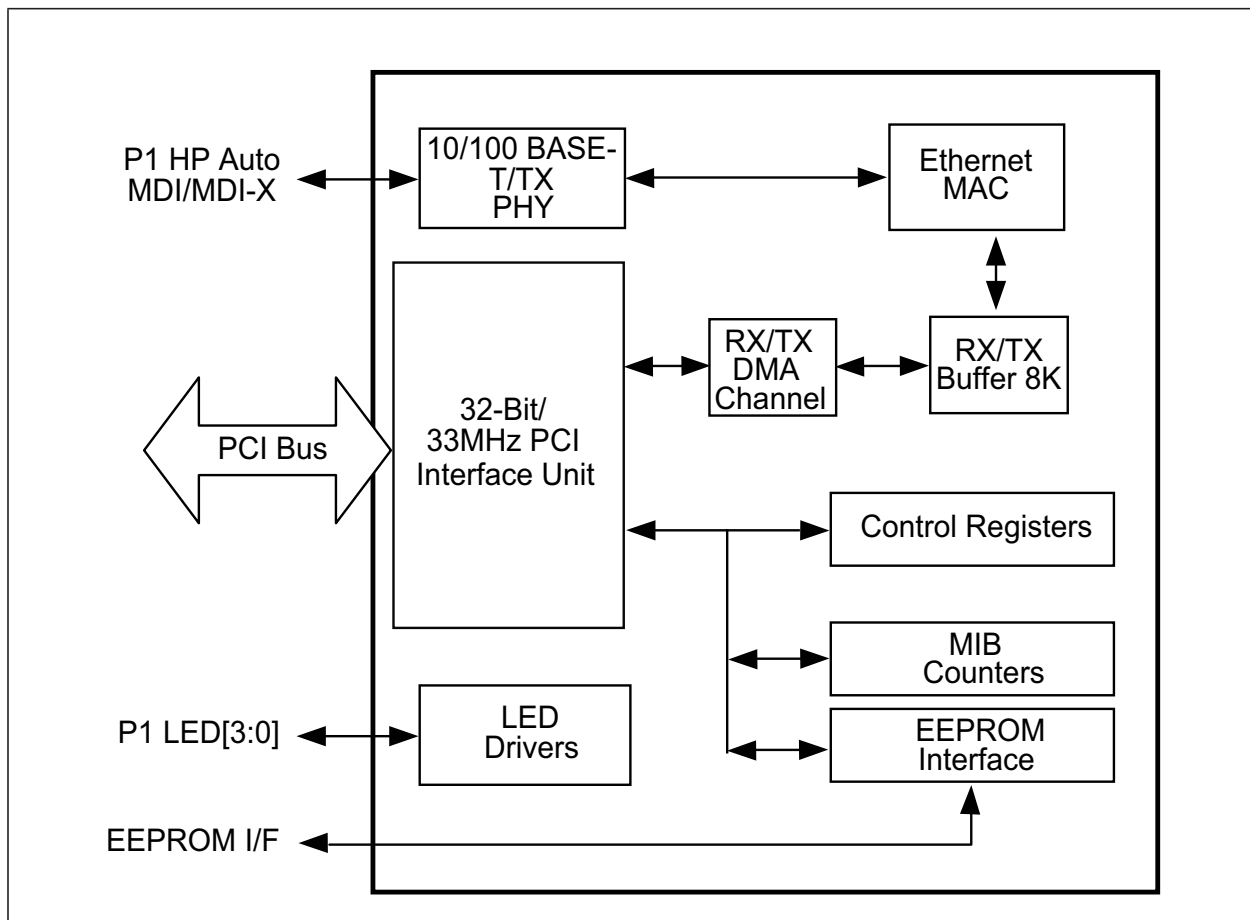
The KSZ8841-PMQL is a single-port Fast Ethernet MAC chip with a 32-bit/33 MHz PCI processor interface. Designed to be fully compliant with the IEEE 802.3u standard, the KSZ8841-PMQL is also available in an industrial temperature-grade version of the KSZ8841-PMQL, the KSZ8841-PMQLI.

Physical signal transmission and reception are enhanced through the use of analog circuitry, making the design more efficient and allowing for lower power consumption. The KSZ8841-PMQL is designed using a low-power CMOS process that features a single 3.3V power supply with 5V tolerant I/O.

The KSZ8841-PMQL is a mixed signal analog/digital device offering Wake-on-LAN technology. Its extensive feature set includes management information base (MIB) counters and CPU control/data interfaces.

The KSZ8841-PMQL includes a unique cable diagnostics feature called LinkMD[®]. This feature calculates the length of the cabling plant and determines if there is an open/short condition in the cable. Accompanying software allows the cable length and cable conditions to be conveniently displayed. In addition, the KSZ8841-PMQL supports Hewlett Packard (HP) Auto-MDIX thereby eliminating the need to differentiate between straight or crossover cables in applications.

FIGURE 1-1: SYSTEM BLOCK DIAGRAM



2.0 PIN DESCRIPTION AND CONFIGURATION

FIGURE 2-1: PIN CONFIGURATION FOR KSZ8841-PMQL (128-PIN PQFP)




TABLE 2-1: SIGNALS

Pin Number	Pin Name	Type (Note 2-1)	Description		
1	TEST_EN	I	Test Enable For normal operation, pull-down this pin to ground.		
2	SCAN_EN	I	Scan Test Scan Mux Enable For normal operation, pull-down this pin to ground.		
3 4 5	P1LED2 P1LED1 P1LED0	OPU	Port 1 LED Indicators, defined as follows		
			Chip Global Control Register: CGCR bit [15,9]		
			[0, 0] Default	[0, 1]	
			P1LED3	—	—
			P1LED2	Link/Activity	100Link/Activity
			P1LED1	Full-Duplex/Col	10Link/Activity
			P1LED0	Speed	Full-Duplex
			Reg. CGCR bit [15,9]		
			[1, 0]	[1, 1]	
			P1LED3	Activity	—
			P1LED2	Link	—
			P1LED1	Full-Duplex/Col	—
			P1LED0	Speed	—
			Note: Link = On; Activity = Blink; Link/Act = On/Blink; Full-Duplex/Col = On/Blink; Full-Duplex = On (Full-duplex); Off (Half-duplex); Speed = On (100BASE-T); Off (10BASE-T)		
Note: P1LED3 is pin 27.					
6	NC	—	No connect.		
7	NC	—	No connect.		
8	NC	—	No connect.		
9	DGND	GND	Digital ground.		
10	VDDIO	P	3.3V digital V _{DDIO}		
11	NC	—	No connect.		
12	PCLK	IPD	PCI Bus Clock This Clock provides the timing for all PCI bus phases. The rising edge defines the start of each phase. The clock maximum frequency is 33 MHz.		
13	NC	—	No connect.		
14	PMEN	OPU	Power Management Enable When asserted (Low), this signal indicates that a Wake-on-LAN packet has been received in this Ethernet MAC chip.		
15	NC	—	No connect.		
16	INTRN	OPD	Interrupt Request Active Low signal to host CPU to request an interrupt when any one of the interrupt conditions occurs in the registers. This pin should be pulled up externally.		
17	NC	—	No connect.		
18	NC	—	No connect.		
19	EECS	OPU	EEPROM Chip Select This signal is used to select an external EEPROM device.		

TABLE 2-1: SIGNALS (CONTINUED)

Pin Number	Pin Name	Type (Note 2-1)	Description
20	NC	—	No connect.
21	NC	—	No connect.
22	NC	—	No connect.
23	DGND	GND	Digital ground.
24	VDDCO	P	1.2V Core Voltage Output. (Internal 1.2V LDO power supply output) This pin provides 1.2V power supply to all 1.2V power pin, VDDC, VDDA, VDDAP.
25	NC	—	No connect.
26	EEEN	IPD	EEPROM Enable EEPROM is enabled and connected when this pin is pulled up. EEPROM is disabled when this pin is pulled down or no connect.
27	P1LED3	OPD	Port 1 LED indicator See the description in pins 3, 4, and 5.
28	EEDO	OPD	EEPROM Data Out This pin is connected to DI input of the serial EEPROM.
29	EESK	OPD	EEPROM Serial Clock 4 μ s serial clock to load configuration data from the serial EEPROM.
30	EEDI	IPD	EEPROM Data In This pin is connected to DO output of the serial EEPROM.
31	NC	—	No connect.
32	VDDIO	P	3.3V digital V_{DDIO}
33	VDDIO	P	3.3V digital V_{DDIO}
34	DGND	GND	Digital ground
35	DGND	GND	Digital ground
36	PWRDN	IPU	Full-chip power-down. Active Low
37	AGND	GND	Analog ground
38	VDDA	P	1.2V analog V_{DD}
39	AGND	GND	Analog ground
40	NC	—	No Connect
41	NC	—	No Connect
42	AGND	GND	Analog ground
43	VDDA	P	1.2V analog V_{DD}
44	NC	—	No Connect
45	RXP1	I/O	Port 1 physical receive (MDI) or transmit (MDIX) signal (+ differential)
46	RXM1	I/O	Port 1 physical receive (MDI) or transmit (MDIX) signal (– differential)
47	AGND	GND	Analog ground
48	TXP1	I/O	Port 1 physical transmit (MDI) or receive (MDIX) signal (+ differential)
49	TXM1	I/O	Port 1 physical transmit (MDI) or receive (MDIX) signal (– differential)
50	VDDATX	P	3.3V analog V_{DD} .
51	VDDARX	P	3.3V analog V_{DD} .
52	NC	—	No Connect
53	NC	—	No Connect
54	AGND	GND	Analog ground
55	NC	—	No Connect
56	NC	—	No Connect

KSZ8841-PMQL

TABLE 2-1: SIGNALS (CONTINUED)

Pin Number	Pin Name	Type (Note 2-1)	Description
57	VDDA	P	1.2 analog V_{DD} .
58	AGND	GND	Analog ground
59	NC	—	No connect
60	NC	—	No connect
61	ISSET	O	Set physical transmit output current. Pull-down this pin with a 3.01 k Ω 1% resistor to ground.
62	AGND	GND	Analog ground
63	VDDAP	P	1.2V analog V_{DD} for PLL.
64	AGND	GND	Analog ground
65	X1	I	25 MHz crystal or oscillator clock connection.
66	X2	O	Pins (X1, X2) connect to a crystal. If an oscillator is used, X1 connects to a 3.3V tolerant oscillator and X2 is a no connect. Note: Clock requirement is ± 50 ppm for either crystal or oscillator.
67	RSTN	IPU	Hardware Reset, Active-Low RSTN will cause the KSZ8841-PMQL to reset all of its functional blocks. RSTN must be asserted for a minimum duration of 10 ms.
68	PAR	I/O	PCI Parity Even parity computed for PAD[31:0] and CBE[3:0]N, master drives PAR for address and write data phase, target drives PAR for read data phase.
69	FRAMEN	I/O	PCI Cycle Frame This signal is asserted low to indicate the beginning of the address phase of the bus transaction and deasserted before the final transfer of the data phase of the transaction in a bus master mode. As a target, the device monitors this signal before decoding the address to check if the current transaction is addressed to it.
70	IRDYN	I/O	PCI Initiator Ready As a bus master, this signal is asserted low to indicate valid data phases on PAD[31:0] during write data phases, indicates it is ready to accept data during read data phases. As a target, it'll monitor this IRDYN signal that indicates the master has put the data on the bus.
71	TRDYN	I/O	PCI Target Ready As a bus target, this signal is asserted low to indicate valid data phases on PAD[31:0] during read data phases, indicates it is ready to accept data during write data phases. As a master, it will monitor this TRDYN signal that indicates the target is ready for data during read/write operation.
72	STOPN	I/O	PCI Stop This signal is asserted low by the target device to request the master device to stop the current transaction.
73	IDSEL	I/O	PCI Initialization Device Select This signal is used to select the KSZ8841-PMQL during configuration read and write transactions. Active-high.
74	DEVSELN	I/O	PCI Device Select This signal is asserted low when it is selected as a target during a bus transaction. As a bus master, the KSZ8841-PMQL samples this signal to ensure that a PCI target recognizes the destination address for the data transfer.
75	REQN	O	PCI Bus Request The KSZ8841-PMQL will assert this signal low to request PCI bus master operation.

TABLE 2-1: SIGNALS (CONTINUED)

Pin Number	Pin Name	Type (Note 2-1)	Description
76	GNTN	I	PCI Bus Grant This signal is asserted low to indicate to the KSZ8841-PMQL that it has been granted the PCI bus master operation.
77	PERRN	I/O	PCI Parity Error The KSZ8841-PMQL as a master or target will assert this signal low to indicate a parity error on any incoming data. As a bus master, it will monitor this signal on all write operations.
78	DGND	GND	Digital ground
79	VDDIO	P	3.3V digital V_{DDIO}
80	SERRN	O	PCI System Error This system error signal is asserted low by the KSZ8841-PMQL. This signal is used to report address parity errors.
81	NC	—	No Connect
82	NC	—	No Connect
83	NC	—	No Connect
84	NC	—	No Connect
85	CBE3N	I/O	Command and Byte Enable These signals are multiplexed on the same PCI pins. During the address phase, these lines define the bus command. During the data phase, these lines are used as Byte Enables, The Byte enables are valid for the entire data phase and determine which byte lanes carry meaningful data.
86	CBE2N	I/O	
87	CBE1N	I/O	
88	CBE0N	I/O	
89	PAD31	I/O	PCI Address/Data 31 Address and data are multiplexed on the all of the PAD pins. The PAD pins carry the physical address during the first clock cycle of a transaction and carry data during the subsequent clock cycles.
90	DGND	GND	Digital core ground
91	VDDC	P	1.2V digital core V_{DD}
92	VDDIO	P	3.3V digital V_{DDIO}
93	PAD30	I/O	PCI Address/Data 30
94	PAD29	I/O	PCI Address/Data 29
95	PAD28	I/O	PCI Address/Data 28
96	PAD27	I/O	PCI Address/Data 27
97	PAD26	I/O	PCI Address/Data 26
98	PAD25	I/O	PCI Address/Data 25
99	PAD24	I/O	PCI Address/Data 24
100	PAD23	I/O	PCI Address/Data 23
101	PAD22	I/O	PCI Address/Data 22
102	PAD21	I/O	PCI Address/Data 21
103	PAD20	I/O	PCI Address/Data 20
104	PAD19	I/O	PCI Address/Data 19
105	PAD18	I/O	PCI Address/Data 18
106	PAD17	I/O	PCI Address/Data 17
107	DGND	GND	Digital ground
108	VDDIO	P	3.3V digital V_{DDIO}
109	PAD16	I/O	PCI Address/Data 16
110	PAD15	I/O	PCI Address/Data 15

KSZ8841-PMQL

TABLE 2-1: SIGNALS (CONTINUED)

Pin Number	Pin Name	Type (Note 2-1)	Description
111	PAD14	I/O	PCI Address/Data 14
112	PAD13	I/O	PCI Address/Data 13
113	PAD12	I/O	PCI Address/Data 12
114	PAD11	I/O	PCI Address/Data 11
115	PAD10	I/O	PCI Address/Data 10
116	PAD9	I/O	PCI Address/Data 9
117	PAD8	I/O	PCI Address/Data 8
118	PAD7	I/O	PCI Address/Data 7
119	PAD6	I/O	PCI Address/Data 6
120	PAD5	I/O	PCI Address/Data 5
121	PAD4	I/O	PCI Address/Data 4
122	PAD3	I/O	PCI Address/Data 3
123	DGND	GND	Digital IO ground
124	DGND	GND	Digital core ground
125	VDDIO	P	3.3V digital V _{DDIO}
126	PAD2	I/O	PCI Address/Data 2
127	PAD1	I/O	PCI Address/Data 1
128	PAD0	I/O	PCI Address/Data 0

Note 2-1 P = power supply; GND = ground; I = input; O = output
I/O = bi-directional
IPU/O = Input with internal pull-up during reset; output pin otherwise.
IPU = Input with internal pull-up.
IPD = Input with internal pull-down.
OPU = Output with internal pull-up.
OPD = Output with internal pull-down.

3.0 FUNCTIONAL DESCRIPTION

The KSZ8841-PMQL is a single-chip Fast Ethernet MAC controller consisting of a 10/100 physical layer transceiver (PHY), a MAC, and a PCI interface unit that controls the KSZ8841-PMQL via a 32-bit/33 MHz PCI processor interface.

The KSZ8841-PMQL is fully compliant to the IEEE802.3u standard.

3.1 PCI Bus Interface Unit

3.1.1 PCI BUS INTERFACE

The PCI Bus Interface implements PCI v2.2 bus protocols and configuration space. The KSZ8841-PMQL supports bus master reads and writes to CPU memory, and CPU access to on-chip register space. When the CPU reads and writes the configuration registers of the KSZ8841-PMQL, it is as a slave. So the KSZ8841-PMQL can be either a PCI bus master or slave. The PCI Bus Interface is also responsible for managing the DMA interfaces and the host processors access. Arbitration logic within the PCI Bus Interface unit accepts bus requests from the TXDMA logic and RXDMA logic.

The PCI bus interface also manages interrupt generation for a host processor.

3.1.2 TXDMA LOGIC AND TX BUFFER MANAGER

The KSZ8841-PMQL supports a multi-frame, multi-fragment DMA gather process. Descriptors representing frames are built and linked in system memory by a host processor. The TXDMA logic is responsible for transferring the multi-fragment frame data from the host memory into the TX buffer.

The KSZ8841-PMQL uses 4K bytes of transmit data buffer between the TXDMA logic and transmit MAC. When the TXDMA logic determines there is enough space available in the TX buffer, the TXDMA logic will move any pending frame data into the TX buffer. The management mechanism depends on the transmit descriptor list.

3.1.3 RXDMA LOGIC AND RX BUFFER MANAGER

The KSZ8841-PMQL supports a multi-frame, multi-fragment DMA scatter process. Descriptors representing frames are built and linked in system memory by the host processor. The RXDMA logic is responsible for transferring the frame data from the RX buffer to the host memory.

The KSZ8841-PMQL uses 4K bytes of receive data buffer between the receive MAC and RXDMA logic. The management mechanism depends on the receive descriptor list.

3.2 Power Management

3.2.1 POWER DOWN

The KSZ8841-PMQL features a port power-down mode. To save power, the user can power down this port that is not in use by setting bit 11 in either P1CR4 or P1MBCR register for this port. To bring the port back up, reset bit 11 in these registers.

In addition, there is a full chip power-down mode by pulling down the PWRDN pin 36. When this pin is pulled down, the entire chip powers down. Transitioning this pin from pull-down to pull-up results in a power up and chip reset.

3.2.2 WAKE-ON-LAN

Wake-up frame events are used to wake the system whenever meaningful data is presented to the system over the network. Examples of meaningful data include the reception of a Magic Packet, a management request from a remote administrator, or simply network traffic directly targeted to the local system. In all of these instances, the network device is pre-programmed by the policy owner or other software with information on how to identify wake frames from other network traffic.

A wake-up event is a request for hardware and/or software external to the network device to put the system into a powered state (working).

A wake-up signal is caused by:

1. Detection of a change in the network link state
2. Receipt of a network wake-up frame
3. Receipt of a Magic Packet

KSZ8841-PMQL

3.2.3 LINK CHANGE

Link status wake events are useful to indicate a change in the network's availability, especially when this change may impact the level at which the system should re-enter the sleeping state. For example, a change from link off to link on may trigger the system to re-enter sleep at a higher level (D2 versus D3) so that wake frames can be detected. Conversely, a transition from link on to link off may trigger the system to re-enter sleep at a deeper level (D3 versus D2) because the network is not currently available.

Note that references to D0, D1, D2, and D3 are power management states defined in a similar fashion to the way they are defined for PCI.

3.2.4 WAKE-UP PACKET

Wake-up packets are certain types of packets with specific CRC values that a system recognizes as a 'wake-up' frame. The KSZ8841-PMQL supports up to four users defined wake-up frames as below:

1. Wake-up frame 0 is defined in registers 0x0220-0x022A and is enabled by bit 0 in wakeup frame control register.
2. Wake-up frame 1 is defined in registers 0x0230-0x023A and is enabled by bit 1 in wakeup frame control register.
3. Wake-up frame 2 is defined in registers 0x0240-0x024A and is enabled by bit 2 in wakeup frame control register.
4. Wake-up frame 3 is defined in registers 0x0250-0x025A and is enabled by bit 3 in wakeup frame control register.

3.2.5 MAGIC PACKET

Magic Packet technology is used to remotely wake up a sleeping or powered-off PC on a LAN. This is accomplished by sending a specific packet of information, called a Magic Packet frame, to a node on the network. When a PC capable of receiving the specific frame goes to sleep, it enables the Magic Packet RX mode in the LAN controller, and when the LAN controller receives a Magic Packet frame, it will alert the system to wake up.

Magic Packet is a standard feature integrated into the KSZ8841-PMQL. The chip implements multiple advanced power-down modes including Magic Packet to conserve power and operate more efficiently.

Once the KSZ8841-PMQL has been put into Magic Packet Enable mode (WFCR[7]=1), it scans all incoming frames addressed to the node for a specific data sequence, which indicates to the chip this is a Magic Packet (MP) frame.

