



CENTAUR™
KS8695PX
Integrated Multi-Port Gateway
Solution

Switch Description

Version 1.0

September 2003



CENTAUR™ KS8695PX Switch Description

Integrated Multi-Port PCI Gateway Solution

Revision History

Revision	Date	Summary of Changes
1.0	9/26/03	Initial Release



CENTAUR™ KS8695PX Switch Description

Integrated Multi-Port PCI Gateway Solution

This page is blank



TABLE OF CONTENTS

1.0 Overview	1
2.0 Functional Description	2
2.1 Physical Layer Transceiver	2
2.1.1 100BaseTX Transmit.....	2
2.1.2 100BaseTX Receive.....	2
2.1.3 Scrambler/De-scrambler (100BaseTX only)	3
2.1.4 100BaseFX operation.....	3
2.1.5 100BaseFX Signal Detection.....	3
2.1.6 100BaseFX Far End Fault.....	3
2.1.7 10BaseT Transmit.....	3
2.1.8 10BaseT Receive.....	3
2.1.9 Power Management.....	4
2.1.10 MDI / MDI-X auto crossover	4
2.1.11 Auto Negotiation	4
2.2 Switch Core.....	6
2.2.1 Address Look Up.....	6
2.2.2 Learning.....	6
2.2.3 Migration.....	6
2.2.4 Aging	6
2.2.5 Forwarding	7
2.2.6 Switching Engine.....	10
2.2.7 MAC operation	10
3.0 KS8695PX Switch Advanced Functionality	13
3.1 Spanning Tree Support:.....	13
3.2 Special Tagging Mode	14
3.3 IGMP Support	16
3.4 Port Mirroring Support.....	16
3.5 VLAN support.....	17
3.6 Rate Limit Support	18



LIST OF TABLES AND FIGURES

Table 1-Special Tagging Mode Format.....	14
Table 2-STPID Egress Rules (Processor to Switch Port 5)	15
Table 3-STPID Egress Rules (Switch to Processor)	16
Table 4-FID+DA look up in the VLAN mode	17
Table 5-FID+SA look up in the VLAN mode.....	18
 Figure 1-Auto Negotiation	 5
Figure 2 DA look up flowchart, stage 1	8
Figure 3 DA resolution flowchart, stage 2	9



1.0 Overview

The KS8695PX Switch contains five 10/100 physical layer transceivers and five MAC (Media Access Control) units with an integrated layer 2 managed switch. On the media side, the KS8695PX Switch supports IEEE 802.3 10BaseT, 100BaseTX on all ports, and 100BaseFX on WAN port and one LAN port. Physical signal transmission and reception are enhanced through the use of patented analog circuitry that makes the design more efficient and allows for lower power consumption and smaller chip die size.

Feature Highlights :

- 5 port 10/100 Integrated switch with 1 WAN and 4 LAN Physical Layer Transceivers
- 10Base-T, 100Base-TX and 100Base-FX modes (FX on WAN port and one LAN port)
- 10Mbps, 100Mbps modes of operations for both full and half duplex
- Automatic MDI / MDI-X crossover for plug-and-play
- 16Kx32 on-chip SRAM for frame buffering.
- 1.4Gbps on-chip memory bandwidth for wire-speed frame switching.
- Supports 802.1q Tag-based VLAN and port-based VLAN
- Supports 802.1p based priority, DiffServ priority and port-based priority.
- Integrated address Look-Up engine, supports 1 K absolute MAC addresses.
- Automatic address learning, address aging and address migration.
- Broadcast storm protection.
- Full duplex IEEE 802.3x flow control.
- Half duplex back pressure flow control.
- Spanning tree protocol support
- Programmable rate limiting 0 to 100Mbps, ingress & egress port, rate options for high & low priority, per port basis
- Extensive MIB counter management support
- IGMP v1/v2 Snooping for multicast packet filtering
- Port mirroring / monitoring / sniffing
- Optimization for fiber-to-copper media conversion



2.0 Functional Description

2.1 Physical Layer Transceiver

2.1.1 100BaseTX Transmit

The 100BaseTX transmit function performs parallel to serial conversion, 4B/5B coding, scrambling, NRZ to NRZI conversion, MLT3 encoding and transmission. The circuit starts with a parallel to serial conversion, which converts the MII data from the MAC 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. The output current is set by an external 1% 3.01 K Ω resistor for the 1:1 transformer ratio. It 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 10BaseT output is also incorporated into the 100BaseTX transmitter.

2.1.2 100BaseTX Receive

The 100BaseTX 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. Since the amplitude loss and phase distortion is a function of the length of the cable, the equalizer has to adjust its characteristics to optimize the performance. In this design, the variable equalizer will make an initial estimation based on comparisons of incoming signal strength against some known cable characteristics, then it tunes itself for optimization. This is an ongoing process and can self adjust against environmental changes such as temperature variations.

The equalized signal then goes through a DC restoration and data conversion block. The DC restoration circuit is used to compensate for the effect of base line wander and 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 125 MHz clock from the edges of the NRZI signal. This recovered clock is then used to convert the NRZI signal into the NRZ format. The signal is then sent through the de-scrambler followed by the 4B/5B decoder. Finally, the NRZ serial data is converted to the MII format and provided as the input data to the MAC.



2.1.3 Scrambler/De-scrambler (100BaseTX only)

The purpose of the scrambler is to spread the power spectrum of the signal in order to reduce EMI and baseline wander. The data is scrambled through the use of an 11-bit wide linear feedback shift register (LFSR). This can generate a 2047-bit non-repetitive sequence. The receiver will then de-scramble the incoming data stream with the same sequence at the transmitter.

2.1.4 100BaseFX operation

100BaseFX operation is very similar to 100BaseTX operation with the differences being that the scrambler / de-scrambler and MLT3 encoder / decoder are bypassed on transmission and reception. In this mode the auto negotiation feature is bypassed since there is no standard that supports fiber auto negotiation.

2.1.5 100BaseFX Signal Detection

The physical port runs in 100BaseFX mode if FXSDx >.6V for WAN ports and one LAN port only. This signal is internally referenced to 1.25V. The fiber module interface should be set by a voltage divider such that FXSDx 'H' is above this 1.25V reference, indicating signal detect, and FXSDx 'L' is below the 1.25V reference to indicate no signal. When FXSDx is below .6V then 100BaseFX mode is disabled. Since there is no auto-negotiation for 100BaseFX mode, WAN port and one LAN port must be forced to either full or half duplex. Note that strap in options exist to set duplex mode for WAN port and one LAN port.

2.1.6 100BaseFX Far End Fault

Far end fault occurs when the signal detection is logically false from the receive fiber module. When this occurs, the transmission side signals the other end of the link by sending 84 1's followed by a zero in the idle period between frames. The far end fault may be disabled through register settings.

2.1.7 10BaseT Transmit

The output 10BaseT driver is incorporated into the 100BaseT driver to allow transmission with the same magnetics. They are internally wave-shaped and pre-emphasized into outputs with a typical 2.3 V amplitude. The harmonic contents are at least 27 dB below the fundamental when driven by an all-ones Manchester-encoded signal.

2.1.8 10BaseT Receive

On the receive side, input buffer and level detecting squelch circuits are employed. A differential input receiver circuit and a 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 in order to prevent noises 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 KS8695PX Switch decodes a data frame. The receiver clock is maintained active during idle periods in between data reception.

2.1.9 Power Management

The KS8695PX Switch features a per port power down mode. To save power the user can power down ports that are not in use by setting port control registers or MII control registers. In addition, it also supports full chip power down mode. When activated, the entire chip will be shut down.

2.1.10 MDI / MDI-X auto crossover

The KS8695PX Switch supports MDI / MDI-X auto crossover. This facilitates the use of either a straight connection CAT-5 cable or a crossover CAT-5 cable. The auto-sense function will detect remote transmit and receive pairs, and correctly assign the transmit and receive pairs from the Micrel device. This can be highly useful when end users are unaware of cable types and can also save on an additional uplink configuration connection. The auto crossover feature may be disabled through the port control registers.

2.1.11 Auto Negotiation

The KS8695PX Switch conforms to the auto negotiation protocol as described by the 802.3 committee. Auto negotiation allows UTP (Unshielded Twisted Pair) 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 KS8695PX Switch is forced to bypass auto negotiation, then 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 flow for the link set up is depicted in Figure 1.