A Magic Packet frame must also meet the basic requirements for the LAN technology chosen, such as Source Address (SA), Destination Address (DA), which may be the receiving station's IEEE address or a multicast or broadcast address and CRC.

The specific sequence consists of 16 duplications of the IEEE address of this node, with no breaks or interruptions. This sequence can be located anywhere within the packet, but must be preceded by a synchronization stream. The synchronization stream allows the scanning state machine to be much simpler. The synchronization stream is defined as 6 bytes of FFh. The device will also accept a broadcast frame, as long as the 16 duplications of the IEEE address match the address of the machine to be awakened.

Example:

If the IEEE address for a particular node on a network is 11h 22h, 33h, 44h, 55h, 66h, the LAN controller would be scanning for the data sequence (assuming an Ethernet frame):

DESTINATION SOURCE MISC: FF FF FF FF FF FF - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - 11 22 33 44 55 66 - MISC -CRC.

There are no further restrictions on a Magic Packet frame. For instance, the sequence could be in a TCP/IP packet or an IPX packet. The frame may be bridged or routed across the network without affecting its ability to wake-up a node at the frame's destination

If the LAN controller scans a frame and does not find the specific sequence shown above, it discards the frame and takes no further action. If the controller (KSZ8841-PMQL) detects the data sequence, however, it then alerts the PC's power management circuitry (asserted the PMEN pin) to wake up the system.

3.3 Physical Layer Transceiver

3.3.1 100BASE-TX TRANSMIT

The 100BASE-TX transmit function performs parallel-to-serial conversion, 4B/5B coding, scrambling, NRZ-to-NRZI conversion, and MLT3 encoding and transmission.

The circuitry starts with a parallel-to-serial conversion, which converts the 25 MHz 4-bit nibbles into a 125 MHz serial bit stream. The data and control stream is then converted into 4B/5B coding, followed by a scrambler. The serialized data is further converted from NRZ-to-NRZI format, and then transmitted in MLT3 current output. An external 1% 3.01 k Ω resistor for the 1:1 transformer ratio sets the output current.

The output signal has a typical rise/fall time of 4 ns and complies with the ANSI TP-PMD standard regarding amplitude balance, overshoot, and timing jitter. The wave-shaped 10BASE-T output driver is also incorporated into the 100BASE-TX driver.

3.3.2 100BASE-TX RECEIVE

The 100BASE-TX receiver function performs adaptive equalization, DC restoration, MLT3-to-NRZI conversion, data and clock recovery, NRZI-to-NRZ conversion, de-scrambling, 4B/5B decoding, and serial-to-parallel conversion.

The receiving side starts with the equalization filter to compensate for inter-symbol interference (ISI) over the twisted pair cable. Because the amplitude loss and phase distortion is a function of the cable length, the equalizer has to adjust its characteristics to optimize performance. In this design, the variable equalizer makes an initial estimation based on comparisons of incoming signal strength against some known cable characteristics, and then tunes itself for optimization. This is an ongoing process and self-adjusts against environmental changes such as temperature variations.

Next, the equalized signal goes through a DC restoration and data conversion block. The DC restoration circuit is used to compensate for the effect of baseline wander and to improve the dynamic range. The differential data conversion circuit converts the MLT3 format back to NRZI. The slicing threshold is also adaptive.

The clock recovery circuit extracts the 125MHz clock from the edges of the NRZI signal. This recovered clock is then used to convert the NRZI signal into the NRZ format. This signal is sent through the de-scrambler followed by the 4B/5B decoder. Finally, the NRZ serial data is converted to an MII format and provided as the input data to the MAC.

3.3.3 PLL CLOCK SYNTHESIZER (RECOVERY)

The internal PLL clock synthesizer generates 125 MHz, 62.5 MHz, 41.66 MHz, and 25 MHz clocks by setting the on-chip bus speed control register OBCR for KSZ8841-PMQL system timing. These internal clocks are generated from an external 25 MHz crystal or oscillator.

Note that the default setting is 25 MHz in the OBCR register; it is recommended that the software driver set it to 125 MHz for best performance.

3.3.4 SCRAMBLER/DE-SCRAMBLER (100BASE-TX ONLY)

The purpose of the scrambler is to spread the power spectrum of the signal to reduce electromagnetic interference (EMI) and baseline wander.

Transmitted data is scrambled through the use of an 11-bit wide linear feedback shift register (LFSR). The scrambler generates a 2047-bit non-repetitive sequence. Then the receiver de-scrambles the incoming data stream using the same sequence as at the transmitter.

3.3.5 10BASE-T TRANSMIT

The 10BASE-T driver is incorporated with the 100BASE-TX driver to allow for transmission using the same magnetics. They are internally wave-shaped and pre-emphasized into outputs with a typical 2.4V amplitude. The harmonic contents are at least 27 dB below the fundamental frequency when driven by an all-ones Manchester-encoded signal.

3.3.6 10BASE-T RECEIVE

On the receive side, input buffers and level detecting squelch circuits are employed. A differential input receiver circuit and a phase-locked loop (PLL) perform the decoding function.

The Manchester-encoded data stream is separated into clock signal and NRZ data. A squelch circuit rejects signals with levels less than 400 mV or with short pulse widths to prevent noise at the RXP-or-RXM input from falsely triggering the decoder. When the input exceeds the squelch limit, the PLL locks onto the incoming signal and the KSZ8841-PMQL decodes a data frame.

The receiver clock is maintained active during idle periods in between data reception.

3.3.7 MDI/MDI-X AUTO CROSSOVER

To eliminate the need for crossover cables between similar devices, the KSZ8841-PMQL supports HP-Auto MDI/MDI-X and IEEE 802.3u standard MDI/MDI-X auto crossover. HP-Auto MDI/MDI-X is the default.

KSZ8841-PMQL

The auto-sense function detects remote transmit and receive pairs and correctly assigns the transmit and receive pairs for the Ksz8841-PMQL device. This feature is extremely useful when end users are unaware of cable types in addition to saving on an additional uplink configuration connection. The auto-crossover feature can be disabled through the port control registers.

The IEEE 802.3u standard MDI and MDI-X definitions are illustrated in [Table 3-1](#).

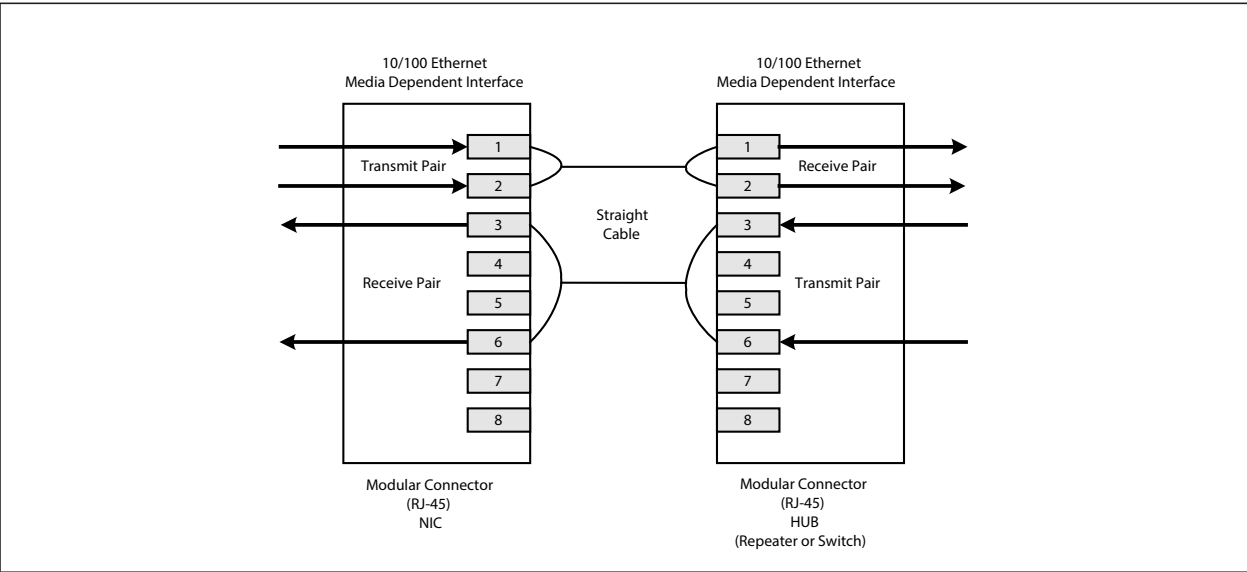
TABLE 3-1: MDI/MDI-X PIN DEFINITIONS

MDI		MDI-X	
RJ-45 Pins	Signals	RJ-45 Pins	Signals
1	TD+	1	RD+
2	TD–	2	RD–
3	RD+	3	TD+
6	RD–	6	TD–

3.3.7.1 Straight Cable

A straight cable connects an MDI device to an MDI-X device, or an MDI-X device to an MDI device. [Figure 3-1](#) depicts a typical straight cable connection between a NIC card (MDI) and a switch or hub (MDI-X).

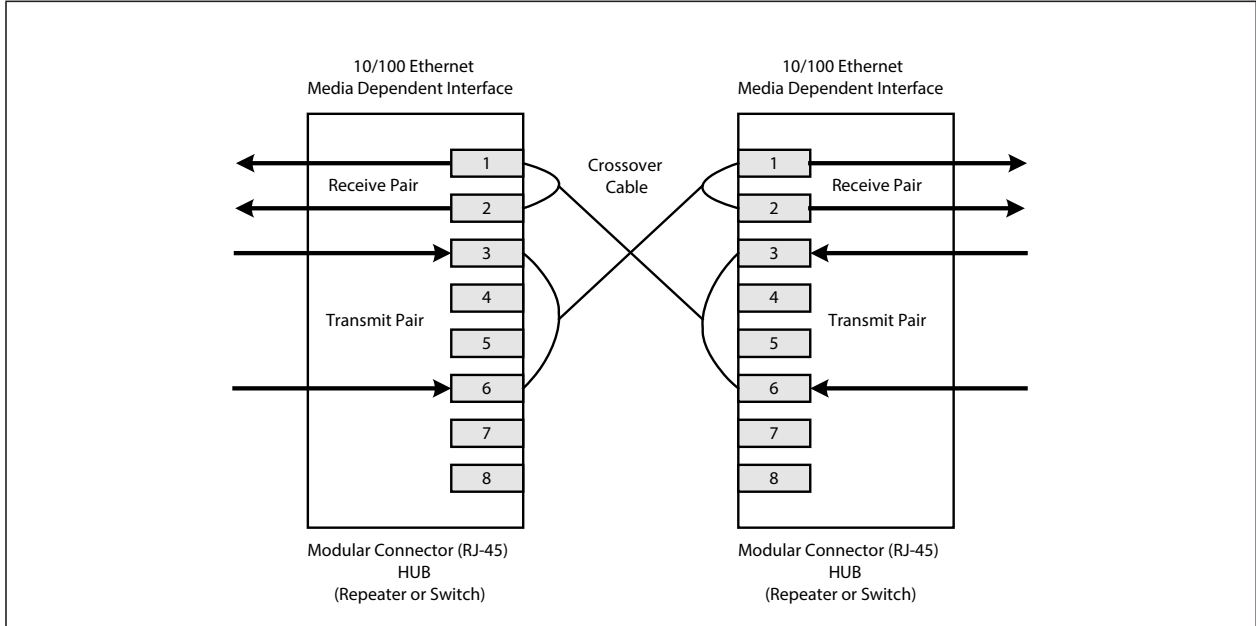
FIGURE 3-1: TYPICAL STRAIGHT CABLE CONNECTION



3.3.7.2 Crossover Cable

A crossover cable connects an MDI device to another MDI device, or an MDI-X device to another MDI-X device. [Figure 3-2](#) shows a typical crossover cable connection between two switches or hubs (two MDI-X devices).

FIGURE 3-2: TYPICAL CROSSOVER CABLE CONNECTION



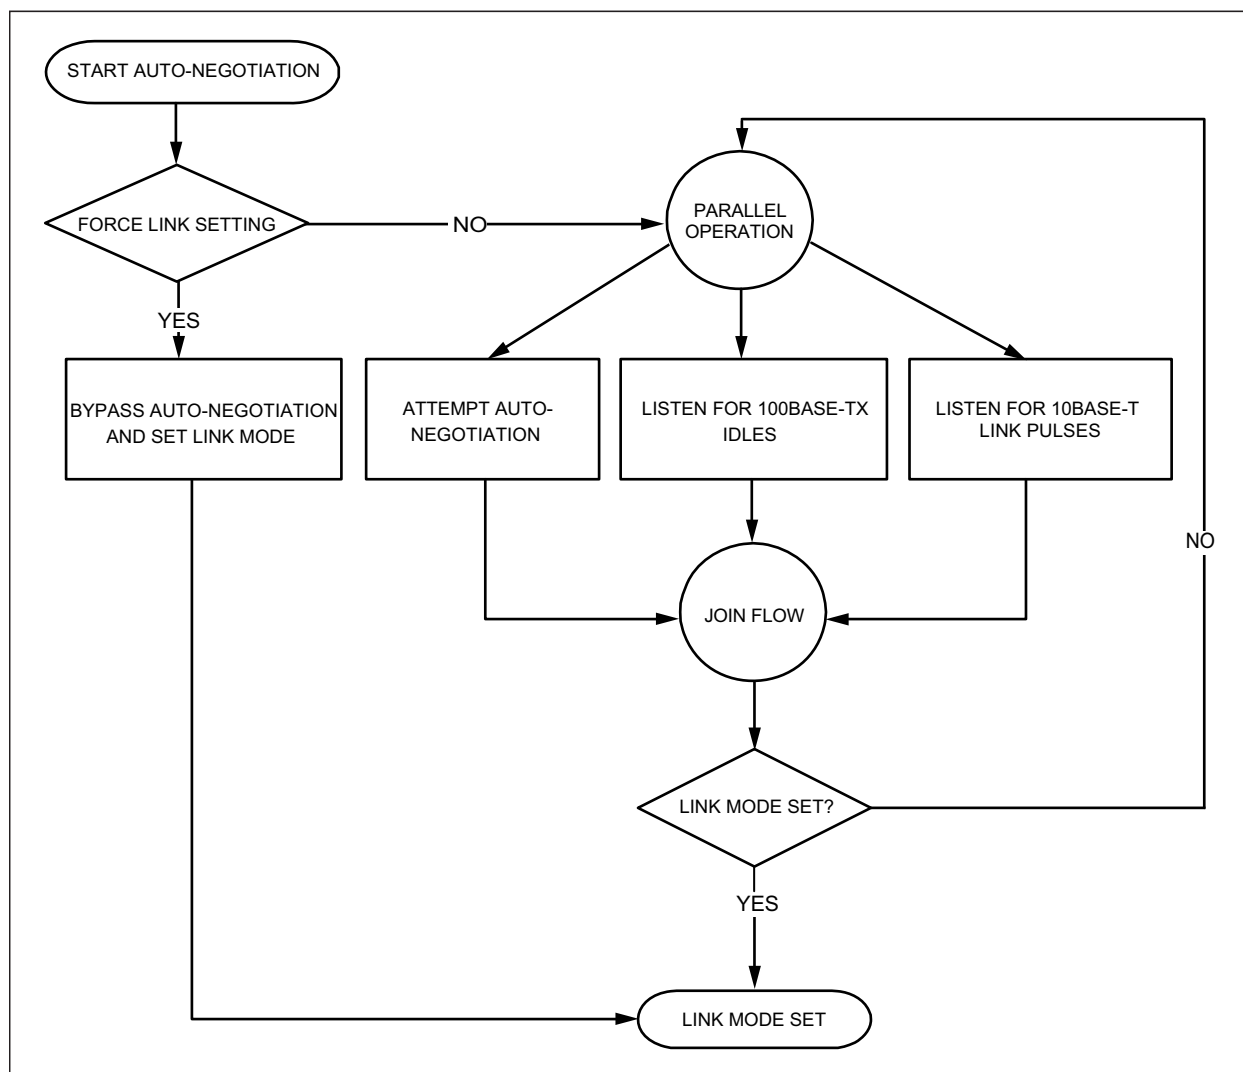
3.3.8 AUTO-NEGOTIATION

The KSZ8841-PMQL conforms to the auto negotiation protocol as described by the 802.3 committee to allow the port to operate at either 10BASE-T or 100BASE-TX.

Auto negotiation allows unshielded twisted pair (UTP) link partners to select the best common mode of operation. In auto negotiation, the link partners advertise capabilities across the link to each other. If auto negotiation is not supported or the link partner to the KSZ8841-PMQL is forced to bypass auto negotiation, the mode is set by observing the signal at the receiver. This is known as parallel mode because while the transmitter is sending auto negotiation advertisements, the receiver is listening for advertisements or a fixed signal protocol.

The link setup process is shown in [Figure 3-3](#).

FIGURE 3-3: AUTO-NEGOTIATION AND PARALLEL OPERATION



3.3.9 LINKMD® CABLE DIAGNOSTICS

The KSZ8841-PMQL LinkMD® uses time domain reflectometry (TDR) to analyze the cabling plant for common cabling problems such as open circuits, short circuits, and impedance mismatches.

LinkMD works by sending a pulse of known amplitude and duration down the MDI and MDI-X pairs and then analyzes the shape of the reflected signal. Timing the pulse duration gives an indication of the distance to the cabling fault with a maximum distance of 200m and an accuracy of $\pm 2\text{m}$.

Note that cable diagnostics are only valid for copper connections. Fiber-optic operation is not supported.

3.3.9.1 Access

LinkMD is initiated by accessing register P1VCT, the LinkMD Control/Status register, in conjunction with register P1CR4, the 100BASE-TX PHY Controller register.

3.3.9.2 Usage

LinkMD can be run at any time. To use LinkMD, disable HP Auto-MDIX by writing a '1' to P1CR4[10] to enable manual control over the pair used to transmit the LinkMD pulse. The self-clearing cable diagnostic test enable bit, P1VCT[15], is set to '1' to start the test on this pair.

When bit P1VCT[15] returns to '0', the test is complete. The test result is returned in bits P1VCT[14:13] and the distance is returned in bits P1VCT[8:0]. The cable diagnostic test results are as follows:

00 = Valid test, normal condition

01 = Valid test, open circuit in cable

10 = Valid test, short circuit in cable

11 = Invalid test, LinkMD failed

If P1VCT[14:13]=11, this indicates an invalid test, and occurs when the KSZ8841-PMQL is unable to shut down the link partner. In this instance, the test is not run, as it is not possible for the KSZ8841-PMQL to determine if the detected signal is a reflection of the signal generated or a signal from another source.

Cable distance can be approximated by the following formula:

Distance = P1VCT[8:0] x 0.4m

This constant may be calibrated for different cabling conditions, including cables with a velocity of propagation that varies significantly from the norm.

3.4 Media Access Control (MAC) Operation

The KSZ8841-PMQL strictly abides by IEEE 802.3 standards to maximize compatibility.

3.4.1 INTER PACKET GAP (IPG)

If a frame is successfully transmitted, then the minimum 96-bit time for IPG is measured between two consecutive packets. If the current packet is experiencing collisions, the minimum 96-bit time for IPG is measured from carrier sense (CRS) to the next transmit packet.

3.4.2 BACK-OFF ALGORITHM

The KSZ8841-PMQL implements the IEEE standard 802.3 binary exponential back-off algorithm in half-duplex mode. After 16 collisions, the packet is dropped.

3.4.3 LATE COLLISION

If a transmit packet experiences collisions after 512 bit times of the transmission, the packet is dropped.

3.4.4 FLOW CONTROL

The KSZ8841-PMQL supports standard 802.3x flow control frames on both transmit and receive sides.

On the receive side, if the KSZ8841-PMQL receives a pause control frame, the KSZ8841-PMQL will not transmit the next normal frame until the timer, specified in the pause control frame, expires. If another pause frame is received before the current timer expires, the timer will be updated with the new value in the second pause frame. During this period (while it is flow controlled), only flow control packets from the KSZ8841-PMQL are transmitted.

KSZ8841-PMQL

On the transmit side, the KSZ8841-PMQL has intelligent and efficient ways to determine when to invoke flow control. The flow control is based on the availability of the system resources.

The KSZ8841-PMQL issues a flow control frame (XON), containing the maximum pause time defined in IEEE standard 802.3x. Once the resource is freed up, the KSZ8841-PMQL sends out another flow control frame (XOFF) with zero pause time to turn off the flow control (turn on transmission to the port). A hysteresis feature is provided to prevent the flow control mechanism from being constantly activated and deactivated.

3.4.5 HALF-DUPLEX BACKPRESSURE

A half-duplex backpressure option (non-IEEE 802.3 standards) is also provided. The activation and deactivation conditions are the same as in full-duplex mode. If backpressure is required, the KSZ8841-PMQL sends preambles to defer the other stations' transmission (carrier sense deference).

To avoid jabber and excessive deference (as defined in the 802.3 standard), after a certain time, the KSZ8841-PMQL discontinues the carrier sense and then raises it again quickly. This short silent time (no carrier sense) prevents other stations from sending out packets thus keeping other stations in a carrier sense deferred state. If the port has packets to send during a backpressure situation, the carrier sense type backpressure is interrupted and those packets are transmitted instead. If there are no additional packets to send, carrier sense type backpressure is reactivated again until chip resources free up. If a collision occurs, the binary exponential back-off algorithm is skipped and carrier sense is generated immediately, thus reducing the chance of further collision and carrier sense is maintained to prevent packet reception.

The backpressure will take effect automatically in auto-negotiation enable and half-duplex mode.

3.4.6 CLOCK GENERATOR

The X1 and X2 pins are connected to a 25 MHz crystal. X1 can also serve as the connector to the 3.3V 25 MHz oscillator (as described in the pin description table).

3.4.7 EEPROM INTERFACE

An external serial EEPROM with a standard microwire bus interface is used for non-volatile storage of information such as the node address and subsystem ID.

As part of the initialization after system reset, the KSZ8841-PMQL reads the external EEPROM and places the data into certain host-accessible registers. If the EEEN pin is pulled up, the KSZ8841-PMQL performs an automatic read of the EEPROM word from 0x0 to 0x6 after the deassertion of Reset. An EEPROM of 1 KB (93C46) or 4 KB (93C66) can be used based on application.

The EEPROM read/write function can also be performed by software reading and writing to the EEPCR register.

The KSZ8841M EEPROM format is given in [Table 3-2](#).

TABLE 3-2: KSZ8841M EEPROM FORMAT

Word	15 - 8	7 - 0
0x0	Base Address	
0x1	Host MAC Address Byte 2	Host MAC Address Byte 1
0x2	Host MAC Address Byte 4	Host MAC Address Byte 3
0x3	Host MAC Address Byte 6	Host MAC Address Byte 5
0x4	Subsystem ID	
0x5	Subsystem Vendor ID	
0x6	ConfigParam	
0x7 - 0x3F	Not used by KSZ8841-PMQL (available for user to use)	

The format for ConfigParam is shown in [Table 3-3](#).

TABLE 3-3: CONFIGPARAM WORD IN EEPROM FORMAT

Bit	Bit Name	Description
15	NEW_CAP	<p>New Capabilities Indicates whether or not the KSZ8841-PMQL implements a list of new capabilities. When set, this bit indicates the presence of new capabilities. When reset, New capabilities are not implemented. The value of this bit is loaded to the NEW_CAP bit in CFCS register.</p>
14	NO_SRST	<p>No Soft Reset When this bit is set, indicates that KSZ8841-PMQL transitioning from D3_hot to D0 because of PowerState commands do not perform an internal reset. Configuration Context is preserved. Upon transition from the D3_hot to the D0 Initialized state, no additional operating system intervention is required to preserve Configuration Context beyond writing the PowerState bits. When this bit is clear, KSZ8841-PMQL performs an internal reset upon transitioning from D3_hot to D0 via software control of the PowerState bits. Configuration Context is lost when performing the soft reset. Upon transition from the D3_hot to the D0 state, full reinitialization sequence is needed to return the device to D0 Initialized. Regardless of this bit, devices that transition from D3_hot to D0 by a system or bus segment reset will return to the device state D0 Uninitialized with only PME context preserved if PME is supported and enabled. This bit is loaded to bit 3 of CPMC register</p>
13	Reserved	—
12	PME_D2	<p>PME -Support D2 When this bit is set, the KSZ8841-PMQL asserts PME event when the KSZ8841-PMQL is in D2 state and PME_EN is set. Otherwise, the KSZ8841-PMQL does not assert PME event when the KSZ8841-PMQL is in D2 state. This bit is loaded to bit 13 of PMCR register, and bit 29 of CCID register.</p>
11	PME_D1	<p>PME Support D1 When this bit is set, the KSZ8841-PMQL asserts PME event when the KSZ8841-PMQL is in D1 state and PME_EN is set. Otherwise, the KSZ8841-PMQL does not assert PME event when the KSZ8841-PMQL is in D1 state. This bit is loaded to bit 12 of PMCR register, and bit 28 of CCID register.</p>
10	D2_SUP	<p>D2 support When this bit is set, the KSZ8841-PMQL supports D2 power state. This bit is loaded to bit 10 of PMCR register, and bit 26 of CCID register.</p>
9	D1_SUP	<p>D1 support When this bit is set, the KSZ8841-PMQL supports D1 power state. This bit is loaded to bit 9 of PMCR register, and bit 25 of CCID register.</p>
8 - 6	Reserved	—
5	DSI	<p>Device Specific Initialization This bit indicates whether special initialization of this function is required (beyond the standard PCI configuration header) before the generic class device driver is able to use it. A “1” indicates that the function requires a device specific initialization sequence following transition to the D0 uninitialized state. This bit is loaded to bit 5 of PMCR register and bit 21 of CCID register.</p>
4	Reserved	—
3	PME_CK	<p>PME Clock When this bit is a “1”, it indicates that the function relies on the presence of the PCI clock for PME# operation. When this bit is a “0”, it indicates that no PCI clock is required for the function to generate PME#. This bit is loaded to bit 3 of PMCR register and bit 19 of CCID register.</p>

KSZ8841-PMQL

TABLE 3-3: CONFIGPARAM WORD IN EEPROM FORMAT (CONTINUED)

Bit	Bit Name	Description
2 - 0	PCI: PME_VER	PCI: Power Management PCI Version. These bits are loaded to bits [2:0] of the PMCR register and bits [18:16] of the CCID register.

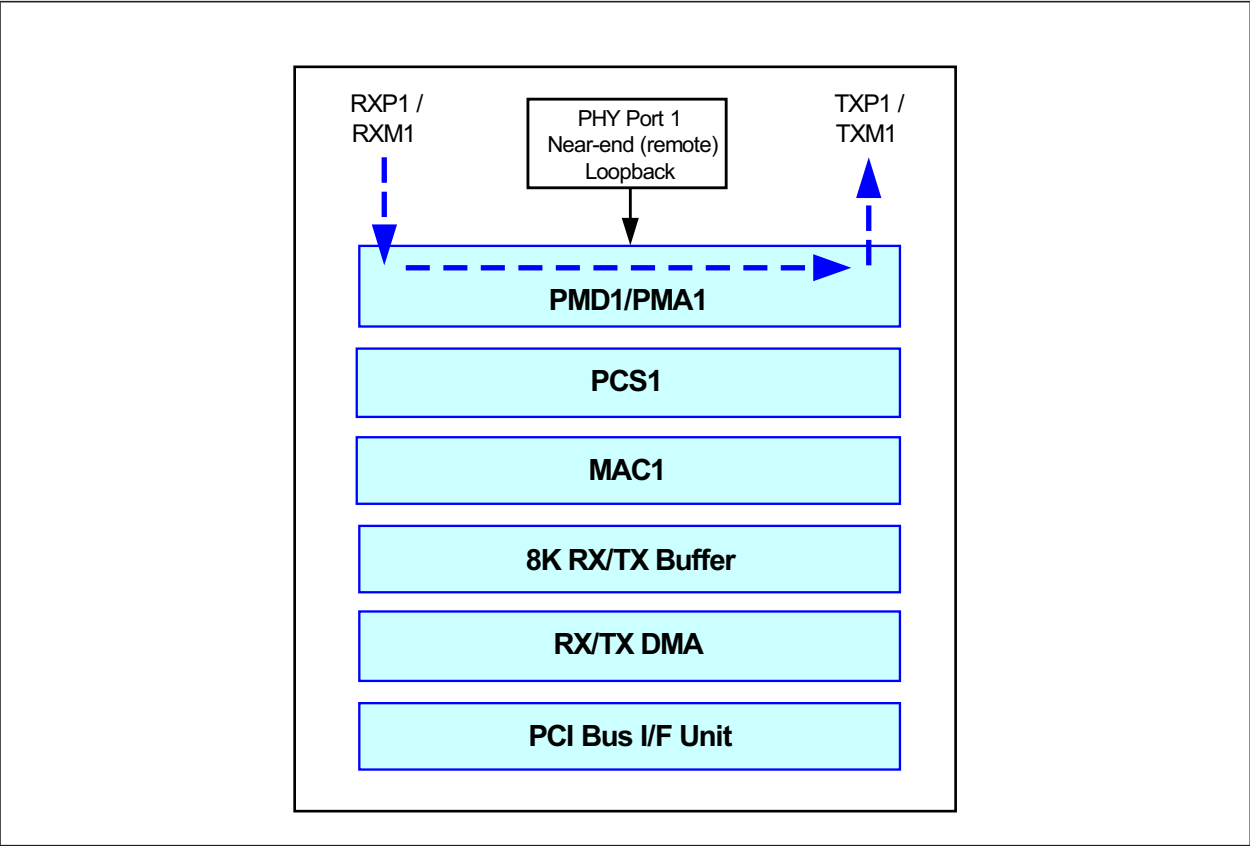
3.4.8 LOOPBACK SUPPORT

The KSZ8841-PMQL provides loopback support for remote diagnostic failure. In loopback mode, the speed at the PHY port will be set to 100BASE-TX full-duplex mode. The KSZ8841-PMQL only supports Near-end (Remote) Loopback.

Near-end (Remote) loopback is conducted at PHY port 1 of the KSZ8841-PMQL. The loopback path starts at the PHY ports receive inputs (RXPx/RXMx), wraps around at the same PHY port's PMD/PMA, and ends at the PHY ports transmit outputs (TXPx/TXMx).

Bit [1] of register P1PHYCTRL is used to enable near-end loopback for port 1. Alternatively, Bit [9] of register P1SCSLMD can also be used to enable near-end loopback. The port's near-end loopback path is illustrated in the following [Figure 3-4](#).

FIGURE 3-4: PORT 1 NEAR-END (REMOTE) LOOPBACK PATH



4.0 REGISTER DESCRIPTIONS

4.1 Host Communication

The descriptor lists and data buffers, collectively called the host communication, manage the actions and status related to buffer management. Commands and signals that control the functional operation of the KSZ8841-PMQL are also described.

The KSZ8841-PMQL and the driver communicate through the two data structures: Command and status registers (CSRs) and Descriptor Lists and Data Buffers.

Note: All unused bits of the data structure in this section are reserved and should be written by the driver as zeros.

4.1.1 HOST COMMUNICATION DESCRIPTOR LISTS AND DATA BUFFERS

The KSZ8841-PMQL transfers received data frames to the receive buffer in host memory and transmits data from the transmit buffers in host memory. Descriptors that reside in the host memory act as pointers to these buffers.

There are two descriptor lists (one for receive and one for transmit) for the MAC DMA. The base address of each list is written in the TDLB register and in the RDLB register, respectively. A descriptor list is forward linked. The last descriptor may point back to the first entry to create a ring structure. Descriptors are chained by setting the next address to the next buffer in both receive and transmit descriptors.

The descriptor lists reside in the host physical memory address space. Each pointer points to one buffer and the second pointer points to the next descriptor. This enables the greatest flexibility for the host to chain any data buffers with discontinuous memory location. This eliminates processor-intensive tasks such as memory copying from the host to memory.

A data buffer contains either an entire frame or part of a frame, but it cannot exceed a single frame. Buffers contain only data; and buffer status is maintained in the descriptor. Data chaining refers to frames that span multiple data buffers. Data chaining can be enabled or disabled. Data buffers reside in host physical memory space.

Receive Descriptors (RDES0 - RDES3)

Receive descriptor and buffer addresses must be Word aligned. Each receive descriptor provides one frame buffer, one byte count field, and control and status bits.

TABLE 4-1: RDES0 REGISTER BIT FIELDS

Bit	Description
31	OWN Own Bit When set, indicates that the descriptor is owned by the KSZ8841-PMQL. When reset, indicates that the descriptor is owned by the host. The KSZ8841-PMQL clears this bit either when it completes the frame reception or when the buffers that are associated with this descriptor are full.
30	FS First Descriptor When set, indicates that this descriptor contains the first buffer of a frame. If the buffer size of the first buffer is 0, the next buffer contains the beginning of the frame.
29	LS Last Descriptor When set, indicates that the buffer pointed by this descriptor is the last buffer of the frame.
28	IPE IP Checksum Error When set, indicates that the received frame is an IP packet and its IP checksum field does not match. This bit is valid only when last descriptor is set.
27	TCPE TCP Checksum Error When set, indicates that the received frame is a TCP/IP packet and its TCP checksum field does not match. This bit is valid only when last descriptor is set.
26	UDPE UDP Checksum Error When set, indicates that the received frame is an UDP/IP packet and its UDP checksum field does not match. This bit is valid only when last descriptor is set.

KSZ8841-PMQL

TABLE 4-1: RDES0 REGISTER BIT FIELDS (CONTINUED)

Bit	Description
25	ES Error Summary Indicates the logical OR of the following RDES0 bits: CRC error Frame too long Runt frame This bit is valid only when last descriptor is set.
24	MF Multicast Frame When set, indicates that this frame has a multicast address. This bit is valid only when last descriptor is set.
23 - 20	SPN Switch Engine Source Port Number This field indicates the source port where the packet originated. If bit 20 is set, it indicates the packet was received from port 1. If bit 21 is set, it indicates the packet was received from port 2. This field is valid only when the last descriptor is set. (Bits 23 and 22 are not used, but reserved for backward compatibility and future expansion.)
19	RE Report on MII Error When set, indicates that a receive error in the physical layer was reported during the frame reception.
18	TL Frame Too Long When set, indicates that the frame length exceeds the maximum size of 1518 bytes. This bit is valid only when last descriptor is set. Note: Frame too long is only a frame length indication and does not cause any frame truncation.
17	RF Runt Frame When set, indicates that this frame was damaged by a collision or premature termination before the collision window has passed. Runt frames are passed on to the host only if the pass bad frame bit is set.
16	CE CRC Error When set, indicates that a CRC error occurred on the received frame. This bit is valid only when last descriptor is set.
15	FT Frame Type When set, indicates that the frame is an Ethernet-type frame (frame length field is greater than 1500 bytes). When clear, indicates that the frame is an IEEE 802.3 frame. This bit is not valid for runt frames. This bit is valid only when last descriptor is set.
14 - 11	Reserved
10 - 0	FL Frame Length Indicates the length, in bytes, of the received frame, including the CRC. This field is valid only when last descriptor is set and descriptor error is reset.

TABLE 4-2: RDES1 REGISTER BIT FIELDS

Bit	Description
31 -26	Reserved
25	RER Receive End of Ring When set, indicates that the descriptor list reached its final descriptor. The KSZ8841-PMQL returns to the base address of the list, thus creating a descriptor ring.
24 -12	Reserved
11 - 0	RBS Receive Buffer Size Indicates the size, in bytes, of the receive data buffer. If the field is 0, the KSZ8841-PMQL ignores this buffer and moves to the next descriptor. The buffer size must be a multiple of 4.

TABLE 4-3: RDES2 REGISTER BIT FIELDS

Bit	Description
31 - 0	Buffer Address Indicates the physical memory address of the buffer. The buffer address must be Word aligned.

TABLE 4-4: RDES3 REGISTER BIT FIELDS

Bit	Description
31 - 0	Next Descriptor Address Indicates the physical memory address of the next descriptor in the descriptor ring. The buffer address must be Word aligned.

Transmit Descriptors (TDES0-TDES3)

Transmit descriptors must be Word aligned. Each descriptor provides one frame buffer, one byte count field, and control and status bits.

TABLE 4-5: TDES0 REGISTER BIT FIELDS

Bit	Description
31	OWN Own Bit When set, indicates that the descriptor is owned by the KSZ8841-PMQL. When cleared, indicates that the descriptor is owned by the host. The KSZ8841-PMQL clears this bit either when it completes the frame transmission or when the buffer allocated in the descriptor is empty. The ownership bit of the first descriptor of the frame should be set after all subsequent descriptors belonging to the same frame have been set. This avoids a possible race condition between the KSZ8841-PMQL fetching a descriptor and the driver setting an ownership bit.
30 - 0	Reserved

TABLE 4-6: TDES1 REGISTER BIT FIELDS

Bit	Description
31	IC Interrupt on Completion When set, the KSZ8841-PMQL sets transmit interrupt after the present frame has been transmitted. It is valid only when last segment is set.
30	FS First Segment When set, indicates that the buffer contains the first segment of a frame.
29	LS Last Segment When set, indicates that the buffer contains the last segment of a frame.
28	IPCKG IP Checksum Generate When set, the KSZ8841-PMQL will generate correct IP checksum for outgoing frames that contains IP protocol header. The KSZ8841-PMQL supports only a standard IP header, i.e., IP with a 20 byte header. When this feature is used, ADD CRC bit in the transmit mode register should always be set. This bit is used as a per-packet control when the IP checksum generate bit in the transmit mode register is not set. This bit should be always set for multiple-segment packets.
27	TCPCKG TCP Checksum Generate When set, the KSZ8841-PMQL will generate correct TCP checksum for outgoing frames that contains IP and TCP protocol header. The KSZ8841-PMQL supports only a standard IP header, i.e., IP with a 20 byte header. When this feature is used, ADD CRC bit in the transmit mode register should always be set. This bit is used as a per-packet control when the TCP checksum generate bit in the transmit mode register is not set. This bit should be always set for multiple-segment packets.

KSZ8841-PMQL

TABLE 4-6: TDES1 REGISTER BIT FIELDS (CONTINUED)

Bit	Description
26	UDPCKG UDP Checksum Generate When set, the KSZ8841-PMQL will generate correct UDP checksum for outgoing frames that contains an IP and UDP protocol header. The KSZ8841-PMQL supports only a standard IP header, i.e., IP with a 20 byte header. When this feature is used, ADD CRC bit in the transmit mode register should always be set. This bit is used as a per-packet control when the UDP checksum generate bit in the transmit mode register is not set.
25	TER Transmit End of Ring When set, indicates that the descriptor pointer has reached its final descriptor. The KSZ8841-PMQL returns to the base address of the list, forming a descriptor ring.
24	Reserved
23 - 20	SPN Switch Engine Destination Port Map When set, this field indicates the destination port(s) where the packet will be forwarded to. If bit 20 is set, it indicates the packet was received from port 1. If bit 21 is set, it indicates the packet was received from port 2. Setting all ports to 1 will cause the controller engine to broadcast the packet. Setting all bits to 0 has no effect. The controller engine forwards the packet according to its internal controller lookup algorithm. This field is valid only when the last descriptor is set. (Bits 23 and 22 are not used, but reserved for backward compatibility and future expansion.)
19 - 11	Reserved
10 - 0	TBS Transmit Buffer Size Indicates the size, in bytes, of the transmit data buffer. If this field is 0, the KSZ8841-PMQL ignores this buffer and moves to the next descriptor.

TABLE 4-7: TDES2 REGISTER BIT FIELDS

Bit	Description
31 - 0	Buffer Address Indicates the physical memory address of the buffer. There is no limitation on the transmit buffer address alignment.

TABLE 4-8: TDES3 REGISTER BIT FIELDS

Bit	Description
31 - 0	Next Descriptor Address Indicates the physical memory address of the next descriptor in the descriptor ring. The buffer address must be Word aligned.

4.2 PCI Configuration Registers

The KSZ8841-PMQL implements 12 configuration registers. These registers are described in the following subsections.

The KSZ8841-PMQL enables a full software-driven initialization and configuration. This allows the software to identify and query the KSZ8841-PMQL. The KSZ8841-PMQL treats configuration space write operations to registers that are reserved as no-ops. That is, the access completes normally on the bus and the data is discarded. Read accesses, to reserved or unimplemented registers, complete normally and a data value of '0' is returned.

Software reset has no effect on the configuration registers. Hardware reset sets the configuration registers to their default values.

TABLE 4-9: LIST OF CONFIGURATION REGISTERS

Configuration Register	Identifier	I/O Address Offset	Default
Identification	CFID	0x00	0x884116C6
Command and Status	CFCS	0x04	0x02*00000
Revision	CFRV	0x08	0x02000010
Latency Timer	CFLT	0x0C	0x00000000
Base Memory Address	CBMA	0x10	0x00000000
Reserved	—	0x14 - 0x28	0x00000000
Subsystem ID	CSID	0x2C	0x*****
Capabilities Pointer	CCAP	0x34	0x*****
Reserved	—	0x38	0x00000000
Interrupt	CFIT	0x3C	0x28140100
Reserved	—	0x40 - 0x4C	0x00000000
Capability ID	CCID	0x50	0x***20001
Power Management Control and Status	CPMC	0x54	0x00000000

Configuration ID Register (CFID Offset 00H)

The CFID register identifies the KSZ8841-PMQL. The following table shows the CFID register bit fields.

TABLE 4-10: CONFIGURATION ID REGISTER (CFID OFFSET 00H)

Bit	Default	Description
31 - 16	0x8841	Device ID
15 - 0	0x16C6	Vendor ID Specifies the manufacturer of the KSZ8841-PMQL.

The following table shows the access rules of the register.

TABLE 4-11: REGISTER ACCESS RULES

Category	Description
Value after hardware reset	0x884116C6
Write access rules	Write has no effect on the KSZ8841-PMQL.