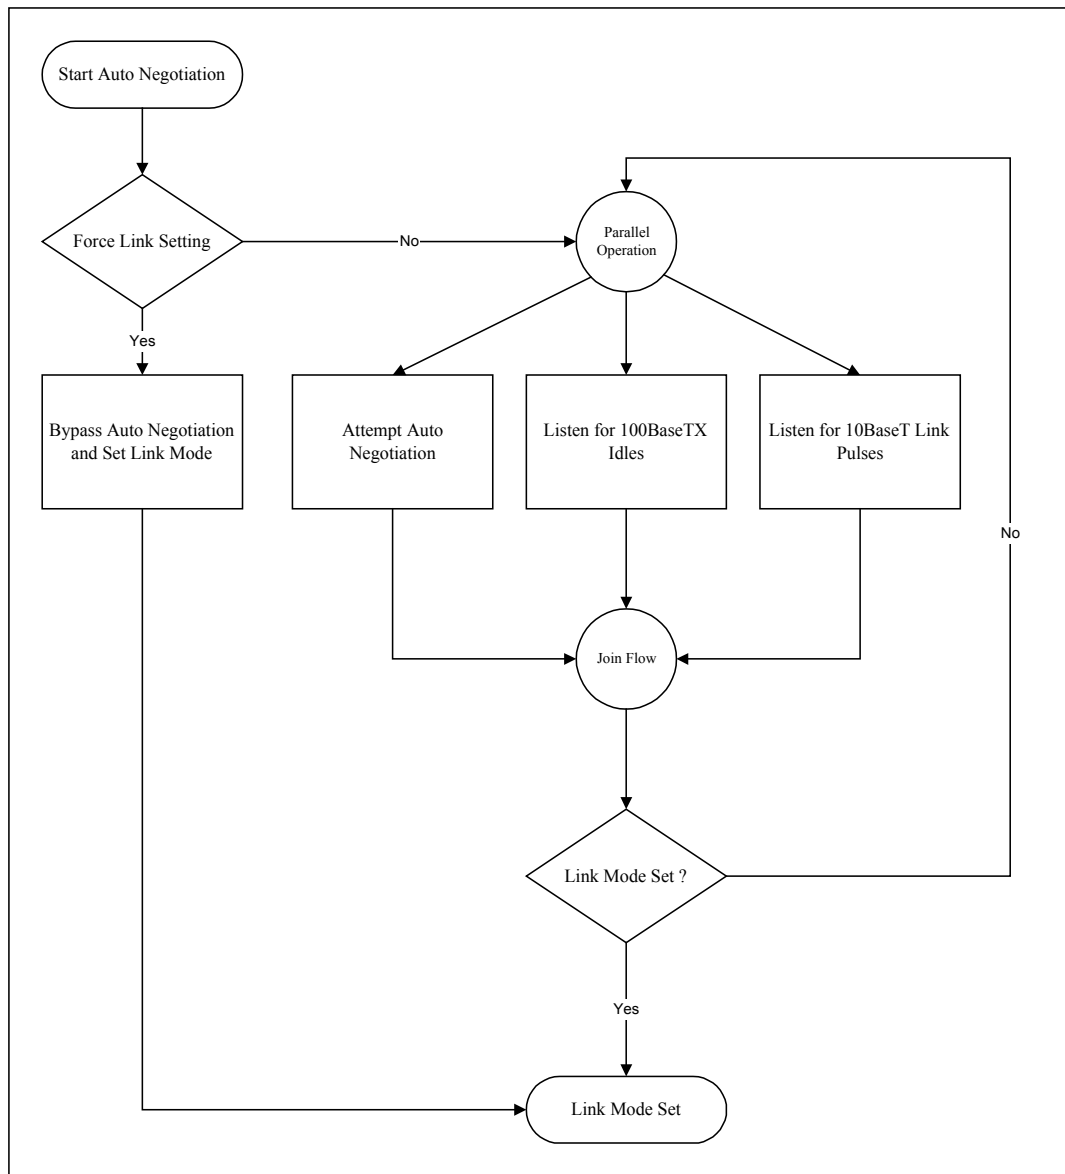


Figure 1-Auto Negotiation



2.2 Switch Core

2.2.1 Address Look Up

The internal look up table stores MAC addresses and their associated information. It contains a 1K unicast address table plus switching information. The KS8695PX Switch is guaranteed to learn 1K addresses and distinguishes itself from hash-based look up tables which, depending on the operating environment and probabilities, may not guarantee the absolute number of addresses it can learn.

2.2.2 Learning

The internal look up engine will update its table with a new entry if the following conditions are met:

- (1). The received packet's SA does not exist in the look up table.
- (2). The received packet is good; the packet has no receiving errors, and is of legal length.

The look up engine will insert the qualified SA into the table, along with the port number and time stamp. If the table is full, the last entry of the table will be deleted to make room for the new entry.

2.2.3 Migration

The internal look up engine also monitors whether a station has moved. If so, it will update the table accordingly. Migration happens when the following conditions are met:

- (1). The received packet's SA is in the table but the associated source port information is different.
- (2). The received packet is good; the packet has no receiving errors, and is of legal length.

The look up engine will update the existing record in the table with the new source port information.

2.2.4 Aging

The look up engine will update the time stamp information of a record whenever the corresponding SA appears. The time stamp is used in the aging process. If a record is not updated for a period of time, the look up engine will remove the



record from the table. The look up engine constantly performs the aging process and will continuously remove aging records. The aging period is 300 ± 75 seconds.

2.2.5 Forwarding

The KS8695PX Switch will forward packets using an algorithm that is depicted in the following flowcharts. Figure 2 shows stage one of the forwarding algorithm where the search engine looks up the VLAN ID, static table, and dynamic table for the destination address, and comes up with "port to forward 1" (PTF1). PTF1 is then further modified by the Spanning Tree, IGMP snooping, port mirroring, and port VLAN processes to come up with "port to forward 2" (PTF2) as shown in Figure 3. This is where the packet will be sent.