Command and Status Configuration Register (CFCS Offset 04H)

The CFCS register is divided into two sections: a command register (CFCS[15:0]) and a status register (CFCS[31:16]). The command register provides control of the KSZ8841-PMQL's ability to generate and respond to PCI cycles. When '0' is written to this register, the KSZ8841-PMQL logically disconnects from the PCI bus for all accesses except configuration accesses.

The status register records status information for the PCI bus-related events. The CFCS status bits are not cleared when they are read. Writing '1' to these bits clears them; writing '0' has no effect.

The following table describes the CFCS register bit fields.

KSZ8841-PMQL

TABLE 4-12: COMMAND AND STATUS CONFIGURATION REGISTER (CFCS OFFSET 04H)

Bit	Type	Default	Description
31	Status	0	Detected Parity Error When set, indicates that the KSZ8841-PMQL detected a parity error, even if parity error handling is disabled in parity error response (CFCS[6]).
30	Status	0	Signal System Error When set, indicates that the KSZ8841-PMQL asserted the system error SERR_N pin.
29	Status	0	Received Master Abort When set, indicates that the KSZ8841-PMQL terminated a master transaction with master abort.
28	Status	0	Received Target Abort When set, indicates that the KSZ8841-PMQL master transaction was terminated due to a target abort.
27	Status	0	Target Abort This bit is set by KSZ8841-PMQL whenever it terminates with a Target Abort. The CSR registers are all 32-bit Little Endian format. For PCI register Read cycles, the KSZ8841-PMQL allows any different combination of CBEN. For PCI register bus cycles, only byte, word (16-bit), or Dword (32-bit) accesses are allowed. Any other combination is illegal and is target aborted.
26 - 25	Status	01	Device Select Timing Indicates the timing of the assertion of device select (DEVSEL_N). These bits are fixed at 01, which indicates a medium assertion of DEVSEL_N.
24	Status	0	Data Parity Report This bit is set when the following conditions are met: The KSZ8841-PMQL asserts parity error PERR_N or it senses the assertion of PERR_N by another device. The KSZ8841-PMQL operates as a bus master for the operation that caused the error. Parity error response (CFCS[6]) is set.
23 - 22	Reserved	00	Reserved
21	Status	0	66 MHz Capable 0 = Not 66 MHz capable
20	Status	—	New Capability Indicates whether or not the KSZ8841-PMQL implements a list of new capabilities. When set, this bit indicates the presence of New capabilities. When reset, New capabilities are not implemented. The value of this bit is loaded from the New_Cap bit in EEPROM.
19 - 9	Reserved	0x000	Reserved
8	Command	0	System Error Enable When set, the KSZ8841-PMQL asserts system error (SERR_N) when it detects a parity error on the address phase.
7	Reserved	0	Reserved
6	Command	0	Parity Error Response When set, the KSZ8841-PMQL asserts fatal bus error after it detects a parity error. When reset, any detected parity error is ignored and the KSZ8841-PMQL continues normal operation. Parity checking is disabled after hardware reset.
5 - 3	Reserved	000	Reserved

TABLE 4-12: COMMAND AND STATUS CONFIGURATION REGISTER (CFCS OFFSET 04H)

Bit	Type	Default	Description
2	Command	0	Master Operation When set, the KSZ8841-PMQL is capable of acting as a bus master. When reset, the KSZ8841-PMQL capability to generate PCI accesses is disabled. For normal operation, this bit must be set.
1	Command	0	Memory Space Access When set, the KSZ8841-PMQL responds to memory space accesses. When reset, the KSZ8841-PMQL does not respond to memory space accesses.
0	Reserved	0	Reserved

Configuration Revision Register (CFRV Offset 08H)

The CFRV register contains the KSZ8841-PMQL revision number. The following table shows the CFRV register bit fields.

TABLE 4-13: CONFIGURATION REVISION REGISTER (CFRV OFFSET 08H)

Bit	Default	Description
31 - 24	0x02	Base Class Indicates the network controller and is equal to 0x2.
23 - 16	0x00	Subclass Indicates the Fast/Gigabit Ethernet chip and is equal to 0x00.
15 - 8	0x00	Reserved
7 - 4	0x1	Revision Number Indicates the KSZ8841-PMQL revision number, and is equal to 0x1. This number is incremented for subsequent revision.
3 - 0	0x0	Step Number Indicates the KSZ8841-PMQL step number, and is equal to 0x0 (chip revision A). This number is incremented for subsequent KSZ8841-PMQL steps within the current revision.

Configuration Latency Timer Register (CFLT Offset 0CH)

This register configures the cache line size field and the latency timer.

The following table shows the CFLT register bit fields.

TABLE 4-14: CONFIGURATION LATENCY TIMER REGISTER (CFLT OFFSET 0CH)

Bit	Default	Description
31 - 16	0x00	Reserved
15 - 8	0x00	Configuration Latency Timer Specifies, in units of PCI bus clocks, the value of the latency timer of the KSZ8841-PMQL. When the KSZ8841-PMQL asserts FRAME_N, it enables its latency timer to count. If the KSZ8841-PMQL deserts FRAME_N prior to count expiration, the content of the latency timer is ignored. Otherwise, after the count expires, the KSZ8841-PMQL initiates transaction termination as soon as its GNT_N is deserted.
7 - 0	0x00	Cache Line Size Specifies, in unit of 32-bit words (Dword), the system cache line size.

KSZ8841-PMQL

Configuration Base Memory Address Register (CBMA Offset 10H)

The CBMA register specifies the base memory address for accessing the KSZ8841-PMQL CSRs. This register must be initialized prior to accessing any CSR with memory access.

The following table shows the CBMA register bit fields.

TABLE 4-15: CONFIGURATION BASE MEMORY ADDRESS REGISTER (CBMA OFFSET 10H)

Bit	Default	Description
31 - 11	0	Configuration Base Memory Address Defines the base address assigned for mapping the KSZ8841-PMQL CSRs.
10 - 1	0	This field value is 0 when read.
0	0	Memory Space Indicator Determines that the register maps into the Memory space. The value in this field is 0. This is a read-only field.

Subsystem ID Register (CSID Offset 2CH)

The CSID register is a read-only 32-bit register. The content of the CSID is loaded from the EEPROM after hardware reset. The loading period lasts at least 27,400 PCI cycles when the system is in 33 MHz mode, and starts 50 cycles after hardware reset desertion. If the host accesses the CSID before its content is loaded from the EEPROM, the KSZ8841-PMQL responds with retry termination on the PCI bus.

The following table shows the CSID register bit fields.

TABLE 4-16: SUBSYSTEM ID REGISTER (CSID OFFSET 2CH)

Bit	Description
31 - 16	Subsystem ID Indicates a 16-bit field containing the subsystem ID.
15 - 0	Subsystem Vendor ID Indicates a 16-bit field containing the subsystem vendor ID.

The following table shows the access rules of the register.

TABLE 4-17: REGISTER ACCESS RULES

Category	Description
Value after hardware reset	Read from EEPROM.
Write access rules	Write has no effect on the KSZ8841-PMQL.

Capabilities Pointer Register (CCAP Offset 34H)

The CCAP register points to the base address of the power management register block in the configuration address space. This pointer is valid only if the new capability bit in CFCS is set.

The following table shows the CCAP register bit fields.

TABLE 4-18: CAPABILITIES POINTER REGISTER (CCAP OFFSET 34H)

Bit	Default	Description
31 - 8	0x000000	Reserved
7 - 0	—	Capabilities Pointer Points to the location of the power management register block in the PCI configuration space. The value of this field is determined by the New Capabilities bit 15 in the EEPROM. If this bit is set, the value of this field is 0x50, which stands for Support Power Management. Otherwise, this field is read as 0x00.

Configuration Interrupt Register (CFIT Offset 3CH)

The CFIT register is divided into two sections: the interrupt line and the interrupt pin. CFIT configures both the system's interrupt and the KSZ8841-PMQL interrupt pin connection.

The following table shows the CFIT register bit fields.

TABLE 4-19: CONFIGURATION INTERRUPT REGISTER (CFIT OFFSET 3CH)

Bit	Default	Description
31 - 24	0x28	MAX_LAT This field indicates how often the device needs to gain access to the PCI bus. Time unit is equal to 0.25 μ s, assuming a PCI clock frequency of 33 MHz. The value after a hardware reset is 0x28 (10 μ s).
23 - 16	0x14	MIN_GNT This field indicates the burst period length that the device needs. Time unit is equal to 0.25 μ s, assuming a PCI clock frequency of 33 MHz. The value after a hardware reset is 0x14 (5 μ s).
15 - 8	0x01	Interrupt Pin Indicates which interrupt pin that the KSZ8841-PMQL uses. The KSZ8841-PMQL uses INTA# and the read value is 0x01.
7 - 0	0x00	Interrupt Line Provides interrupt line routing information. The basic input/output system (BIOS) writes the routing information into this field when it initialized and configures the system. The value in this field indicates which input of the system interrupt controller is connected to the KSZ8841-PMQL's interrupt pin. The driver can use this information to determine priority and vector information. Values in this field are system architecture specific.

The following table shows the access rules of the register.

TABLE 4-20: REGISTER ACCESS RULES

Category	Description
Value after hardware reset	0x281401XX

Capabilities ID Register (CCID Offset 50H)

The CCID register is a read-only register that provides information on the KSZ8841-PMQL power management capabilities. The following table shows the CCID register bit fields. The CCID register bits [31:16] are mirrored with PMCR register bits [15:0]. References to D0, D1, D2, and D3 are power management states defined in a similar fashion to the way they are defined for PCI.

TABLE 4-21: CAPABILITIES ID REGISTER (CCID OFFSET 50H)

Bit	Default	Description
31	0	PME Support D3 (cold) If this bit is set, the KSZ8841-PMQL asserts PME in D3 (cold) power state. Otherwise, the KSZ8841-PMQL does not assert PME in D3(cold). The value of this bit is loaded from the PME_D3_cold bit in the EEPROM.
30	1	PME Support D3 (hot) The value of this bit is 1, indicating that the KSZ8841-PMQL may assert PME in D3 (hot) power state.
29	0	PME Support D2 If this bit is set, the KSZ8841-PMQL asserts PME in D2 power state. Otherwise, the KSZ8841-PMQL does not assert PME in D2 state. The value of this bit is loaded from the PME_D2 bit in the EEPROM.
28	0	PME Support D1 If this bit is set, the KSZ8841-PMQL asserts PME in D1 power state. Otherwise, the KSZ8841-PMQL does not assert PME in D1 state. The value of this bit loaded from the PME_D1 bit in the EEPROM.

KSZ8841-PMQL

TABLE 4-21: CAPABILITIES ID REGISTER (CCID OFFSET 50H) (CONTINUED)

Bit	Default	Description
27	0	PME Support D0 The value of this bit is 0, indicating that the KSZ8841-PMQL does not assert PME in D0 power state.
26	0	D2 Support If this bit is set, it indicates that the KSZ8841-PMQL support D2 power state. The value of this bit is loaded from the D2_SUP bit in the EEPROM.
25	0	D1 Support If this bit is set, it indicates that the KSZ8841-PMQL support D1 power state. The value of this bit loaded from the D1_SUP bit in the EEPROM.
24 - 22	000	Auxiliary Current This 3-bit field reports the 3.3V _{AUX} auxiliary current requirements for the PCI function. If PME# generation from D3_cold is not supported by the function, this field must return a value of 000 when read.
21	0	Device Specific Initialization Indicates whether special initialization of this function is required (beyond the standard PCI configuration header) before the generic class device driver is able to use it. Note that this bit is not used by some operating systems. Microsoft Windows and Windows NT, for instance, do not use this bit to determine whether to use D3. Instead, they use the driver's capabilities to determine this. A "1" indicates that the function requires a device specific initialization sequence following transition to the D0 uninitialization state. The value of this bit is loaded from the PME_DSI bit in the EEPROM.
20	0	Reserved Should be set to 0.
19	0	PME Clock When this bit is a "1", it indicates that the function relies on the presence of the PCI clock for PME# operation. When this bit is a "0", it indicates that no PCI clock is required for the function to generate PME#. The value of this bit is loaded from the PME_CK bit in the EEPROM.
18 - 16	0	Power Management PCI Version The value of this bit is loaded from the PME_VER[2:0] bits in the EEPROM.
15 - 8	0x00	Next Item Pointer Points to the location of the next block of the capabilities list in the PCI Configuration Space. The value of this field is 0x00, indicating that this is the last item of the Capability linked list.
7 - 0	0x01	Capabilities ID PCI Power Management Registers ID. The value of this field is 01h, indicating that this is the power-management register block.

The following table shows the access rules of the register.

TABLE 4-22: REGISTER ACCESS RULES

Category	Description
Value after hardware reset	0x40000001 & EEPROM
Write access rules	Write has no effect on the KSZ8841-PMQL

Power Management Control and Status Register (CPMC Offset 54H)

The CMPC register is a power management control and status register. This register can control and sate power management events. The following table shows the CMPC register bit fields.

TABLE 4-23: POWER MANAGEMENT CONTROL AND STATUS REGISTER (CPMC OFFSET 54H)

Bit	Default	Description
31 - 16	0x0000	Reserved
15	0	PME_Status This bit indicates that the KSZ8841-PMQL has detected a power management event. If bit PME_Enable is set, the KSZ8841-PMQL also asserts the PME_N pin. This bit is cleared on power-up reset or by write 1. It is not modified by either hardware or software reset. When this bit is cleared, the KSZ8841-PMQL deserts the PME_N pin.
14 - 9	0x00	Reserved
8	0	PME_Enable If this bit is set, the KSZ8841-PMQL can assert the PME_N pin. Otherwise, assertion of the PME_N pin is disabled. This bit is cleared on power-up reset only and is not modified by either hardware or software reset.
7 - 4	0x0	Reserved
3	0	No Soft Reset If this bit is set ("1"), the KSZ8841-PMQL does not perform an internal reset when transitioning from D3_hot to D0 because of Power State commands. Configuration context is preserved. Upon transition from D3_hot to the D0 Initialized state, no additional operating system intervention is required to preserve configuration context beyond writing the Power State bits. If this bit is cleared ("0"), the KSZ8841-PMQL does perform an internal reset when transitioning from D3_hot to D0 via software control of the Power State bits. Configuration context is lost when performing the soft reset. Upon transition from D3_hot to the D0 state, full reinitialization sequence is needed to return the device to D0 Initialized. Regardless of this bit, devices that transition from D3_hot to D0 by a system or bus segment reset will return to the device state D0 Uninitialized with only PME context preserved if PME is supported and enabled.
2	0	Reserved
1 - 0	00	Power State This field is used to set the current power state of the KSZ8841-PMQL and to determine its power state. The definitions of the field values are: 0 = D0 1 = D1 2 = D2 3 = D3 (hot) This field gets a value of 0 after power up.

The following table shows the access rules of the register.

TABLE 4-24: REGISTER ACCESS RULES

Category	Description
Bit 15	Read/Write 1 Clear (RW1C)
Bit 8	Read/Write (RW)
Bit 3	Read Only (RO)
Bit1:0	Read Write (RW)

4.3 PCI Control & Status Registers

The PCI CSR registers are all 32-bit in Little Endian format. For PCI register Read cycle, the KSZ8841-PMQL allows any different combination of CBEN. For PCI register bus cycles, only byte, word (16-bit), or Dword (32-bit) accesses are allowed. Any other combinations are illegal and will be target aborted.

All other registers not included below are reserved.

MAC DMA Transmit Control Register (MDTXC Offset 0x0000)

The MAC DMA transmit control register establishes the transmit operating modes and commands for the port. This register should be one of the last CSRs to be written as part of the transmit initialization.

The following table shows the register bit fields.

TABLE 4-25: MAC DMA TRANSMIT CONTROL REGISTER (MDTXC OFFSET 0X0000)

Bit	Default	R/W	Description
31 - 30	—	RO	Reserved
29 - 24	0x00	R/W	MTBS DMA Transmit Burst Size This field indicates the maximum number of words to be transferred in one DMA transaction. If reset, the MAC DMA burst size is limited only by the amount of data stored in the transmit buffer before issuing a bus request. The MTBS can be programmed with permissible values 0, 1, 2, 4, 8, 16, or 32. After reset, the MTBS default is 0, i.e. unlimited.
23 - 19	0x00	RO	Reserved
18	0	R/W	MTUCG MAC Transmit UDP Checksum Generate When set, the KSZ8841-PMQL will generate correct UDP checksum for outgoing UDP/IP frames at port. When this bit is set, ADD CRC should also turn on.
17	0	R/W	MTTCG MAC Transmit TCP Checksum Generate When set, the KSZ8841-PMQL will generate correct TCP checksum for outgoing TCP/IP frames at port. When this bit is set, ADD CRC should also turn on.
16	0	R/W	MTICG MAC Transmit IP Checksum Generate When set, the KSZ8841-PMQL will generate correct IP checksum for outgoing IP frames at port. When this bit is set, ADD CRC should also turn on.
15 - 10	0x00	RO	Reserved
9	0	R/W	MTFCE MAC Transmit Flow Control Enable When this bit is set and the KSZ8841-PMQL is in Full-Duplex mode, flow control is enabled and the KSZ8841-PMQL will transmit a PAUSE frame when the Receive Buffer capacity has reached a level that may cause the buffer to overflow. When this bit is set and the KSZ8841-PMQL is in Half-Duplex mode, back-pressure flow control is enabled. When this bit is cleared, no transmit flow control is enabled.
8 - 3	0x0	RO	Reserved
2	0	R/W	MTEP MAC DMA Transmit Enable Padding When set, the KSZ8841-PMQL automatically adds a padding field to a packet shorter than 64 bytes. Note: Setting this bit automatically enables Add CRC feature.
1	0	R/W	MTAC MAC DMA Transmit Add CRC When set, the KSZ8841-PMQL appends the CRC to the end of the transmission frame.

TABLE 4-25: MAC DMA TRANSMIT CONTROL REGISTER (MDTXC OFFSET 0X0000)

Bit	Default	R/W	Description
0	0	R/W	MTE MAC DMA TX Enable When the bit is set, the MDMA TX block is enabled and placed in a running state. When reset, the transmission process is placed in the stopped state after completing the transmission of the current frame. The stop transmission command is effective only when the transmission process is in the running state.

MAC DMA Receive Control Register (MDRXC Offset 0x0004)

The MAC DMA receive control register establishes the receive operating modes and commands for the port. This register should be one of the last CSRs to be written as part of the receive initialization.

The following table shows the register bit fields.

TABLE 4-26: MAC DMA RECEIVE CONTROL REGISTER (MDRXC OFFSET 0X0004)

Bit	Default	R/W	Description
31 - 30	00	RO	Reserved
29 - 24	0x00	R/W	MRBS DMA Receive Burst Size This field indicates the maximum number of words to be transferred in one DMA transaction. If reset, the MAC DMA burst size is limited only by the amount of data stored in the receive buffer before issuing a bus request. The MRBS can be programmed with permissible values 0, 1, 2, 4, 8, 16, or 32. After reset, the MRBS default is 0, i.e. unlimited.
23 - 20	0x0	RO	Reserved
19	0	R/W	IP Header Alignment Enable 1 = Enable alignment of IP header to dWord address. Layer 2 header will not be dWord aligned anymore. Please look at RX descriptor 0 for the Layer 2 header address shift. 0 = IP Header alignment disabled.
18	0	R/W	MRUCC MAC Receive UDP Checksum Check When set, the KSZ8841-PMQL will check for correct UDP checksum for incoming UDP/IP frames at port. Packets received with incorrect UDP checksum will be discarded.
17	0	R/W	MRTCG MAC Receive TCP Checksum Check When set, the KSZ8841-PMQL will check for correct TCP checksum for incoming TCP/IP frames at port. Packets received with incorrect TCP checksum will be discarded.
16	0	R/W	MRICG MAC Receive IP Checksum Check When set, the KSZ8841-PMQL will check for correct IP checksum for incoming IP frames at port. Packets received with incorrect IP checksum will be discarded.
15 - 10	0x00	RO	Reserved
9	0	R/W	MRFCE MAC Receive Flow Control Enable When this bit is set and the KSZ8841-PMQL is in Full-Duplex mode, flow control is enabled and the KSZ8841-PMQL will acknowledge a PAUSE frame from MAC of the controller, the outgoing packets will be pending in the transmit buffer until the PAUSE control timer expires. This field has no meaning in half-duplex mode and should be programmed to 0. When this bit is cleared, no flow control is enabled.
8 - 7	00	RO	Reserved
6	0	R/W	MRB MAC Receive Broadcast When set, the MAC receive all broadcast frames.

TABLE 4-26: MAC DMA RECEIVE CONTROL REGISTER (MDRXC OFFSET 0X0004) (CONTINUED)

Bit	Default	R/W	Description
5	0	R/W	MRM MAC Receive Multicast When set, the MAC receive all multicast frames (including broadcast).
4	0	R/W	MRU MAC Receive Unicast When set, the MAC receive unicast frames that match the 48-bit Station Address of the MAC.
3	0	R/W	MRE MAC DMA Receive Error Frame When set, the KSZ8841-PMQL will pass the errors frames received to the host. Error frames include runt frames, oversized frames, CRC errors.
2	0	R/W	MRA MAC DMA Receive All When set, the KSZ8841-PMQL receives all incoming frames, regardless of its destination address.
1	0	R/W	DMA Receive Multicast Hash-Table Enable Setting this bit enables the RX function to receive multicast frames that pass the CRC Hash filtering mechanism.
0	0	R/W	MRE MAC DMA RX Enable When the bit is set, the DMA RX block is enabled and placed in a running state. When reset, the receive process is placed in the stopped state after completing the reception of the current frame. The stop transmission command is effective only when the reception process is in the running state.

MAC DMA Transmit Start Command Register (MDTSC Offset 0x0008)

This register is written by the CPU when packets in the data buffer need to be transmitted. The following table shows the register bit fields.

TABLE 4-27: MAC DMA TRANSMIT START COMMAND REGISTER (MDTSC OFFSET 0X0008)

Bit	Default	R/W	Description
31 - 0	0x00000000	WO	WTSC Transmit Start Command When written with any value, the Transmit DMA checks for frames to be transmitted. If no descriptor is available, the transmit process returns to suspended state. If descriptors are available, the transmit process starts or resumes. This bit is self-clearing.

MAC DMA Receive Start Command Register (MDRSC Offset 0x000C)

This register is written by the CPU when there are frame data in receive buffer to be processed.

The following table shows the register bit fields.

TABLE 4-28: MAC DMA RECEIVE START COMMAND REGISTER (MDRSC OFFSET 0X000C)

Bit	Default	R/W	Description
31 - 0	0x00000000	WO	WRSC Receive Start Command When written with any value, the Receive DMA checks for descriptors to be acquired. If no descriptor is available, the receive process returns to suspended state and wait for the next receive restart command. If descriptors are available, the receive process resumes. This bit is self-clearing.

Transmit Descriptor List Base Address Register (TDLB Offset 0x0010)

This register is used for Transmit descriptor list base address register. The register is used to point to the start of the appropriate descriptor list. Writing to this register is permitted only when its respective process is in the stopped state. When stopped, the register must be written before the respective START command is given.

Note that the descriptor lists must be Word (32-bit) aligned. The KSZ8841-PMQL behavior is unpredictable when the lists are not word-aligned.

The following table shows the register bit fields.

TABLE 4-29: TRANSMIT DESCRIPTOR LIST BASE ADDRESS REGISTER (TDLB OFFSET 0X0010)

Bit	Default	R/W	Description
31 - 0	0x00000000	R/W	WSTL Start of Transmit List Note: Write can only occur when the transmit process stopped.

Receive Descriptor List Base Address Register (RDLB Offset 0x0014)

This register is used for Receive descriptor list base address register. The register is used to point to the start of the appropriate descriptor list. Writing to this register is permitted only when its respective process is in the stopped state. When stopped, the register must be written before the respective START command is given.

Note that the descriptor lists must be Word (32-bit) aligned. The KSZ8841-PMQL behavior is unpredictable when the lists are not word-aligned.

The following table shows the register bit fields.

TABLE 4-30: RECEIVE DESCRIPTOR LIST BASE ADDRESS REGISTER (RDLB OFFSET 0X0014)

Bit	Default	R/W	Description
31 - 0	0x00000000	R/W	WSRL Start of Receive List Note: Write can only occur when the transmit process stopped.

MAC Multicast Table 0 Register (MTR0 Offset 0x0020)

The 64-bit multicast table is used for group address filtering. The value is defined as the six most significant bits of the CRC of the DA. The two most significant bits select the register to be used, while the other determines the bit within the register.

TABLE 4-31: MAC MULTICAST TABLE 0 REGISTER (MTR0 OFFSET 0X0020)

Bit	Default	R/W	Description
31 - 0	0x00000000	R/W	MTR0 Multicast Table 0 When appropriate bit is set, the packet received with DA matches the CRC hashing function is received without being filtered. Note: when receive all (RXRA) or receive multicast (RXRM) bit is set in the RXCR then all multicast addresses are received regardless of the multicast table value.

MAC Multicast Table 1 Register (MTR1 Offset 0x0024)

The 64-bit multicast table is used for group address filtering. The value is defined as the six most significant bits of the CRC of the DA. The two most significant bits select the register to be used, while the other determines the bit within the register.

TABLE 4-32: MAC MULTICAST TABLE 1 REGISTER (MTR1 OFFSET 0X0024)

Bit	Default	R/W	Description
31 - 0	0x00000000	R/W	MTR0 Multicast Table 1 When appropriate bit is set, the packet received with DA matches the CRC hashing function is received without being filtered. Note: When receive all (RXRA) or receive multicast (RXRM) bit is set in the RXCR then all multicast addresses are received regardless of the multicast table value.

Interrupt Enable Register (INTEN Offset 0x0028)

This register enables the interrupts from the internal or external sources.

The following table shows the register bit fields.

TABLE 4-33: INTERRUPT ENABLE REGISTER (INTEN OFFSET 0X0028)

Bit	Default	R/W	Description
31	0	R/W	DMLCIE DMA MAC Link Changed Interrupt Enable When this bit is set, the DMA MAC Link Changed Interrupt is enabled. When this bit is reset, the DMA MAC Link Changed Interrupt is disabled.
30	0	R/W	DMTIE DMA MAC Transmit Interrupt Enable When this bit is set, the DMA MAC Transmit Interrupt is enabled. When this bit is reset, the DMA MAC Transmit Interrupt is disabled.
29	0	R/W	DMRIE DMA MAC Receive Interrupt Enable When this bit is set, the DMA MAC Receive Interrupt is enabled. When this bit is reset, the DMA MAC Receive Interrupt is disabled.
28	0	R/W	DMTBUIE DMA MAC Transmit Buffer Unavailable Interrupt Enable When this bit is set, the DMA MAC Transmit Buffer Unavailable Interrupt is enabled. When this bit is reset, the DMA MAC Transmit Buffer Unavailable Interrupt is disabled.
27	0	R/W	DMRBUIE DMA MAC Receive Buffer Unavailable Interrupt Enable When this bit is set, the DMA MAC Receive Buffer Unavailable Interrupt is enabled. When this bit is reset, the DMA MAC Receive Buffer Unavailable Interrupt is disabled.
26	0	R/W	DMTPSIE DMA MAC Transmit Process Stopped Interrupt Enable When this bit is set, the DMA MAC Transmit Process Stopped Interrupt is enabled. When this bit is reset, the DMA MAC Transmit Process Stopped Interrupt is disabled.
25	0	R/W	DMRPSIE DMA MAC Receive Process Stopped Interrupt Enable When this bit is set, the DMA MAC Receive Process Stopped Interrupt is enabled. When this bit is reset, the DMA MAC Receive Process Stopped Interrupt is disabled.
24 - 0	—	RO	Reserved

Interrupt Status Register (INTST Offset 0x002C)

This register contains all the status bits for the ARM CPU. When corresponding enable bit is set, it causes the CPU to be interrupted. This register is usually read by the driver during interrupt service routine or polling. The register bits are not cleared when read. Each field can be masked.

The following table shows the register bit fields.

TABLE 4-34: INTERRUPT STATUS REGISTER (INTST OFFSET 0X002C)

Bit	Default	R/W	Description
31	0	R/W	DMLCS DMA MAC Link Changed Status When this bit is set, it indicates that the DMA MAC link status has changed from link up to link down or from link down to link up. This edge-triggered interrupt status is cleared by writing 1 to this bit.
30	0	R/W	DMTS DMA MAC Transmit Status When this bit is set, it indicates that the DMA MAC has transmitted at least a frame on the DMA port and the MAC is ready for new frames from the host. This edge-triggered interrupt status is cleared by writing 1 to this bit.

TABLE 4-34: INTERRUPT STATUS REGISTER (INTST OFFSET 0X002C) (CONTINUED)

Bit	Default	R/W	Description
29	0	R/W	DMRS DMA MAC Receive Status When this bit is set, it indicates that the DMA MAC has received a frame from the DMA port and it is ready for the host to process This edge-triggered interrupt status is cleared by writing 1 to this bit.
28	0	R/W	DMTBUS DMA MAC Transmit Buffer Unavailable Status When this bit is set, it indicates that the next descriptor on the transmit list is owned by the host and cannot be acquired by the KSZ8841-PMQL. The transmission process is suspended. To resume processing transmit descriptors, the host should change the ownership bit of the descriptor and then issue a transmit start command. This edge-triggered interrupt status is cleared by writing 1 to this bit.
27	0	R/W	DMRBUS DMA MAC Receive Buffer Unavailable Status When this bit is set, it indicates that the descriptor list is owned by the host and cannot be acquired by the KSZ8841-PMQL. The receiving process is suspended. To resume processing receive descriptors, the host should change the ownership of the descriptor and may issue a receive start command. If no receive start command is issued, the receiving process resumes when the next recognized incoming frame is received. After the first assertion, this bit is not asserted for any subsequent not owned receive descriptors fetches. This bit is asserted only when the previous receive descriptor was owned by the KSZ8841-PMQL. This edge-triggered interrupt status is cleared by writing 1 to this bit.
26	0	R/W	DMPSS DMA MAC Transmit Process Stopped Status Asserted when the DMA MAC transmit process enters the stopped state. This edge-triggered interrupt status is cleared by writing 1 to this bit.
25	0	R/W	DMRPSS DMA MAC Receive Process Stopped Status Asserted when the DMA MAC receive process enters the stopped state. This edge-triggered interrupt status is cleared by writing 1 to this bit.
24 - 0	—	RO	Reserved

MAC Additional Station Address Low Register (MAAL0-15)

The KSZ8841-PMQL supports 16 additional MAC addresses for MAC address filtering. This MAC address is used to define one of the 16 destination addresses that the KSZ8841-PMQL will respond to when receiving frames on the port. Network addresses are generally expressed in the form of 01:23:45:67:89:AB, where the bytes are received left to right, and the bits within each byte are received right to left (LSB to MSB). The actual transmitted and received bits are in the order of 10000000 11000100 10100010 11100110 10010001 11010101.

The following table shows the register bit fields.

TABLE 4-35: MAC ADDITIONAL STATION ADDRESS LOW REGISTER (MAAL0-15)

Bit	Default	R/W	Description
31 - 0	—	R/W	MAAL0 MAC Additional Station Address 0 Low 4 bytes The least significant word of the additional MAC 0 station address.

KSZ8841-PMQL

MAC Additional Station Address High Register (MAAH0-15)

The KSZ8841-PMQL supports 16 additional MAC addresses for MAC address filtering. This MAC address is used to define one of the 16 destination addresses that the KSZ8841-PMQL will respond to when receiving frames on the port. Network addresses are generally expressed in the form of 01:23:45:67:89:AB, where the bytes are received left to right, and the bits within each byte are received right to left (LSB to MSB). The actual transmitted and received bits are in the order of 10000000 11000100 10100010 11100110 10010001 11010101.

The following table shows the register bit fields.

TABLE 4-36: MAC ADDITIONAL STATION ADDRESS HIGH REGISTER (MAAH0-15)

Bit	Default	R/W	Description
31	0	R/W	MAA0E MAC Additional Station Address 0 Enable When set, the additional MAC address is enabled for received frames. When reset, the additional MAC address is disabled.
30 - 16	0x0	RO	Reserved
15 - 0	—	R/W	MAAH0 MAC Additional Station Address 0 High 2 bytes The most significant word of the additional MAC 0 station address.

The following table shows the register map for all 16 additional MAC address registers.

TABLE 4-37: REGISTER MAP FOR ALL 16 MAC ADDRESS REGISTERS

Register	Identifier	Offset
ADD MAC Low 0	MAAL0	0x0080
ADD MAC High 0	MAAH0	0x0084
ADD MAC Low 1	MAAL1	0x0088
ADD MAC High 1	MAAH1	0x008C
ADD MAC Low 2	MAAL2	0x0090
ADD MAC High 2	MAAH2	0x0094
ADD MAC Low 3	MAAL3	0x0098
ADD MAC High 3	MAAH3	0x009C
ADD MAC Low 4	MAAL4	0x00A0
ADD MAC High 4	MAAH4	0x00A4
ADD MAC Low 5	MAAL5	0x00A8
ADD MAC High 5	MAAH5	0x00AC
ADD MAC Low 6	MAAL6	0x00B0
ADD MAC High 6	MAAH6	0x00B4
ADD MAC Low 7	MAAL7	0x00B8
ADD MAC High 7	MAAH7	0x00BC
ADD MAC Low 8	MAAL8	0x00C0
ADD MAC High 8	MAAH8	0x00C4
ADD MAC Low 9	MAAL9	0x00C8
ADD MAC High 9	MAAH9	0x00CC
ADD MAC Low 10	MAAL10	0x00D0
ADD MAC High 10	MAAH10	0x00D4
ADD MAC Low 11	MAAL11	0x00D8
ADD MAC High 11	MAAH11	0x00DC
ADD MAC Low 12	MAAL12	0x00E0
ADD MAC High 12	MAAH12	0x00E4
ADD MAC Low 13	MAAL13	0x00E8
ADD MAC High 13	MAAH13	0x00EC
ADD MAC Low 14	MAAL14	0x00F0

TABLE 4-37: REGISTER MAP FOR ALL 16 MAC ADDRESS REGISTERS (CONTINUED)

Register	Identifier	Offset
ADD MAC High 14	MAAH14	0x00F4
ADD MAC Low 15	MAAL15	0x00F8
ADD MAC High 15	MAAH15	0x00FC

4.4 MAC/PHY and Control Registers

MAC Address Register Low (0x0200): MARL

This register along with other 2 MAC address registers are loaded starting at word location 0x10 of the EEPROM upon hardware reset. The register can be modified by software driver, but will not modify the original MAC address value in the EEPROM. These six bytes of MAC address in external EEPROM are loaded to these three registers as mapping below:

MARL[15:0] = EEPROM 0x1(MAC Byte 2 and 1)

MARM[15:0] = EEPROM 0x2(MAC Byte 4 and 3)

MARH[15:0] = EEPROM 0x3(MAC Byte 6 and 5)

The MAC address is used to define the individual destination address the KSZ8841-PMQL responds to when receiving frames. Network addresses are generally expressed in the form of 01:23:45:67:89:AB, where the bytes are received from left to right, and the bits within each byte are received from right to left (LSB to MSB). For example, the actual transmitted and received bits are on the order of 10000000 11000100 10100010 11100110 10010001 11010101.

These three registers value for MAC address 01:23:45:67:89:AB will be held as below:

MARL[15:0] = 0x89AB

MARM[15:0] = 0x4567

MARH[15:0] = 0x0123

The following table shows the register bit fields for low word of MAC address.

TABLE 4-38: MAC ADDRESS REGISTER LOW (0X0200): MARL

Bit	Default	R/W	Description
15 - 0	—	R/W	MARL MAC Address Low The least significant word of the MAC address

MAC Address Register Middle (0x0202): MARM

The following table shows the register bit fields for middle word of MAC address.

TABLE 4-39: MAC ADDRESS REGISTER MIDDLE (0X0202): MARM

Bit	Default	R/W	Description
15 - 0	—	R/W	MARM MAC Address Middle The middle word of the MAC address

MAC Address Register High (0x0204): MARH

The following table shows the register bit fields for high word of MAC address.

TABLE 4-40: MAC ADDRESS REGISTER HIGH (0X0204): MARH

Bit	Default	R/W	Description
15 - 0	—	R/W	MARH MAC Address High The Most significant word of the MAC address

KSZ8841-PMQL

On-Chip Bus Control Register (Offset 0x0210): OBCR

This register controls the on-chip bus speed for the KSZ8841-PMQL operations. It's used for power management when the external host CPU is running a slow frequency. The default of the on-chip bus speed is 25 MHz. When the external host CPU is running at a higher clock rate, it's recommended the on-chip bus is adjusted accordingly for the best performance.

TABLE 4-41: ON-CHIP BUS CONTROL REGISTER (OFFSET 0X0210): OBCR

Bit	Default	R/W	Description
15 - 2	—	RO	Reserved
1 - 0	0x3	R/W	OBSC On-Chip Bus Speed Control 00 = 125 MHz 01 = 62.5 MHz 10 = 41.66 MHz 11 = 25 MHz

EEPROM Control Register (Offset 0x0212): EEPCR

KSZ8841-PMQL supports both with and without EEPROM system design. To support external EEPROM, tie the EEPROM Enable (EEEN) pin to high; otherwise, tie it to Low (or no connect). Also, KSZ8841-PMQL allows software to access (read and write) EEPROM directly. That is, the EEPROM access timing can be fully controlled by software if EEPROM Software Access bit is set.

TABLE 4-42: EEPROM CONTROL REGISTER (OFFSET 0X0212): EEPCR

Bit	Default	R/W	Description
15 - 5	0	RO	Reserved
4	0	R/W	EESA EEPROM Software Access 1 = Enable software to access EEPROM through bit 14 to bit 11. 0 = Disable software to access EEPROM.
3	00	RO	EECB EEPROM Status Bits Bit 3: Data receive from EEPROM. This bit directly reflects the value of the EEDI pin.
2	00	R/W	EECB EEPROM Control Bits Bit 2: Data In to EEPROM. This bit directly controls the device's the EEDO pin.
1	00	R/W	EECB EEPROM Control Bits Bit 1: Serial Clock. This bit directly controls the device's the EESK pin.
0	00	R/W	EECB EEPROM Control Bits Bit 0: Chip Select. This bit directly controls the device's the EECS pin.

Memory BIST Info Register (Offset 0x0214): MBIR

The following table shows the register bit fields.

TABLE 4-43: MEMORY BIST INFO REGISTER (OFFSET 0X0214): MBIR

Bit	Default	R/W	Description
15 - 13	0x0	RO	Reserved
12	—	RO	TXMBF TX Memory Bits Finish When set, it indicates the Memory Built In Self Test has completed for the TX Memory.
11	—	RO	TXMBFA TX Memory Bits Fail When set, it indicates the Memory Built In Self Test has failed.
10 - 5	—	RO	Reserved
4	—	RO	RXMBF RX Memory Bits Finish When set, it indicates the Memory Built In Self Test has completed for the RX Memory.

TABLE 4-43: MEMORY BIST INFO REGISTER (OFFSET 0X0214): MBIR (CONTINUED)

Bit	Default	R/W	Description
3	—	RO	RXMBFA RX Memory Bits Fail When set, it indicates the Memory Built In Self Test has failed.
2 - 0	—	RO	Reserved

Global Reset Register (Offset 0x0216): GRR

This register holds control information programmed by the CPU to control the global soft reset function.

TABLE 4-44: GLOBAL RESET REGISTER (OFFSET 0X0216): GRR

Bit	Default	R/W	Description
15 - 1	0x00	RO	Reserved
0	0	R/W	Global Soft Reset 1 = Software reset active 0 = Software reset inactive Soft reset will affect all of the registers except PCI configuration registers.

Power Management Capabilities Register (Offset 0x0218): PMCR

This register is a read-only register that provides information on the KSZ8841-PMQL power management capabilities. These bits are automatically downloaded from the Configparam word of EEPROM, if pin EEEN is pulled high (enable EEPROM). The PMCR register bits [15-0] are mirrored to CCID register bits [31-16].

TABLE 4-45: POWER MANAGEMENT CAPABILITIES REGISTER (OFFSET 0X0218): PMCR

Bit	Default	R/W	Description
15	0	RO	PME Support D3 (cold) This bit is 0 only; the KSZ8841-PMQL does not support PME in D3(cold) power state.
14	1	RO	PME Support D3 (hot) This bit is 1 only, it is indicating that the KSZ8841-PMQL can assert PME event (PMEN pin 14) in D3(hot) power state.
13	0	RO	PME Support D2 If this bit is set, the KSZ8841-PMQL asserts PME event (PMEN pin 14) when the KSZ8841PMQL is in D2 power state and PME_EN (see bit8 in PMCS register) is set. Otherwise, the KSZ8841PMQL does not assert PME event (PMEN pin 14) when the KSZ8841PMQL is in D2 power state. The value of this bit is loaded from the PME_D2 bit in the EEPROM 0x6 word.
12	0	RO	PME Support D1 If this bit is set, the KSZ8841-PMQL asserts PME event (PMEN pin 14) when the KSZ8841-PMQL is in D1 power state and PME_EN (see bit8 in PMCS register) is set. Otherwise, the KSZ8841M does not assert PME event (PMEN pin 14) when the KSZ8841M is in D1 power state. The value of this bit loaded from the PME_D1 bit in the EEPROM 0x6 word.
11	0	RO	PME Support D0 This bit is 0 only, it indicates that the KSZ8841-PMQL does not assert PME event (PMEN pin 14) in D0 power state.
10	0	RO	D2 Support If this bit is set, it indicates that the KSZ8841-PMQL support D2 power state. The value of this bit is loaded from the D2_SUP bit in the EEPROM 0x6 word. (This bit is 0 only if without EEPROM).
9	0	RO	D1 Support If this bit is set, it indicates that the KSZ8841-PMQL support D1 power state. The value of this bit loaded from the D1_SUP bit in the EEPROM 0x6 word. (This bit is 0 only if without EEPROM).