Figure 2 DA look up flowchart, stage 1

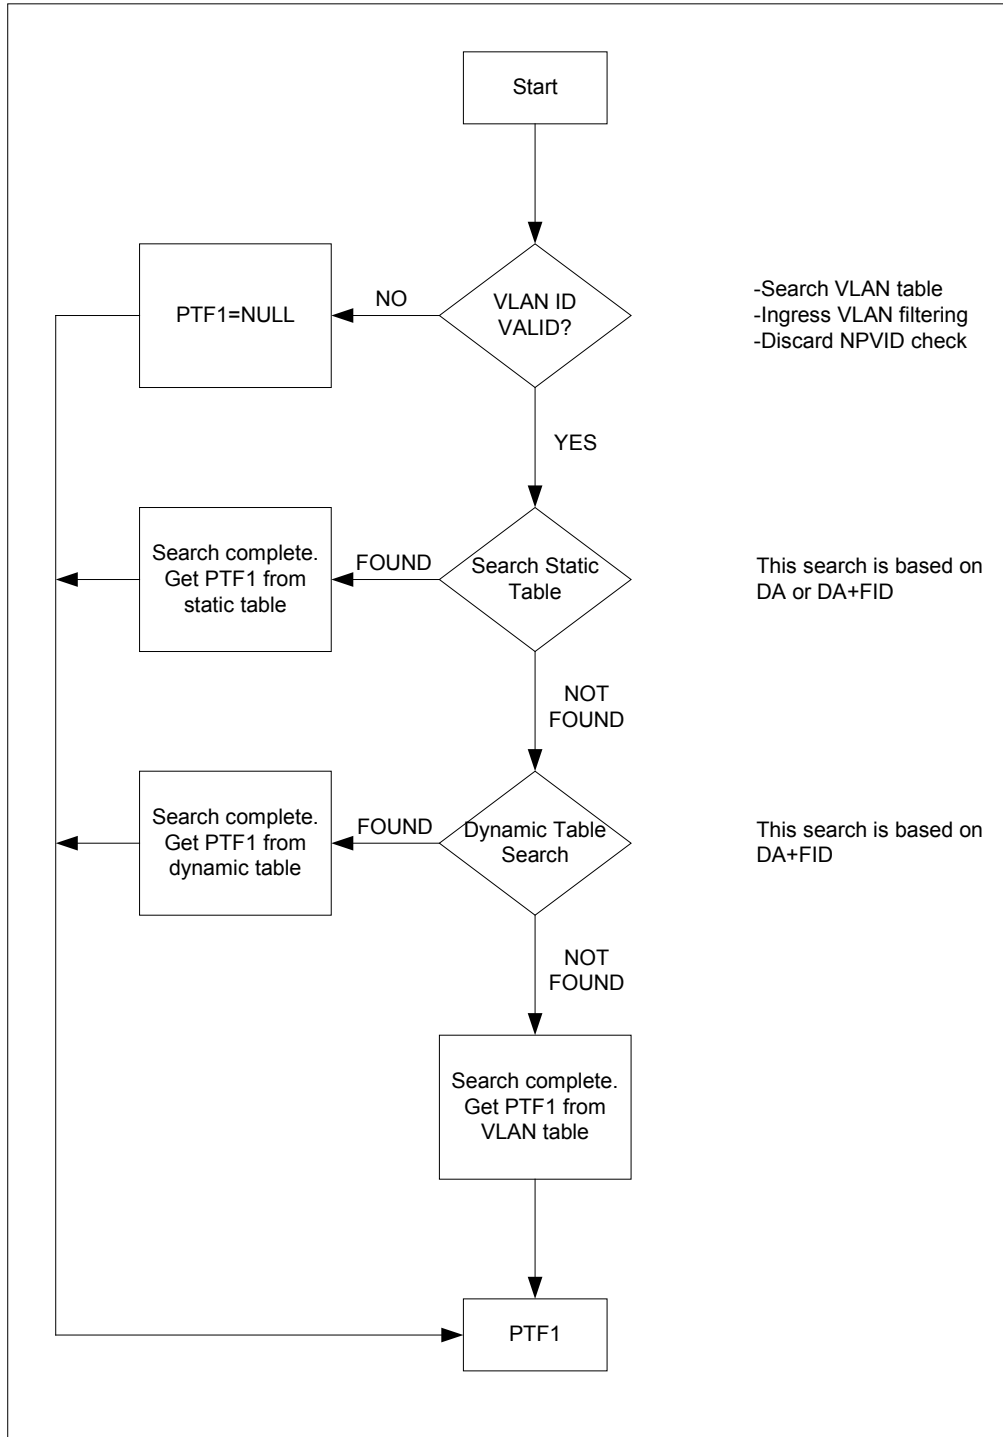
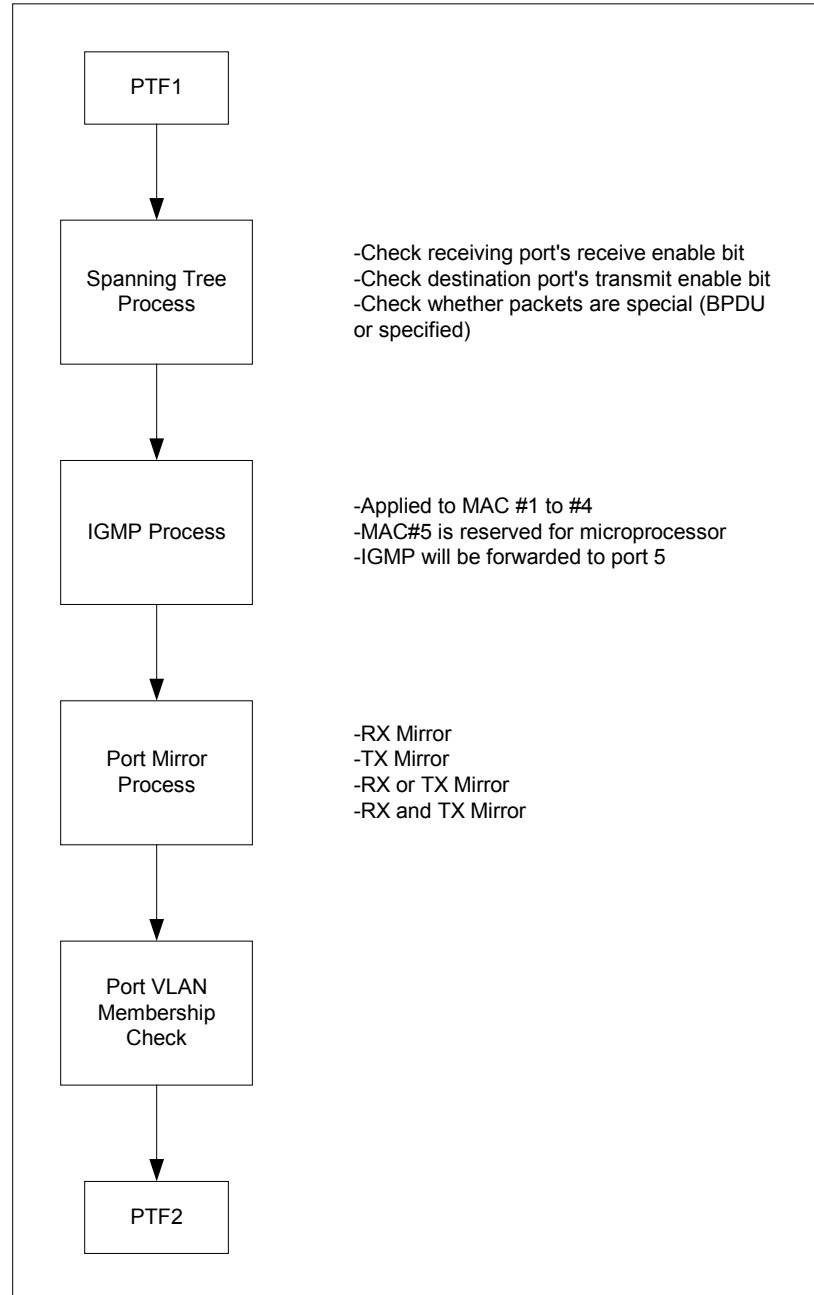