TABLE 4-45: POWER MANAGEMENT CAPABILITIES REGISTER (OFFSET 0X0218): PMCR

Bit	Default	R/W	Description
8 - 6	000	RO	Auxiliary Current This 3-bit field reports the 3.3Vaux auxiliary current requirements for the PCI function. If PME# generation from D3_cold is not supported by the function, this field must return a value of 000 when read.
5	0	RO	Device Specific Initialization Indicates whether special initialization of this function is required (beyond the standard PCI configuration header) before the generic class device driver is able to use it. Note that this bit is not used by some operating systems. Microsoft Windows and Windows NT, for instance, do not use this bit to determine whether to use D3. Instead, they use the driver's capabilities to determine this. A "1" indicates that the function requires a device specific initialization sequence following transition to the D0 uninitialization state. The value of this bit is loaded from the PME_DSI bit in the EEPROM 0X6 word.
4	0	RO	Reserved
3	0	RO	PME Clock When this bit is a "1", it indicates that the function relies on the presence of the PCI clock for PME# operation. When this bit is a "0", it indicates that no PCI clock is required for the function to generate PME#. The value of this bit is loaded from the PME_CK bit in the EEPROM 0x6 word.
2 - 0	0	RO	Power Management PCI Version The value of this bit is loaded from the PME_VER[2:0] bits in the EEPROM 0x6 word.

Wakeup Frame Control Register (Offset 0x021A): WFCR

This register holds control information programmed by the CPU to control the transmit module function.

TABLE 4-46: WAKEUP FRAME CONTROL REGISTER (OFFSET 0X021A): WFCR

Bit	Default	R/W	Description
15 - 8	0x00	RO	Reserved
7	0	R/W	MPRXE Magic Packet RX Enable When set, it enables the magic packet pattern detection. When reset, the magic packet pattern detection is disabled.
6 - 4	0x0	RO	Reserved
3	0	R/W	WF3E Wake up Frame 3 Enable When set, it enables the wake up frame 3 pattern detection. When reset, the wake up frame pattern detection is disabled.
2	0	R/W	WF2E Wake up Frame 2 Enable When set, it enables the wake up frame 2 pattern detection. When reset, the wake up frame pattern detection is disabled.
1	0	R/W	WF1E Wake up Frame 1 Enable When set, it enables the wake up frame 1 pattern detection. When reset, the wake up frame pattern detection is disabled.
0	0	R/W	WF0E Wake up Frame 0 Enable When set, it enables the wake up frame 0 pattern detection. When reset, the wake up frame pattern detection is disabled.

Wakeup Frame 0 CRC0 Register (Offset 0x0220): WF0CRC0

This register contains the expected CRC values of the Wake up frame 0 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, taken over the bytes specified in the wake up byte mask registers.

TABLE 4-47: WAKEUP FRAME 0 CRC0 REGISTER (OFFSET 0X0220): WF0CRC0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF0CRC0 Wake up Frame 0 CRC (lower 16 bits) The expected CRC value of a wake up frame 0 pattern.

Wakeup Frame 0 CRC1 Register (Offset 0x0222): WF0CRC1

This register contains the expected CRC values of the Wake up frame 0 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, taken over the bytes specified in the wake up byte mask registers.

TABLE 4-48: WAKEUP FRAME 0 CRC1 REGISTER (OFFSET 0X0222): WF0CRC1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF0CRC1 Wake up Frame 0 CRC (upper 16 bits) The expected CRC value of a wake up frame 0 pattern.

Wakeup Frame 0 Byte Mask 0 Register (Offset 0x0224): WF0BM0

This register contains the first 16 bytes mask values of the Wake up frame 0 pattern. Setting bit 0 selects the first byte of the Wake up frame 0; setting bit 15 selects the 16th byte of the Wake up frame 0.

TABLE 4-49: WAKEUP FRAME 0 BYTE MASK 0 REGISTER (OFFSET 0X0224): WF0BM0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF0BM0 Wake up Frame 0 Byte Mask 0 The first 16 bytes mask of a wake up frame 0 pattern.

Wakeup Frame 0 Byte Mask 1 Register (Offset 0x0226): WF0BM1

This register contains the next 16 bytes mask values of the Wake up frame 0 pattern. Setting bit 0 selects the 17th byte of the Wake up frame 0; setting bit 15 selects the 32nd byte of the Wake up frame 0.

TABLE 4-50: WAKEUP FRAME 0 BYTE MASK 1 REGISTER (OFFSET 0X0226): WF0BM1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF0BM1 Wake up Frame 0 Byte Mask 1 The next 16 bytes mask covering bytes 17 to 32 of a wake up frame 0 pattern.

Wakeup Frame 0 Byte Mask 2 Register (Offset 0x0228): WF0BM2

This register contains the next 16 bytes mask values of the Wake up frame 0 pattern. Setting bit 0 selects the 33rd byte of the Wake up frame 0; setting bit 15 selects the 48th byte of the Wake up frame 0.

TABLE 4-51: WAKEUP FRAME 0 BYTE MASK 2 REGISTER (OFFSET 0X0228): WF0BM2

Bit	Default	R/W	Description
15 - 0	—	R/W	WF0BM2 Wake up Frame 0 Byte Mask 2 The next 16 bytes mask covering bytes 33 to 48 of a wake up frame 0 pattern.

KSZ8841-PMQL

Wakeup Frame 0 Byte Mask 3 Register (Offset 0x022A): WF0BM3

This register contains the last 16 bytes mask values of the Wake up frame 0 pattern. Setting bit 0 selects the 49th byte of the Wake up frame 0; setting bit 15 selects the 64th byte of the Wake up frame 0.

TABLE 4-52: WAKEUP FRAME 0 BYTE MASK 3 REGISTER (OFFSET 0X022A): WF0BM3

Bit	Default	R/W	Description
15 - 0	—	R/W	WF0BM3 Wake up Frame 0 Byte Mask 3 The last 16 bytes mask covering bytes 49 to 64 of a wake up frame 0 pattern.

Wakeup Frame 1 CRC0 Register (Offset 0x0230): WF1CRC0

This register contains the expected CRC values of the Wake up frame 1 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, it is taken over the bytes specified in the wake up byte mask registers.

TABLE 4-53: WAKEUP FRAME 1 CRC0 REGISTER (OFFSET 0X0230): WF1CRC0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF1CRC0 Wake up Frame 1 CRC (lower 16 bits) The expected CRC value of a wake up frame 1 pattern.

Wakeup Frame 1 CRC1 Register (Offset 0x0232): WF1CRC1

This register contains the expected CRC values of the Wake up frame 1 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, it is taken over the bytes specified in the wake up byte mask registers.

TABLE 4-54: WAKEUP FRAME 1 CRC1 REGISTER (OFFSET 0X0232): WF1CRC1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF1CRC1 Wake up Frame 1 CRC (upper 16 bits) The expected CRC value of a wake up frame 1 pattern.

Wakeup Frame 1 Byte Mask 0 Register (Offset 0x0234): WF1BM0

This register contains the first 16 bytes mask values of the Wake up frame 1 pattern. Setting bit 0 selects the first byte of the Wake up frame 1; setting bit 15 selects the 16th byte of the Wake up frame 1.

TABLE 4-55: WAKEUP FRAME 1 BYTE MASK 0 REGISTER (OFFSET 0X0234): WF1BM0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF1BM0 Wake up Frame 1 Byte Mask 0 The first 16 bytes mask of a wake up frame 1 pattern.

Wakeup Frame 1 Byte Mask 1 Register (Offset 0x0236): WF1BM1

This register contains the next 16 bytes mask values of the Wake up frame 1 pattern. Setting bit 0 selects the 17th byte of the Wake up frame 1; setting bit 15 selects the 32nd byte of the Wake up frame 1.

TABLE 4-56: WAKEUP FRAME 1 BYTE MASK 1 REGISTER (OFFSET 0X0236): WF1BM1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF1BM1 Wake up Frame 1 Byte Mask 1 The next 16 bytes mask covering bytes 17 to 32 of a wake up frame 1 pattern.

Wakeup Frame 1 Byte Mask 2 Register (Offset 0x0238): WF1BM2

This register contains the next 16 bytes mask values of the Wake up frame 1 pattern. Setting bit 0 selects the 33rd byte of the Wake up frame 1; setting bit 15 selects the 48th byte of the Wake up frame 1.

TABLE 4-57: WAKEUP FRAME 1 BYTE MASK 2 REGISTER (OFFSET 0X0238): WF1BM2

Bit	Default	R/W	Description
15 - 0	—	R/W	WF1BM2 Wake up Frame 1 Byte Mask 2 The next 16 byte mask covering bytes 33 to 48 of a wake up frame1 pattern.

Wakeup Frame 1 Byte Mask 3 Register (Offset 0x023A): WF1BM3

This register contains the last 16 bytes mask values of the Wake up frame 1 pattern. Setting bit 0 selects the 49th byte of the Wake up frame 1; setting bit 15 selects the 64th byte of the Wake up frame 1.

TABLE 4-58: WAKEUP FRAME 1 BYTE MASK 3 REGISTER (OFFSET 0X023A): WF1BM3

Bit	Default	R/W	Description
15 - 0	—	R/W	WF1BM3 Wake up Frame 1 Byte Mask 3 The last 16 bytes mask covering bytes 49 to 64 of a wake up frame 1 pattern.

Wakeup Frame 2 CRC0 Register (Offset 0x0240): WF2CRC0

This register contains the expected CRC values of the Wake up frame 2 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, it is taken over the bytes specified in the wake up byte mask registers.

TABLE 4-59: WAKEUP FRAME 2 CRC0 REGISTER (OFFSET 0X0240): WF2CRC0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF2CRC0 Wake up Frame 2 CRC (lower 16 bits) The expected CRC value of a wake up frame 2 pattern.

Wakeup Frame 2 CRC1 Register (Offset 0x0242): WF2CRC1

This register contains the expected CRC values of the Wake up frame 2 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, it is taken over the bytes specified in the wake up byte mask registers.

TABLE 4-60: WAKEUP FRAME 2 CRC1 REGISTER (OFFSET 0X0242): WF2CRC1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF2CRC1 Wake up Frame 2 CRC (upper 16 bits) The expected CRC value of a wake up frame 2 pattern.

Wakeup Frame 2 Byte Mask 0 Register (Offset 0x0244): WF2BM0

This register contains the first 16 bytes mask values of the Wake up frame 2 pattern. Setting bit 0 selects the first byte of the Wake up frame 2; setting bit 15 selects the 16th byte of the Wake up frame 2.

TABLE 4-61: WAKEUP FRAME 2 BYTE MASK 0 REGISTER (OFFSET 0X0244): WF2BM0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF2BM0 Wake up Frame 2 Byte Mask 0 The first 16 byte mask of a wake up frame 2 pattern.

KSZ8841-PMQL

Wakeup Frame 2 Byte Mask 1 Register (Offset 0x0246): WF2BM1

This register contains the next 16 bytes mask values of the Wake up frame 2 pattern. Setting bit 0 selects the 17th byte of the Wake up frame 2; setting bit 15 selects the 32nd byte of the Wake up frame 2.

TABLE 4-62: WAKEUP FRAME 2 BYTE MASK 1 REGISTER (OFFSET 0X0246): WF2BM1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF2BM1 Wake up Frame 2 Byte Mask 1 The next 16 bytes mask covering bytes 17 to 32 of a wake up frame 2 pattern.

Wakeup Frame 2 Byte Mask 2 Register (Offset 0x0248): WF2BM2

This register contains the next 16 bytes mask values of the Wake up frame 2 pattern. Setting bit 0 selects the 33rd byte of the Wake up frame 2; setting bit 15 selects the 48th byte of the Wake up frame 2.

TABLE 4-63: WAKEUP FRAME 2 BYTE MASK 2 REGISTER (OFFSET 0X0248): WF2BM2

Bit	Default	R/W	Description
15 - 0	—	R/W	WF2BM2 Wake up Frame 2 Byte Mask 2 The next 16 bytes mask covering bytes 33 to 48 of a wake up frame 2 pattern.

Wakeup Frame 2 Byte Mask 3 Register (Offset 0x024A): WF2BM3

This register contains the last 16 bytes mask values of the Wake up frame 2 pattern. Setting bit 0 selects the 49th byte of the Wake up frame 2; setting bit 15 selects the 64th byte of the Wake up frame 2.

TABLE 4-64: WAKEUP FRAME 2 BYTE MASK 3 REGISTER (OFFSET 0X024A): WF2BM3

Bit	Default	R/W	Description
15 - 0	—	R/W	WF2BM3 Wake up Frame 2 Byte Mask 3 The last 16 bytes mask covering bytes 49 to 64 of a wake up frame 2 pattern.

Wakeup Frame 3 CRC0 Register (Offset 0x0250): WF3CRC0

This register contains the expected CRC values of the Wake up frame 3 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, it is taken over the bytes specified in the wake up byte mask registers.

TABLE 4-65: WAKEUP FRAME 3 CRC0 REGISTER (OFFSET 0X0250): WF3CRC0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF3CRC0 Wake up Frame 3 CRC (lower 16 bits) The expected CRC value of a wake up frame 3 pattern.

Wakeup Frame 3 CRC1 Register (Offset 0x0252): WF3CRC1

This register contains the expected CRC values of the Wake up frame 3 pattern.

The value of the CRC calculated is based on the IEEE 802.3 Ethernet standard, it is taken over the bytes specified in the wake up byte mask registers.

TABLE 4-66: WAKEUP FRAME 3 CRC1 REGISTER (OFFSET 0X0252): WF3CRC1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF3CRC1 Wake up Frame 3 CRC (upper 16 bits) The expected CRC value of a wake up frame 3 pattern.

Wakeup Frame 3 Byte Mask 0 Register (Offset 0x0254): WF3BM0

This register contains the first 16 bytes mask values of the Wake up frame 3 pattern. Setting bit 0 selects the first byte of the Wake up frame 3; setting bit 15 selects the 16th byte of the Wake up frame 3.

TABLE 4-67: WAKEUP FRAME 3 BYTE MASK 0 REGISTER (OFFSET 0X0254): WF3BM0

Bit	Default	R/W	Description
15 - 0	—	R/W	WF3BM0 Wake up Frame 3 Byte Mask 0 The first 16 bytes mask of a wake up frame 3 pattern.

Wakeup Frame 3 Byte Mask 1 Register (Offset 0x0256): WF3BM1

This register contains the next 16 bytes mask values of the Wake up frame 3 pattern. Setting bit 0 selects the 17th byte of the Wake up frame 3; setting bit 15 selects the 32nd byte of the Wake up frame 3.

TABLE 4-68: WAKEUP FRAME 3 BYTE MASK 1 REGISTER (OFFSET 0X0256): WF3BM1

Bit	Default	R/W	Description
15 - 0	—	R/W	WF3BM1 Wake up Frame 3 Byte Mask 1 The next 16 bytes mask covering bytes 17 to 32 of a wake up frame 3 pattern.

Wakeup Frame 3 Byte Mask 2 Register (Offset 0x0258): WF3BM2

This register contains the next 16 bytes mask values of the Wake up frame 3 pattern. Setting bit 0 selects the 33rd byte of the Wake up frame 3; setting bit 15 selects the 48th byte of the Wake up frame 3.

TABLE 4-69: WAKEUP FRAME 3 BYTE MASK 2 REGISTER (OFFSET 0X0258): WF3BM2

Bit	Default	R/W	Description
15 - 0	—	R/W	WF3BM2 Wake up Frame 3 Byte Mask 2 The next 16 bytes mask covering bytes 33 to 48 of a wake up frame 3 pattern.

Wakeup Frame 3 Byte Mask 3 Register (Offset 0x025A): WF3BM3

This register contains the last 16 bytes mask values of the Wake up frame 3 pattern. Setting bit 0 selects the 49th byte of the Wake up frame 3; setting bit 15 selects the 64th byte of the Wake up frame 3.

TABLE 4-70: WAKEUP FRAME 3 BYTE MASK 3 REGISTER (OFFSET 0X025A): WF3BM3

Bit	Default	R/W	Description
15 - 0	—	R/W	WF3BM3 Wake up Frame 3 Byte Mask 3 The last 16 bytes mask covering bytes 49 to 64 of a wake up frame 3 pattern.

Chip ID and Enable Register (Offset 0x0400): CIDER

This register contains the chip ID, and the chip enables control.

TABLE 4-71: CHIP ID AND ENABLE REGISTER (OFFSET 0X0400): CIDER

Bit	Default	R/W	Description
15 - 8	0x88	RO	Family ID Chip family ID
7 - 4	0x05	RO	Chip ID 0x05 is assigned to KSZ8841-PMQL
3 - 1	000	RO	Revision ID
0	—	R/W	Start Controller 1 = Start the chip operation 0 = Stop the chip operation

KSZ8841-PMQL

Chip Global Control Register (Offset 0x040A): CGCR

This register contains the global control for the chip function.

TABLE 4-72: CHIP GLOBAL CONTROL REGISTER (OFFSET 0X040A): CGCR

Bit	Default	R/W	Description
15	0	R/W	LEDSEL1 See description for bit 9.
14 - 10	—	R/W	Reserved
9	0	R/W	LEDSEL0 The two register bits LEDSEL1 and LEDSEL0, are used to select the LED mode. The LED Indicators, are defined as below:
			— [LEDSEL1, LEDSEL0] [0, 0] [0, 1]
			P1LED3 — —
			P1LED2 Link/Activity 100Link/Activity
			P1LED1 Full-Duplex/Col 10Link/Activity
			P1LED0 Speed Full-Duplex
			— [LEDSEL1, LEDSEL0] [1, 0] [1, 1]
			P1LED3 Activity —
			P1LED2 Link —
			P1LED1 Full-Duplex/Col —
			P1LED0 Speed —
8 - 0	0	R/W	Reserved

Indirect Access Control Register (Offset 0x04A0): IACR

This register contains the indirect control for the MIB counter. Write IACR will actually trigger a command. Read or write access is determined by this register bit 12.

TABLE 4-73: INDIRECT ACCESS CONTROL REGISTER (OFFSET 0X04A0): IACR

Bit	Default	R/W	Description
15 - 13	000	R/W	Reserved
12	0	R/W	Read High. Write Low 1 = Read cycle 0 = Write cycle
11 - 10	00	R/W	Table select 11 = MIB counter selected
9 - 0	0x000	R/W	Indirect address Bit 9 - 0 of indirect address

Indirect Access Data Register 1 (Offset 0x04A2): IADR1

This register contains the indirect data for the chip function.

TABLE 4-74: INDIRECT ACCESS DATA REGISTER 1 (OFFSET 0X04A2): IADR1

Bit	Default	R/W	Description
15 - 8	0x00	RO	Reserved
7	0	RO	CPU Read Status Only for dynamic and statistics counter reads. 1 = Read is still in progress 0 = Read has completed
6 - 3	0x0	RO	Reserved
2 - 0	000	RO	Reserved

Indirect Access Data Register 2 (Offset 0x04A4): IADR2

This register contains the indirect data for the chip function.

TABLE 4-75: INDIRECT ACCESS DATA REGISTER 2 (OFFSET 0X04A4): IADR2

Bit	Default	R/W	Description
15 - 0	0x0000	R/W	Indirect data Bit 47 - 32 of indirect data

Indirect Access Data Register 3 (Offset 0x04A6): IADR3

This register contains the indirect data for the chip function.

TABLE 4-76: INDIRECT ACCESS DATA REGISTER 3 (OFFSET 0X04A6): IADR3

Bit	Default	R/W	Description
15 - 0	0x0000	R/W	Reserved

Indirect Access Data Register 4 (Offset 0x04A8): IADR4

This register contains the indirect data for the chip function.

TABLE 4-77: INDIRECT ACCESS DATA REGISTER 4 (OFFSET 0X04A8): IADR4

Bit	Default	R/W	Description
15 - 0	0x0000	R/W	Indirect data Bit 15 - 0 of indirect data

Indirect Access Data Register 5 (Offset 0x04AA): IADR5

This register contains the indirect data for the chip function.

TABLE 4-78: INDIRECT ACCESS DATA REGISTER 5 (OFFSET 0X04AA): IADR5

Bit	Default	R/W	Description
15 - 0	0x0000	R/W	Indirect data Bit 31 - 16 of indirect data

Reserved (Offset 0x04C0-0x04CF)

PHY 1 MII Register Basic Control Register (Offset 0x04D0): P1MBCR

This register contains the MII register control for the chip function.

TABLE 4-79: PHY 1 MII REGISTER BASIC CONTROL REGISTER (OFFSET 0X04D0): P1MBCR

Bit	Default	R/W	Description	Bit Same As
15	0	RO	Soft reset NOT SUPPORTED	—
14	0	R/W	Reserved	—

KSZ8841-PMQL

TABLE 4-79: PHY 1 MII REGISTER BASIC CONTROL REGISTER (OFFSET 0X04D0): P1MBCR

Bit	Default	R/W	Description	Bit Same As
13	0	R/W	Force 100 1 = Force 100 Mbps if AN is disabled (bit12) 0 = Force 10 Mbps if AN is disabled (bit12)	P1CR4, bit 6
12	1	R/W	AN enable 1 = Auto-negotiation enabled 0 = Auto-negotiation disabled	P1CR4, bit 7
11	0	R/W	Power down 1 = Power down 0 = Normal operation	P1CR4, bit 11
10	0	RO	Isolate NOT SUPPORTED	—
9	0	R/W	Restart AN 1 = Restart auto-negotiation 0 = Normal operation	P1CR4, bit 13
8	0	R/W	Force full-duplex 1 = Force full-duplex if AN is disabled (bit12) 0 = Force half-duplex if AN is disabled (bit12)	P1CR4, bit 5
7	0	RO	Reserved	—
6	0	RO	Reserved	—
5	0	R/W	HP_mdix 1 = HP Auto MDIX mode 0 = Microchip Auto MDIX mode	P1SR, bit 15
4	0	R/W	Force MDIX 1 = Force MDIX 0 = Normal operation	P1CR4, bit 9
3	0	R/W	Disable MDIX 1 = Disable auto MDIX 0 = Normal operation	P1CR4, bit 10
2	0	R/W	Disable far end fault 1 = Disable far end fault detection 0 = Normal operation	P1CR4, bit 12
1	0	R/W	Disable transmit 1 = Disable transmit 0 = Normal operation	P1CR4, bit 14
0	0	R/W	Disable LED 1 = Disable LED 0 = Normal operation	P1CR4, bit 15

PHY 1 MII Register Basic Status Register (Offset 0x04D2): P1MBSR

This register contains the MII register control for the chip function.

TABLE 4-80: PHY 1 MII REGISTER BASIC STATUS REGISTER (OFFSET 0X04D2): P1MBSR

Bit	Default	R/W	Description	Bit Same As
15	0	RO	T4 capable 1 = 100BASE-T4 capable 0 = Not 100BASE-T4 capable	—
14	1	RO	100 Full capable 1 = 100BASE-TX full-duplex capable 0 = Not 100BASE-TX full-duplex capable	Always 1
13	1	RO	100 Half capable 1 = 100BASE-TX half-duplex capable 0 = Not 100BASE-TX half-duplex capable	Always 1

TABLE 4-80: PHY 1 MII REGISTER BASIC STATUS REGISTER (OFFSET 0X04D2): P1MBSR

Bit	Default	R/W	Description	Bit Same As
12	1	RO	10 Full capable 1 = 10BASE-T full-duplex capable 0 = Not 10BASE-T full-duplex capable	Always 1
11	1	RO	10 Half capable 1 = 10BASE-T half-duplex capable 0 = Not 10BASE-T half-duplex capable	Always 1
10 - 7	0	RO	Reserved	—
6	0	RO	Preamble suppressed NOT SUPPORTED	—
5	0	RO	AN complete 1 = Auto-negotiation complete 0 = Auto-negotiation not completed	P1SR, bit 6
4	0	RO	Far end fault 1 = Far end fault detected 0 = No far end fault detected	P1SR, bit 8
3	1	RO	AN capable 1 = Auto-negotiation capable 0 = Not auto-negotiation capable	P1CR4, bit 7
2	0	RO	Link status 1 = Link is up 0 = Link is down	P1SR, bit 5
1	0	RO	Reserved	—
0	0	RO	Extended capable 1 = Extended register capable 0 = Not extended register capable	—

PHY 1 PHYID Low Register (Offset 0x04D4): PHY1ILR

This register contains the PHY ID (low) for the chip function.

TABLE 4-81: PHY 1 PHYID LOW REGISTER (OFFSET 0X04D4): PHY1ILR

Bit	Default	R/W	Description
15 - 0	0x1430	RO	PHYID low Low order PHYID bits

PHY 1 PHYID High Register (Offset 0x04D6): PHY1IHR

This register contains the PHY ID (high) for the chip function.

TABLE 4-82: PHY 1 PHYID HIGH REGISTER (OFFSET 0X04D6): PHY1IHR

Bit	Default	R/W	Description
15 - 0	0x0022	RO	PHYID high High order PHYID bits

KSZ8841-PMQL

PHY 1 Auto-Negotiation Advertisement Register (Offset 0x04D8): P1ANAR

This register contains the auto-negotiation advertisement for the chip function.

TABLE 4-83: PHY 1 AUTO-NEGOTIATION ADVERTISEMENT REGISTER (OFFSET 0X04D8): P1ANAR

Bit	Default	R/W	Description	Bit Same As
15	0	RO	Next page NOT SUPPORTED	—
14	0	RO	Reserved	—
13	0	RO	Remote fault NOT SUPPORTED	—
12 - 11	0	RO	Reserved	—
10	1	R/W	Pause (follow control capability) 1 = Advertise pause ability 0 = Do not advertise pause ability	P1CR4, bit 4
9	0	R/W	Reserved	—
8	1	R/W	Adv 100 Full 1 = Advertise 100 full-duplex ability 0 = Do not advertise 100 full-duplex ability	P1CR4, bit 3
7	1	R/W	Adv 100 Half 1 = Advertise 100 half-duplex ability 0 = Do not advertise 100 half-duplex ability	P1CR4, bit 2
6	1	R/W	Adv 10 Full 1 = Advertise 10 full-duplex ability 0 = Do not advertise 10 full-duplex ability	P1CR4, bit 1
5	1	R/W	Adv 10 Half 1 = Advertise 10 half-duplex ability 0 = Do not advertise 10 half-duplex ability	P1CR4, bit 0
4 - 0	0_0001	RO	Selector field 802.3	—

PHY 1 Auto-Negotiation Link Partner Ability Register (Offset 0x04DA): P1ANLPR

This register contains the auto-negotiation link partner ability for the chip function.

TABLE 4-84: PHY 1 AUTO-NEGOTIATION LINK PARTNER ABILITY REGISTER (OFFSET 0X04DA): P1ANLPR

Bit	Default	R/W	Description	Bit Same As
15	0	RO	Next page NOT SUPPORTED	—
14	0	RO	LP ACK NOT SUPPORTED	—
13	0	RO	Remote fault NOT SUPPORTED	—
12 - 11	0	RO	Reserved	—
10	0	RO	Pause Link partner pause capability	P1SR, bit 4
9	0	RO	Reserved	—
8	0	RO	Adv 100 Full Link partner 100 full capability	P1SR, bit 3
7	0	RO	Adv 100 Half Link partner 100 half capability	P1SR, bit 2
6	0	RO	Adv 10 Full Link partner 10 full capability	P1SR, bit 1

TABLE 4-84: PHY 1 AUTO-NEGOTIATION LINK PARTNER ABILITY REGISTER (OFFSET 0X04DA): P1ANLPR (CONTINUED)

Bit	Default	R/W	Description	Bit Same As
5	0	RO	Adv 10 Half Link partner 10 half capability	P1SR, bit 0
4 - 0	0_0000	RO	Reserved	—

PHY1 LinkMD® Control/Status (Offset 0x04F0): P1VCT

This register contains the LinkMD control and status of PHY 1.

TABLE 4-85: PHY1 LINKMD CONTROL/STATUS (OFFSET 0X04F0): P1VCT

Bit	Default	R/W	Description	Bit Same As
15	0	R/W SC	Vct_enable 1 = The cable diagnostic test is enabled. It'll be self-cleared after VCT test is done 0 = Indicates the cable diagnostic test is completed and the status information is valid for read	P1SCSLMD, bit 12
14 - 13	00	RO	Vct_result [00] = Normal condition [01] = Open condition has been detected in cable [10] = Short condition has been detected in cable [11] = Cable diagnostic test is failed	P1SCSLMD, bit 14:13
12	—	RO	Vct 10M short 1 = Less than 10 meter short	P1SCSLMD, bit 15
11 - 9	0	RO	Reserved	—
8 - 0	0	RO	Vct_fault_count Distance to the fault. The distance is approximately 0.4m x vct_fault_count	P1SCSLMD, bits 8:0

PHY1 Special Control/Status Register (Offset 0x04F2): P1PHYCTRL

This register contains the control and status information of PHY1.

TABLE 4-86: PHY1 SPECIAL CONTROL/STATUS REGISTER (OFFSET 0X04F2): P1PHYCTRL

Bit	Default	R/W	Description	Bit Same As
15 - 6	0	RO	Reserved	—
5	0	RO	Polarity reverse (polrvs) 1 = Polarity is reversed 0 = Polarity is not reversed	P1SR, bit 13
4	0	RO	MDIX Status (mdix_st) 1 = MDIX 0 = MDI	P1SR, bit 7
3	0	R/W	Force Link (force_lnk) 1 = Force link pass 0 = Normal Operation	P1SCSLMD, bit 11
2	1	R/W	Power Saving (pwrsave) 1 = Disable 0 = Enable power saving	P1SCSLMD, bit 10
1	0	R/W	Remote loopback (rlb) 1 = Loop back at PMD/PMA of port's PHY (RXP1/RXM1 -> TXP1/TXM1) 0 = Normal operation.	P1SCSLMD, bit 9
0	0	R/W	Reserved	—

KSZ8841-PMQL

Reserved (Offset 0x04F8 - 0x04FA)

TABLE 4-87: RESERVED (OFFSET 0X04F8 - 0X04FA)

Bit	Default	R/W	Description
15 - 0	0x0000	RO	Reserved

Port 1 PHY Special Control/Status, LinkMD® (Offset 0x0510): P1SCSLMD

This register contains the port LinkMD control register for the chip function.

TABLE 4-88: PORT 1 PHY SPECIAL CONTROL/STATUS, LINKMD (OFFSET 0X0510): P1SCSLMD

Bit	Default	R/W	Description	Bit Same As
15	0	RO	Vct 10M short Less than 10 meter short	P1VCT, bit 12
14 - 13	00	RO	Vct result [00] = Normal condition [01] = Open condition has been detected in cable [10] = Short condition has been detected in cable [11] = Cable diagnostic test is failed	P1VCT, bits 14:13
12	0	R/W SC	Vct enable 1 = The cable diagnostic test is enabled. It'll be self-cleared after VCT test is done 0 = It indicates the cable diagnostic test is completed and the status information is valid for read	P1VCT, bit 15
11	0	R/W	Force Link 1 = Force link pass 0 = Normal Operation	P1PHYCTRL, bit 3
10	1	R/W	Power Saving 1 = Disable 0 = Enable power saving	P1PHYCTRL, bit 2
9	0	R/W	Remote loopback 1 = Loop back at PMD/PMA of port's PHY (RXP1/RXM1 -> TXP1/TXM1) 0 = Normal operation.	P1PHYCTRL, bit 1
8 - 0	0x000	RO	VCT fault count Distance to the fault. The distance is approximately 0.4m x vct_fault_count	P1VCT, bits 8:0

Port 1 Control Register 4 (Offset 0x0512): P1CR4

This register contains the global per port control for the chip function.

TABLE 4-89: PORT 1 CONTROL REGISTER 4 (OFFSET 0X0512): P1CR4

Bit	Default	R/W	Description	Bit Same As
15	0	R/W	LED off 1 = Turn off all port's LEDs (LED1_3, LED1_2, LED1_1, LED1_0. These pins will be driven high if this bit is set to one 0 = Normal operation	P1MBCR, bit 0
14	0	R/W	Txids 1 = Disable port's transmitter 0 = Normal operation	P1MBCR, bit 1
13	0	R/W	Restart AN 1 = Restart auto-negotiation 0 = Normal operation	P1MBCR, bit 9

TABLE 4-89: PORT 1 CONTROL REGISTER 4 (OFFSET 0X0512): P1CR4 (CONTINUED)

Bit	Default	R/W	Description	Bit Same As
12	0	R/W	Disable Far end fault 1 = Disable far end fault detection & pattern transmission. 0 = Enable far end fault detection & pattern transmission.	P1MBCR, bit 2
11	0	R/W	Power down 1 = Power down 0 = Normal operation	P1MBCR, bit 11
10	0	R/W	Disable auto MDI/MDIX 1 = Disable auto MDI/MDIX function 0 = Enable auto MDI/MDIX function	P1MBCR, bit 3
9	0	R/W	Force MDIX 1 = If auto MDI/MDIX is disabled, force PHY into MDIX mode 0 = Do not force PHY into MDIX mode	P1MBCR, bit 4
8	0	—	Reserved	P1MBCR, bit 14
7	1	R/W	Auto-Negotiation Enable 1 = Auto-negotiation is enable 0 = Disable auto-negotiation, speed and duplex are decided by bit 6 and 5 of the same register.	P1MBCR, bit 12
6	0	R/W	Force Speed 1 = Force 100BT if AN is disabled (bit 7) 0 = Force 10BT if AN is disabled (bit 7)	P1MBCR, bit 13
5	0	R/W	Force duplex 1 = Force full-duplex if (1) AN is disabled or (2) AN is enabled but failed. 0 = Force half-duplex if (1) AN is disabled or (2) AN is enabled but failed.	P1MBCR, bit 9
4	1	R/W	Advertised flow control capability 1 = Advertise flow control (pause) capability 0 = Suppress flow control (pause) capability from transmission to link partner	P1ANAR, bit 4
3	1	R/W	Advertised 100BT Full-duplex capability 1 = Advertise 100BT Full-duplex capability 0 = Suppress 100BT Full-duplex capability from transmission to link partner	P1ANAR, bit 3
2	1	R/W	Advertised 100BT half-duplex capability 1 = Advertise 100BT Half-duplex capability 0 = Suppress 100BT Half-duplex capability from transmission to link partner	P1ANAR, bit 2
1	1	R/W	Advertised 10BT Full-duplex capability 1 = Advertise 10BT Full-duplex capability 0 = Suppress 10BT Full-duplex capability from transmission to link partner	P1ANAR, bit 1
0	1	R/W	Advertised 10BT half-duplex capability 1 = Advertise 10BT Half-duplex capability 0 = Suppress 10BT Half-duplex capability from transmission to link partner	P1ANAR, bit 0

KSZ8841-PMQL

Port 1 Status Register (Offset 0x0514): P1SR

This register contains the global per port status for the chip function.

TABLE 4-90: PORT 1 STATUS REGISTER (OFFSET 0X0514): P1SR

Bit	Default	R/W	Description	Bit Same As
15	0	R/W	HP_mdix 1 = HP Auto MDIX mode 0 = Microchip Auto MDIX mode	P1MBCR, bit 5
14	0	RO	Reserved	—
13	0	RO	Polarity reverse 1 = Polarity is reversed 0 = Polarity is not reversed	P1PHYCTRL, bit 5
12	0	RO	Receive flow control enable 1 = Receive flow control feature is active 0 = Receive flow control feature is inactive	—
11	0	RO	Transmit flow control enable 1 = Transmit flow control feature is active 0 = Transmit flow control feature is inactive	—
10	0	RO	Operation Speed 1 = Link speed is 100 Mbps 0 = Link speed is 10 Mbps	—
9	0	RO	Operation duplex 1 = Link duplex is full 0 = Link duplex is half	—
8	0	RO	Far end fault 1 = Far end fault status detected 0 = No Far end fault status detected	P1MBSR, bit 4
7	0	RO	MDIX status 1 = MDIX 0 = MDI	P1PHYCTRL, bit 4
6	0	RO	AN done 1 = AN done 0 = AN not done	P1MBSR, bit 5
5	0	RO	Link good 1 = Link good 0 = Link not good	P1MBSR, bit 2
4	0	RO	Partner flow control capability 1 = Link partner flow control (pause) capable 0 = Link partner not flow control (pause) capable	P1ANLPR, bit 10
3	0	RO	Partner 100BT full-duplex capability 1 = Link partner 100BT full-duplex capable 0 = Link partner not 100BT full-duplex capable	P1ANLPR, bit 8
2	0	RO	Partner 100BT half-duplex capability 1 = Link partner 100BT half-duplex capable 0 = Link partner not 100BT half-duplex capable	P1ANLPR, bit 7
1	0	RO	Partner 10BT full-duplex capability 1 = Link partner 10BT full-duplex capable 0 = Link partner not 10BT full-duplex capable	P1ANLPR, bit 6
0	0	RO	Partner 10BT half-duplex capability 1 = Link partner 10BT half-duplex capable 0 = Link partner not 10BT half-duplex capable	P1ANLPR, bit 5

Reserved (Offset 0x0516 – 0x0560)

TABLE 4-91: RESERVED (OFFSET 0X0516 – 0X0560)

Bit	Default	R/W	Description
15 - 0	0x0000	RO	Reserved

4.5 Management Information Base (MIB) Counters

The KSZ8841-PMQL provides 32 MIB counters to monitor the port activity for network management. The MIB counters are formatted as shown in the following table.

TABLE 4-92: FORMAT OF PORT MIB COUNTERS

Bit	Name	R/W	Description	Default
31	Overflow	RO	1 = Counter overflow. 0 = No counter overflow.	0
30	Count Valid	RO	1 = Counter value is valid. 0 = Counter value is not valid.	0
29 - 0	Counter Values	RO	Counter value	0

The port MIB counters are read using indirect memory access. The base address offsets is 0x00 and address ranges is 0x00 - 0x1F as shown in [Table 4-93](#).

The Port MIB counters read/write functions use Access Control register IACR (0x04A0) bit 12. The base address offset and address range for port 1 is 0x00 and range is (0x00-0x1F) that can be changed in register IACR (0x04A0) bits[9:0]. The data of MIB counters are from the Indirect Access data register IADR4 (0x04A8) and IADR5 (0x04AA) based on [Table 4-92](#).

TABLE 4-93: PORT 1'S MIB COUNTERS INDIRECT MEMORY OFFSETS

Offset	Counter Name	Description
0x0 (base address)	RxByte	Rx octet count including bad packets.
0x1	Reserved	Reserved. Do not write to this register.
0x2	RxUndersizePkt	Rx undersize packets w/ good CRC.
0x3	RxFragments	Rx fragment packets w/ bad CRC, symbol errors or alignment errors.
0x4	RxOversize	Rx oversize packets w/ good CRC (max: 1536 or 1522 bytes).
0x5	RxJabbers	Rx packets longer than 1522 bytes w/ either CRC errors, alignment errors, or symbol errors (depends on max packet size setting).
0x6	RxSymbolError	Rx packets w/ invalid data symbol and legal packet size.
0x7	RxCRCError	Rx packets within (64,1522) bytes w/ an integral number of bytes and a bad CRC (upper limit depends on max packet size setting).
0x8	RxAlignmentError	Rx packets within (64,1522) bytes w/ a non-integral number of bytes and a bad CRC (upper limit depends on max packet size setting).
0x9	RxControl8808Pkts	Number of MAC control frames received by a port with 88-08h in EtherType field.
0xA	RxPausePkts	Number of PAUSE frames received by a port. PAUSE frame is qualified with EtherType (88-08h), DA, control opcode (00-01), data length (64B min), and a valid CRC.
0xB	RxBroadcast	Rx good broadcast packets (not including error broadcast packets or valid multicast packets).
0xC	RxMulticast	Rx good multicast packets (not including MAC control frames, error multicast packets or valid broadcast packets).
0xD	RxUnicast	Rx good unicast packets.
0xE	Rx64Octets	Total Rx packets (bad packets included) that were 64 octets in length.

TABLE 4-93: PORT 1'S MIB COUNTERS INDIRECT MEMORY OFFSETS (CONTINUED)

Offset	Counter Name	Description
0xF	Rx65to127Octets	Total Rx packets (bad packets included) that are between 65 and 127 octets in length.
0x10	Rx128to255Octets	Total Rx packets (bad packets included) that are between 128 and 255 octets in length.
0x11	Rx256to511Octets	Total Rx packets (bad packets included) that are between 256 and 511 octets in length.
0x12	Rx512to1023Octets	Total Rx packets (bad packets included) that are between 512 and 1023 octets in length.
0x13	Rx1024to1522Octets	Total Rx packets (bad packets included) that are between 1024 and 1522 octets in length (upper limit depends on max packet size setting).
0x14	TxByte	Tx good octet count, including PAUSE packets.
0x15	Reserved	Reserved. Do not write to this register.
0x16	TxLateCollision	The number of times a collision is detected later than 512 bit-times into the Tx of a packet.
0x17	TxPausePkts	Number of PAUSE frames transmitted by a port.
0x18	TxBroadcastPkts	Tx good broadcast packets (not including error broadcast or valid multicast packets).
0x19	TxMulticastPkts	Tx good multicast packets (not including error multicast packets or valid broadcast packets).
0x1A	TxUnicastPkts	Tx good unicast packets.
0x1B	TxDeferred	Tx packets by a port for which the 1st Tx attempt is delayed due to the busy medium.
0x1C	TxTotalCollision	Tx total collision, half-duplex only.
0x1D	TxExcessiveCollision	A count of frames for which Tx fails due to excessive collisions.
0x1E	TxSingleCollision	Successfully Tx frames on a port for which Tx is inhibited by exactly one collision.
0x1F	TxMultipleCollision	Successfully Tx frames on a port for which Tx is inhibited by more than one collision.

Example: MIB Counter Read (read “Rx64Octets” counter at indirect address offset 0x0E)

Write to reg. IACR with 0x1C0E (set indirect address and trigger a read MIB counters operation)

Then:

Read reg. IADR5 (MIB counter value 31-16) // If bit 31 = 1, there was a counter overflow

// If bit 30 = 0, restart (reread) from this register

Read reg. IADR4 (MIB counter value 15-0)

Additional MIB Information

In the heaviest condition, the byte counter will overflow in two minutes. It is recommended that the software read all the counters at least every 30 seconds.