Figure 3 DA resolution flowchart, stage 2





The KS8695PX Switch will not forward the following packets:

- (1). Error packets. These include framing errors, FCS errors, alignment errors, and illegal size packet errors.
- (2). 802.3x pause frames. The KS8695PX Switch will intercept these packets and perform the appropriate actions.
- (3). "Local" packets. Based on destination address (DA) look up. If the destination port from the look up table matches the port where the packet was from, the packet is defined as "local".

2.2.6 Switching Engine

The KS8695PX Switch features a high performance switching engine to move data to and from the MAC's, packet buffers. It operates in store and forward mode, while the efficient switching mechanism reduces overall latency.

The KS8695PX Switch has a 64kB internal frame buffer. This resource is shared between all five ports. In one mode, ports are allowed to use any free buffers in the buffer pool. In the second mode, each port is only allowed to use 1/5 of the total buffer pool. There are a total of 512 buffers available. Each buffer sized at 128B.

2.2.7 MAC operation

The KS8695PX Switch strictly abides by IEEE 802.3 standards to maximize compatibility.

Inter Packet Gap (IPG)

If a frame is successfully transmitted, the 96 bit time IPG is measured between the two consecutive MTXEN. If the current packet is experiencing collision, the 96 bit time IPG is measured from MCRS and the next MTXEN.