MIB counters are designed as “read clear”. That is, these counters will be cleared after they are read.

5.0 OPERATIONAL CHARACTERISTICS

5.1 Absolute Maximum Ratings*

Supply Voltage (V _{DDATX} , V _{DDARX} , V _{DDIO}).....	–0.5V to +4.0V
Input Voltage (all inputs).....	–0.5V to +5.0V
Output Voltage (all outputs).....	–0.5V to +4.0V
Lead Temperature (soldering, 10s).....	+270°C
Storage Temperature (T _S).....	–55°C to +150°C

*Exceeding the absolute maximum rating may damage the device. Stresses greater than those listed in the table above may cause permanent damage to the device. Operation of the device at these or any other conditions above those specified in the operating sections of this specification is not implied. Maximum conditions for extended periods may affect reliability. Unused inputs must always be tied to an appropriate logic voltage level.

5.2 Operating Ratings**

Supply Voltage (V _{DDATX} , V _{DDARX} , V _{DDIO}).....	+3.1V to +3.5V
Ambient Operating Temperature for Commercial Options (T _A).....	0°C to +70°C
Maximum Junction Temperature (T _J).....	+125°C
Thermal Resistance (Note 5-1) (Θ _{JA}).....	+42.91°C/W
Thermal Resistance (Note 5-1) (Θ _{JC}).....	+19.6°C/W

**The device is not guaranteed to function outside its operating ratings. Unused inputs must always be tied to an appropriate logic voltage level (Ground to VDD).

Note 5-1 No heat spreader (HS) in this package. The Θ_{JC}/Θ_{JA} is under air velocity 0 m/s.

Note: Do not drive input signals without power supplied to the device.

KSZ8841-PMQL

6.0 ELECTRICAL CHARACTERISTICS

Specification is for packaged product only. Single port's transformer consumes an additional 45 mA @ 3.3V for 100BASE-TX and 70 mA @ 3.3V for 10BASE-T.

TABLE 6-1: ELECTRICAL CHARACTERISTICS

Parameters	Symbol	Min.	Typ.	Max.	Units	Note
Supply Current for 100BASE-TX Operation (All Ports @ 100% Utilization)						
100BASE-TX (analog core + digital core + transceiver + digital I/O)	I_{DDXIO}	—	100	—	mA	$V_{DDATX}, V_{DDARX}, V_{DDIO} = 3.3V$
Supply Current for 10BASE-T Operation (All Ports @ 100% Utilization)						
10BASE-T (analog core + digital core + transceiver + digital I/O)	I_{DDXIO}	—	85	—	mA	$V_{DDATX}, V_{DDARX}, V_{DDIO} = 3.3V$
CMOS Inputs						
Input High Voltage	V_{IH}	2.0	—	—	V	—
Input Low Voltage	V_{IL}	—	—	0.8	V	—
Input Current	I_{IN}	-10	—	10	μA	$V_{IN} = GND \sim V_{DDIO}$
CMOS Outputs						
Output High Voltage	V_{OH}	2.4	—	—	V	$I_{OH} = -8 \text{ mA}$
Output Low Voltage	V_{OL}	—	—	0.4	V	$I_{OL} = 8 \text{ mA}$
Output Tri-State Leakage	I_{OZ}	—	—	10	μA	—
100BASE-TX Transmit (measured differentially after 1:1 transformer) $V_{DDATX} = 3.3V$ only						
Peak Differential Output Voltage	V_O	0.95	—	1.05	V	100 Ω termination on the differential output.
Output Voltage Imbalance	V_{IMB}	—	—	2	%	100 Ω termination on the differential output.
Rise/Fall Time	t_r/t_f	3	—	5	ns	—
Rise/Fall Time Imbalance	—	0	—	0.5	ns	—
Duty Cycle Distortion	—	—	—	± 0.5	ns	—
Overshoot	—	—	—	5	%	—
Reference Voltage of I_{SET}	V_{SET}	—	0.5	—	V	—
Output Jitter	—	—	0.7	1.4	ns	Peak-to-peak
10BASE-T Receive						
Squelch Threshold	V_{SQ}	—	400	—	mV	5 MHz square wave
10BASE-T Transmit (measured differentially after 1:1 transformer) $V_{DDATX} = 3.3V$ only						
Peak Differential Output Voltage	V_P	—	2.4	—	V	100 Ω termination on the differential output.
Jitter Added	—	—	1.8	± 3.5	ns	100 Ω termination on the differential output

7.0 TIMING SPECIFICATIONS

For PCI timing, please refer to PCI Specification version 2.2.

7.1 EEPROM Timing

FIGURE 7-1: EEPROM READ CYCLE TIMING DIAGRAM




TABLE 7-1: EEPROM TIMING PARAMETERS

Symbol	Parameter	Min.	Typ.	Max.	Units
t_{cyc}	Clock cycle	—	4000	—	ns
t_s	Setup time	20	—	—	ns
t_h	Hold time	20	—	—	ns

7.2 Auto-Negotiation Timing

FIGURE 7-2: AUTO-NEGOTIATION TIMING




TABLE 7-2: AUTO-NEGOTIATION TIMING PARAMETERS

Symbol	Parameter	Min.	Typ.	Max.	Units
t_{BTB}	FLP burst to FLP burst	8	16	24	ms
t_{FLPW}	FLP burst width	—	2	—	ms
t_{PW}	Clock/Data pulse width	—	100	—	ns
t_{CTD}	Clock pulse to data pulse	55.5	64	69.5	μ s
t_{CTC}	Clock pulse to clock pulse	111	128	139	μ s
—	Number of Clock/Data pulses per burst	17	—	33	—

7.3 Reset Timing

As long as the stable supply voltages to reset High timing (minimum of 10 ms) are met, there is no power-sequencing requirement for the KSZ8841-PMQL supply voltages (3.3V).

The reset timing requirement is summarized in [Figure 7-3](#) and [Table 7-3](#).

FIGURE 7-3: RESET TIMING




TABLE 7-3: RESET TIMING PARAMETERS

Parameter	Description	Min.	Typ.	Max.	Units
t_{SR}	Stable supply voltages to reset high	10	—	—	ms

8.0 SELECTION OF ISOLATION TRANSFORMERS

A 1:1 isolation transformer is required at the line interface. An isolation transformer with integrated common-mode choke is recommended for exceeding FCC requirements.

Table 8-1 lists recommended transformer characteristics.

TABLE 8-1: TRANSFORMER SELECTION CRITERIA

Parameter	Value	Test Conditions
Turns Ratio	1 CT : 1 CT	—
Open-Circuit Inductance (min.)	350 μ H	100 mV, 100 kHz, 8 mA
Leakage Inductance (max.)	0.4 μ H	1 MHz (min.)
Interwinding Capacitance (max.)	12 pF	—
D.C. Resistance (max.)	0.9 Ω	—
Insertion Loss (max.)	1.0 dB	0 MHz to 65 MHz
HIPOT (min.)	1500 V _{RMS}	—

TABLE 8-2: QUALIFIED SINGLE-PORT MAGNETICS

Manufacturer	Part Number	Auto MDI-X
Bel Fuse	S558-5999-U7	Yes
Delta	LF8505	Yes
LanKom	LF-H41S	Yes
Pulse	H1102	Yes
Pulse (Low Cost)	H1260	Yes
Transpower	HB726	Yes
TDK (Mag Jack)	TLA-6T718	Yes

TABLE 8-3: TYPICAL REFERENCE CRYSTAL CHARACTERISTICS

Characteristic	Value
Frequency	25 MHz
Frequency Tolerance (max.)	± 50 ppm
Load Capacitance (max.)	20 pF
Series Resistance	25 Ω

9.0 PACKAGE OUTLINE

9.1 Package Marking Information

128-Lead PQFP*

MICREL
XXXXXXX
XXXX
YYWWA7
XXXXXXYYWWNNN
● YYWWNNN

Example

MICREL
KSZ8841
PMQL
1936A7
G00001936287
● 1936287

Legend:	XX...X	Product code or customer-specific information
	Y	Year code (last digit of calendar year)
	YY	Year code (last 2 digits of calendar year)
	WW	Week code (week of January 1 is week '01')
	NNN	Alphanumeric traceability code
	(e3)	Pb-free JEDEC [®] designator for Matte Tin (Sn)
	*	This package is Pb-free. The Pb-free JEDEC designator ((e3)) can be found on the outer packaging for this package.
	●, ▲, ▼	Pin one index is identified by a dot, delta up, or delta down (triangle mark).

Note:	In the event the full Microchip part number cannot be marked on one line, it will be carried over to the next line, thus limiting the number of available characters for customer-specific information. Package may or may not include the corporate logo.
	Underbar (_) and/or Overbar (¯) symbol may not be to scale.



KSZ8841-PMQL

FIGURE 9-1: 128-LEAD PQFP 14 MM X 20 MM PACKAGE OUTLINE AND RECOMMENDED LAND PATTERN

TITLE

128 LEAD PQFP 14x20mm PACKAGE OUTLINE & RECOMMENDED LAND PATTERN

DRAWING #	PQFP14x20-128LD-PL-1	UNIT	MM [INCHES]
-----------	----------------------	------	-------------




SYMBOL	MILLIMETER			INCH		
	MIN.	NOM.	MAX.	MIN.	NOM.	MAX.
A	—	—	3.40	—	—	0.134
A1	0.25	—	—	0.010	—	—
A2	2.50	2.72	2.90	0.098	0.107	0.114
D	23.20	BASIC	—	0.913	BASIC	—
D1	20.00	BASIC	—	0.787	BASIC	—
E	17.20	BASIC	—	0.677	BASIC	—
E1	14.00	BASIC	—	0.551	BASIC	—
R1	0.13	—	0.30	0.005	—	0.012
R2	0.13	—	—	0.005	—	—
θ	0°	—	7°	0°	—	7°
θ1	0°	—	—	0°	—	—
θ2, θ3	15°	REF	—	15°	REF	—

SYMBOL	MILLIMETER			INCH		
	MIN.	NOM.	MAX.	MIN.	NOM.	MAX.
c	0.11	0.15	0.23	0.004	0.006	0.009
L	0.73	0.88	1.03	0.029	0.035	0.041
L1	1.60	REF	—	0.063	REF	—
S	0.20	—	—	0.008	—	—
b	0.170	0.200	0.270	0.007	0.008	0.011
e	0.50	BSC.	—	0.020	BSC	—
D2	18.50	—	—	0.728	—	—
E2	12.50	—	—	0.492	—	—
TOLERANCES OF FORM AND POSITION						
aaa	0.20	—	—	0.008	—	—
bbb	0.20	—	—	0.008	—	—
ccc	0.08	—	—	0.003	—	—
ddd	0.08	—	—	0.003	—	—

CONTROL DIMENSIONS ARE IN MILLIMETERS.

NOTES :

1. DIMENSION D1 AND E1 DO NOT INCLUDE MOLD PROTRUSION. ALLOWABLE PROTRUSION IS 0.25mm PER SIDE. DIMENSIONS D1 AND E1 DO INCLUDE MOLD MISMATCH AND ARE DETERMINED AT DATUM PLANE $\square-H$.
2. DIMENSION b DOES NOT INCLUDE DAMBAR PROTRUSION. ALLOWABLE DAMBAR PROTRUSION SHALL BE 0.08mm TOTAL IN EXCESS OF THE b DIMENSION AT MAXIMUM MATERIAL CONDITION. DAMBAR CANNOT BE LOCATED ON THE LOWER RADIUS OR THE LEAD FOOT.
3. THE DIAGRAMS DO NOT REPRESENT THE ACTUAL PIN COUNT.
4. ALL UNITS IN mm. TOLERANCE ± 0.05 IF NOT NOTED.



Note: For the most current package drawings, please see the Microchip Packaging Specification located at <http://www.microchip.com/packaging>.

APPENDIX A: DATA SHEET REVISION HISTORY

TABLE A-1: REVISION HISTORY

Revision	Section/Figure/Entry	Correction
DS00003284A (10-28-19)	—	Converted Micrel data sheet KSZ8841-PMQL to Microchip DS00003284A. Minor text changes throughout.

THE MICROCHIP WEB SITE

Microchip provides online support via our WWW site at www.microchip.com. This web site is used as a means to make files and information easily available to customers. Accessible by using your favorite Internet browser, the web site contains the following information:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQ), technical support requests, online discussion groups, Microchip consultant program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

CUSTOMER CHANGE NOTIFICATION SERVICE

Microchip's customer notification service helps keep customers current on Microchip products. Subscribers will receive e-mail notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, access the Microchip web site at www.microchip.com. Under "Support", click on "Customer Change Notification" and follow the registration instructions.

CUSTOMER SUPPORT

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Field Application Engineer (FAE)
- Technical Support

Customers should contact their distributor, representative or field application engineer (FAE) for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in the back of this document.

Technical support is available through the web site at: <http://microchip.com/support>

PRODUCT IDENTIFICATION SYSTEM

To order or obtain information, e.g., on pricing or delivery, refer to the factory or the listed sales office.

<u>PART NO.</u>	<u>-XX</u>	<u>X</u>	<u>X</u>	<u>X</u>	<u>[X]</u>	<u>[-XX]</u>
Device	Bus Design	Interface	Package	Supply Voltage	Temperature	Media Type
Device:	KSZ8841: Single-Port Ethernet MAC Controller with PCI Interface					
Bus Design:	P = PCI					
Interface:	M = Management Interface					
Package:	Q = 128-lead PQFP					
Supply Voltage:	L = Single 3.3V Power Supply Supported with Internal 1.8V LDO					
Temperature:	<blank> = 0°C to +70°C (Commercial) I = -40°C to +85°C (Industrial)					
Media Type:	<blank> = 66/Tray (PQFP option)					

Examples:

- KSZ8841-PMQL
PCI Management Interface
128-lead PQFP, Single 3.3V Power Supply
Commercial Temperature Range
66/Tray
- KSZ8841-PMQLI
PCI Management Interface
128-lead PQFP, Single 3.3V Power Supply
Industrial Temperature Range
66/Tray

Note 1: Tape and Reel identifier only appears in the catalog part number description. This identifier is used for ordering purposes and is not printed on the device package. Check with your Microchip Sales Office for package availability with the Tape and Reel option.

KSZ8841-PMQL

NOTES:

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specification contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is one of the most secure families of its kind on the market today, when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods used to breach the code protection feature. All of these methods, to our knowledge, require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Most likely, the person doing so is engaged in theft of intellectual property.
- Microchip is willing to work with the customer who is concerned about the integrity of their code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of their code. Code protection does not mean that we are guaranteeing the product as "unbreakable."

Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

Information contained in this publication regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications. MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION, INCLUDING BUT NOT LIMITED TO ITS CONDITION, QUALITY, PERFORMANCE, MERCHANTABILITY OR FITNESS FOR PURPOSE. Microchip disclaims all liability arising from this information and its use. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AnyRate, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, chipKIT, chipKIT logo, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Klear, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PackeTime, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TempTrackr, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, FlashTec, Hyper Speed Control, HyperLight Load, IntelliMOS, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, Vite, WinPath, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, BlueSky, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, EtherGREEN, In-Circuit Serial Programming, ICSP, INICnet, Inter-Chip Connectivity, JitterBlocker, KlearNet, KlearNet logo, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICKit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, SAM-ICE, Serial Quad I/O, SMART-I.S., SQL, SuperSwitcher, SuperSwitcher II, Total Endurance, TSHARC, USBCheck, VariSense, ViewSpan, WiperLock, Wireless DNA, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2019, Microchip Technology Incorporated, All Rights Reserved.

ISBN: 978-1-5224-5200-3

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.

Worldwide Sales and Service

AMERICAS

Corporate Office
2355 West Chandler Blvd.
Chandler, AZ 85224-6199
Tel: 480-792-7200
Fax: 480-792-7277
Technical Support:
<http://www.microchip.com/support>
Web Address:
www.microchip.com

Atlanta
Duluth, GA
Tel: 678-957-9614
Fax: 678-957-1455

Austin, TX
Tel: 512-257-3370

Boston
Westborough, MA
Tel: 774-760-0087
Fax: 774-760-0088

Chicago
Itasca, IL
Tel: 630-285-0071
Fax: 630-285-0075

Dallas
Addison, TX
Tel: 972-818-7423
Fax: 972-818-2924

Detroit
Novi, MI
Tel: 248-848-4000

Houston, TX
Tel: 281-894-5983

Indianapolis
Noblesville, IN
Tel: 317-773-8323
Fax: 317-773-5453
Tel: 317-536-2380

Los Angeles
Mission Viejo, CA
Tel: 949-462-9523
Fax: 949-462-9608
Tel: 951-273-7800

Raleigh, NC
Tel: 919-844-7510

New York, NY
Tel: 631-435-6000

San Jose, CA
Tel: 408-735-9110
Tel: 408-436-4270

Canada - Toronto
Tel: 905-695-1980
Fax: 905-695-2078

ASIA/PACIFIC

Australia - Sydney
Tel: 61-2-9868-6733

China - Beijing
Tel: 86-10-8569-7000

China - Chengdu
Tel: 86-28-8665-5511

China - Chongqing
Tel: 86-23-8980-9588

China - Dongguan
Tel: 86-769-8702-9880

China - Guangzhou
Tel: 86-20-8755-8029

China - Hangzhou
Tel: 86-571-8792-8115

China - Hong Kong SAR
Tel: 852-2943-5100

China - Nanjing
Tel: 86-25-8473-2460

China - Qingdao
Tel: 86-532-8502-7355

China - Shanghai
Tel: 86-21-3326-8000

China - Shenyang
Tel: 86-24-2334-2829

China - Shenzhen
Tel: 86-755-8864-2200

China - Suzhou
Tel: 86-186-6233-1526

China - Wuhan
Tel: 86-27-5980-5300

China - Xian
Tel: 86-29-8833-7252

China - Xiamen
Tel: 86-592-2388138

China - Zhuhai
Tel: 86-756-3210040

ASIA/PACIFIC

India - Bangalore
Tel: 91-80-3090-4444

India - New Delhi
Tel: 91-11-4160-8631

India - Pune
Tel: 91-20-4121-0141

Japan - Osaka
Tel: 81-6-6152-7160

Japan - Tokyo
Tel: 81-3-6880-3770

Korea - Daegu
Tel: 82-53-744-4301

Korea - Seoul
Tel: 82-2-554-7200

Malaysia - Kuala Lumpur
Tel: 60-3-7651-7906

Malaysia - Penang
Tel: 60-4-227-8870

Philippines - Manila
Tel: 63-2-634-9065

Singapore
Tel: 65-6334-8870

Taiwan - Hsin Chu
Tel: 886-3-577-8366

Taiwan - Kaohsiung
Tel: 886-7-213-7830

Taiwan - Taipei
Tel: 886-2-2508-8600

Thailand - Bangkok
Tel: 66-2-694-1351

Vietnam - Ho Chi Minh
Tel: 84-28-5448-2100

EUROPE

Austria - Wels
Tel: 43-7242-2244-39
Fax: 43-7242-2244-393

Denmark - Copenhagen
Tel: 45-4450-2828
Fax: 45-4485-2829

Finland - Espoo
Tel: 358-9-4520-820

France - Paris
Tel: 33-1-69-53-63-20
Fax: 33-1-69-30-90-79

Germany - Garching
Tel: 49-8931-9700

Germany - Haan
Tel: 49-2129-3766400

Germany - Heilbronn
Tel: 49-7131-72400

Germany - Karlsruhe
Tel: 49-721-625370

Germany - Munich
Tel: 49-89-627-144-0
Fax: 49-89-627-144-44

Germany - Rosenheim
Tel: 49-8031-354-560

Israel - Ra'anana
Tel: 972-9-744-7705

Italy - Milan
Tel: 39-0331-742611
Fax: 39-0331-466781

Italy - Padova
Tel: 39-049-7625286

Netherlands - Drunen
Tel: 31-416-690399
Fax: 31-416-690340

Norway - Trondheim
Tel: 47-7288-4388

Poland - Warsaw
Tel: 48-22-3325737

Romania - Bucharest
Tel: 40-21-407-87-50

Spain - Madrid
Tel: 34-91-708-08-90
Fax: 34-91-708-08-91

Sweden - Gothenberg
Tel: 46-31-704-60-40

Sweden - Stockholm
Tel: 46-8-5090-4654

UK - Wokingham
Tel: 44-118-921-5800
Fax: 44-118-921-5820

Mouser Electronics

Authorized Distributor

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

[Microchip:](#)

[KSZ8841-16MQL](#) [KSZ8841-16MVL](#) [KSZ8841-PMQLI](#) [KSZ8841-32MVLI](#) [KSZ8841-32MQL](#) [KSZ8841-16MVLI](#)
[KSZ8841-32MVL](#) [KSZ8841-PMQL](#) [KSZ8841-16MVL-TR](#) [KSZ8841-16MVLI-TR](#) [KSZ8841-PMQLI-TR](#)