2.2.7.1 Backoff Algorithm

The KS8695PX Switch implements the IEEE Std 802.3 binary exponential back-off algorithm, and optional "aggressive mode" back off. After 16 collisions, the packet will be optionally dropped depending on the register configuration.

2.2.7.2 Late Collision



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

2.2.7.3 Illegal Frames

The KS8695PX Switch discards frames less than 64 bytes and can be programmed to accept frames up to 1536 bytes. For special applications, the KS8695PX Switch can also be programmed to accept frames up to 1916 bytes. Since the KS8695PX Switch supports VLAN tags, the maximum sizing is adjusted when these tags are present.

2.2.7.4 Flow Control

The KS8695PX Switch supports standard 802.3x flow control frames on both transmit and receive sides.

On the receive side, if the KS8695PX Switch receives a pause control frame, the KS8695PX Switch 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 (being flow controlled), only flow control packets from the KS8695PX Switch will be transmitted.

On the transmit side, the KS8695PX Switch has intelligent and efficient ways to determine when to invoke flow control. The flow control is based on availability of the system resources, including available buffers, available transmit queues and available receive queues.

The KS8695PX Switch will flow control a port, which just received a packet, if the destination port resource is being used up. The KS8695PX Switch will issue a flow control frame (XOFF), containing the maximum pause time defined in IEEE standard 802.3x. Once the resource is freed up, the KS8695PX Switch will send out the other flow control frame (XON) 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 activated and deactivated too many times.

The KS8695PX Switch will flow control all ports if the receive queue becomes full.

2.2.7.5 Half Duplex Back Pressure

A half duplex back pressure option (Note: not in 802.3 standards) is also provided. The activation and deactivation conditions are the same as the above



in full duplex mode. If back pressure is required, the KS8695PX Switch will send preambles to defer the other stations' transmission (carrier sense deference). To avoid jabber and excessive deference defined in 802.3 standard, after a certain time it will discontinue the carrier sense but it will raise the carrier sense quickly. This short silent time (no carrier sense) is to prevent other stations from sending out packets and keeps other stations in carrier sense deferred state. If the port has packets to send during a back pressure situation, the carrier sense type back pressure will be interrupted and those packets will be transmitted instead. If there are no more packets to send, carrier sense type back pressure will be active again until switch resources free up. If a collision occurs, the binary exponential back-off algorithm is skipped and carrier sense is generated immediately, reducing the chance of further colliding and maintaining carrier sense to prevent reception of packets.

To ensure no packet loss in 10 Base T or 100 Base TX half duplex modes, the user must enable the following:

1. Aggressive Backoff (Switch Engine Control 0 Register, bit 13)
2. No Excessive collision drop (Switch Engine Control 0 Register, bit 8)
3. Back Pressure (Switch Engine Control 0 Register, bit 10)

These bits are not set as the default because this is not the IEEE standard.

2.2.7.6 Broadcast Storm Protection

The KS8695PX Switch has an intelligent option to protect the switch system from receiving too many broadcast packets. Broadcast packets will be forwarded to all ports except the source port, and thus use too many switch resources (bandwidth and available space in transmit queues). The KS8695PX Switch has the option to include "multicast packets" for storm control. The broadcast storm rate parameters are programmed globally, and can be enabled or disabled on a per port basis. The rate is based on a 50ms interval for 100BT and a 500 ms interval for 10BT. At the beginning of each interval, the counter is cleared to zero, and the rate limit mechanism starts to count the number of bytes during the interval. The rate definition is described in "Switch Engine Control 1 Register". The default setting is 0x4A, which is 74 decimal. This is equal to a rate of 1 %, calculated as follows:

$$148,800 \text{ frames/sec} * 50 \text{ ms/interval} * 1\% = 74 \text{ frames/interval (approx.)} = 0x4A$$



3.0 KS8695PX Switch Advanced Functionality

3.1 Spanning Tree Support:

To support spanning tree, port 5 is the designated port for the processor.

The other ports (port 1 – port 4) can be configured in one of the five spanning tree states via “transmit enable”, “receive enable” and “learning disable” register settings. The following description shows the port setting and software actions taken for each of the five spanning tree states.

Disable state: The port should not forward or receive any packets. Learning is disabled.

Port setting: “transmit enable = 0, receive enable = 0, learning disable =1”

Software action: the processor should not send any packets to the port. The switch may still send specific packets to the processor (packets that match some entries in the static table with “overriding bit” set) and the processor should discard those packets.

Blocking state: only packets to the processor are forwarded. Learning is disabled.

Port setting: “transmit enable = 0, receive enable = 0, learning disable =1”

Software action: the processor should not send any packets to the port(s) in this state. The processor should program the “Static Mac table” with the entries that it needs to receive (e.g. BPDU packets). The “overriding” bit should also be set so that the switch will forward those specific packets to the processor. Address learning is disabled on the port in this state.

Listening state: only packets to and from the processor are forwarded. Learning is disabled.

Port setting: “transmit enable = 0, receive enable = 0, learning disable =1”

Software action: The processor should program the “Static MAC table” with the entries that it needs to receive (e.g. BPDU packets). The “overriding” bit should be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state, see “special tagging” mode (section 3.2) for details. Address learning is disabled on the port in this state.

Learning state: only packets to and from the processor are forwarded. Learning is enabled

Port setting: “transmit enable = 0, receive enable = 0, learning disable = 0”

Software action: The processor should program the “Static MAC table” with the entries that it needs to receive (e.g. BPDU packets). The “overriding” bit should



be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state, see “special tagging” mode for details. Address learning is enabled on the port in this state.

Forwarding state: packets are forwarded and received normally. Learning is enabled.

Port setting: “transmit enable = 1, receive enable = 1, learning disable = 0”

Software action: The processor should program the “Static MAC table” with the entries that it needs to receive (e.g. BPDU packets). The “overriding” bit should be set so that the switch will forward those specific packets to the processor. The processor may send packets to the port(s) in this state, see “special tagging” mode for details. Address learning is enabled on the port in this state.

3.2 Special Tagging Mode

The special tagging mode is designed for Spanning Tree protocol IGMP snooping and is flexible for use in other applications. The special tagging mode, similar to 802.1Q, requires software to change network drivers to insert/modify/strip/interpret the special tag.

Table 1-Special Tagging Mode Format

802.1Q tag format	Special tag format
TPID (tag protocol identifier, 0x8100) + TCI.	STPID(special tag identifier, 0x810 + 4 bit for “port mask”) + TCI

The STPID will only be seen and used on the port 5 interface, which should be connected to a processor. Packets from the processor to the switch should be tagged with STPID and the port mask defined as below:

“0001”, packet to port 1 only.

“0010”, packet to port 2 only

“0100”, packet to port 3 only

“1000”, packet to port 4 only

“0011”, packet broadcast to port 1 and port 2.

.....

“1111” packet broadcast to port 1, 2, 3 and 4.

“0000” normal tag, will use KS8695PX Switch internal look up result. Normal packets should use this setting. If packets from the processors do not have a tag, the KS8695PX Switch will treat them as normal packets and an internal look up will be performed.

The KS8695PX Switch uses a non-zero “port mask” to bypass the look up result and override any port setting, regardless of port states (blocking, disable,



listening, learning). The table below shows the egress rules when dealing with STPID.

Table 2-STPID Egress Rules (Processor to Switch Port 5)

Ingress tag field	Tx port "tag insertion"	Tx port "tag removal"	Egress Action to tag field
(0x810+ port mask)	0	0	-Modify tag field to 0x8100 -recalculate CRC -no change to TCI if not null VID -replace VID with ingress (port 5) port VID if null VID
(0x810+ port mask)	0	1	-(STPID + TCI) will be removed. -padding to 64 bytes if necessary -recalculate CRC
(0x810+ port mask)	1	0	-Modify tag field to 0x8100 -recalculate CRC -no change to TCI if not null VID -replace VID with ingress (port 5) port VID if null VID
(0x810+ port mask)	1	1	-Modify tag field to 0x8100 -recalculate CRC -no change to TCI if not null VID -replace VID with ingress (port 5) port VID if null VID
Not Tagged	Don't care	Don't care	Determined by the dynamic MAC address table.

For packets from regular ports (port 1 – port 4) to port 5, the port mask is used to tell the processor which port the packet was received on, defined as

“0001” from port 1,
“0010” from port 2,
“0100” from port 3,
“1000” from port 4.

No values other than the previous four defined should be received in this direction in the special mode. The egress rule for this direction is defined as,

Table 3-STPID Egress Rules (Switch to Processor)

Ingress packets	Egress Action to tag field
Tagged with 0x8100 + TCI	-Modify TPID to 0x810 + “port mask”, which indicates source port. -no change to TCI, if VID is not Null -replace Null VID with ingress port VID. -recalculate CRC
Not tagged.	-Insert TPID to 0x810 + “port mask”, which indicates source port. -Insert TCI with ingress port VID -recalculate CRC

3.3 IGMP Support

There are two parts involved to support IGMP in layer 2. The first part is “IGMP” snooping. The switch will trap IGMP packets and forward them only to the processor port. The IGMP packets are identified as IP packets (either Ethernet IP packets or IEEE 802.3 SNAP IP packets) AND IP version = 0x4 AND protocol number = 0x2. The second part is “multicast address insertion” in the static MAC table. Once the multicast address is programmed in the static MAC table, the multicast session will be trimmed to the subscribed ports, instead of broadcasting to all ports.

3.4 Port Mirroring Support

KS8695PX Switch supports “port mirror” comprehensively as:

(1), “receive only” mirror on a port. All the packets received on the port will be mirrored on the sniffer port. For example, port 1 is programmed to be “rx sniff”, and port 5 is programmed to be the “sniffer port”. A packet, received on port 1, is destined to port 4 after the internal look up. The KS8695PX Switch will forward the packet to both port 4 and port 5. KS8695PX Switch can optionally forward even “bad” received packets to port 5.

(2), “transmit only” mirror on a port. All the packets transmitted on the port will be mirrored on the sniffer port. For example, port 1 is programmed to be “tx sniff”,

and port 5 is programmed to be the “sniffer port”. A packet, received on any of the ports, is destined to port 1 after the internal look up. The KS8695PX Switch will forward the packet to both port 1 and port 5.

(3), “receive and transmit” mirror on two ports. All the packets received on port A AND transmitted on port B will be mirrored on the sniffer port. To turn on the “AND” feature, set register 5 bit 0 to 1. For example, port 1 is programmed to be “rx sniff”, port 2 is programmed to be “transmit sniff” and port 5 is programmed to be the “sniffer port”. A packet, received on port 1, is destined to port 4 after the internal look up. The KS8695PX Switch will forward the packet to port 4 only, since it does not meet the “AND” condition. A packet, received on port 1, is destined to port 2 after the internal look up. The KS8695PX Switch will forward the packet to both port 2 and port 5.

Multiple ports can be selected to be “rx sniffed” or “tx sniffed”. And any port can be selected to be the “sniffer port”. All these per port features can be selected through “Advance Feature Control Register”.

3.5 VLAN support

KS8695PX Switch supports 16 active VLANs out of 4096 possible VLANs specified in IEEE 802.1Q. KS8695PX Switch provides a 16-entry VLAN table, which converts VID (12 bits) to FID (4bits) for address look up. If a non-tagged or null-VID-tagged packet is received, the ingress port VID is used for look up. In the VLAN mode, the look up process starts with VLAN table look up to determine whether the VID is valid. If the VID is not valid, the packet will be dropped and its address will not be learned. If the VID is valid, FID is retrieved for further look up. FID+DA is used to determine the destination port. FID+SA is used for learning purposes.

Table 4-FID+DA look up in the VLAN mode

DA found in Static MAC table	USE FID flag?	FID match?	DA+FID found in dynamic MAC table	Action
No	Don't care	Don't care	No	Broadcast to the membership ports defined in the VLAN table bit [20:16]
No	Don't care	Don't care	Yes	Send to the destination port defined in the dynamic MAC table bit[54:52]



Yes	0	Don't care	Don't care	Send to the destination port(s) defined in the static MAC table bit[52:48]
Yes	1	No	No	Broadcast to the membership ports defined in the VLAN table bit [20:16]
Yes	1	No	Yes	Send to the destination port defined in the dynamic MAC table bit[54:52]
Yes	1	Yes	Don't care	Send to the destination port(s) defined in the static MAC table bit[52:48]

Table 5-FID+SA look up in the VLAN mode

SA+FID found in dynamic MAC table	Action
No	The SA+FID will be learned into the dynamic table.
Yes	Time stamp will be updated.

Advanced VLAN features are also supported in KS8695PX Switch, such as “VLAN ingress filtering” and “discard non PVID” defined in “Port 1-4 Configuration Register 2 “. These features can be controlled on a port basis.

3.6 Rate Limit Support

KS8695PX Switch supports hardware rate limiting on “receive” and “transmit” independently on a per port basis. It also supports rate limiting in a priority or non-priority environment. The rate limit starts from 0 kbps and goes up to the line rate in steps of 32 kbps. The KS8695PX Switch uses one second as an interval. At the beginning of each interval, the counter is cleared to zero, and the rate limit mechanism starts to count the number of bytes during this interval.

For receive, if the number of bytes exceeds the programmed limit, the switch will stop receiving packets on the port until the “one second” interval expires. There is an option provided for flow control to prevent packet loss. If the rate limit is programmed greater than or equal to 128kbps and the byte counter is 8Kbytes



below the limit, the flow control will be triggered. If the rate limit is programmed lower than 128kbps and the byte counter is 2Kbytes below the limit, the flow control will be triggered.

For transmit, if the number of bytes exceeds the programmed limit, the switch will stop transmitting packets on the port until the "one second" interval expires.

If priority is enabled, the KS8695PX Switch can support different rate controls for both high priority and low priority packets.

Micrel Semiconductor, Inc.
1849 Fortune Drive



CENTAUR™ KS8695PX Switch Description

Integrated Multi-Port PCI Gateway Solution

San Jose, CA 95131
© Micrel, Inc. 2003
All rights reserved

Micrel is a registered trademarks of Micrel and its subsidiaries in the United States and certain other countries. All other trademarks are the property of their respective owners.