OreX

Ethernet Layer 2 Switch IP

Preliminary Data Sheet

Overview

OreX is a 3 port 10/100/1000 Ethernet layer 2 switch engine with a host interface port. The switch core provides wire speed forwarding at 1G x 4 ports.

32k bytes of buffering is provided on transmit for each port, split between multiple priorities. Store and forward architecture.

The host port can operate up to 1G bit/sec via a descriptor style DMA engine with an AHB master interface.

The core is compatible with Synopsys' Universal 10/100/1000 MAC.

The 'Switch and MAC wrapper' integrates the MAC's with the 'Layer 2 Switch Core'.

A configurable number of Ethernet MAC address are supported in a search engine that operates at wire rate for 4 ports.

Core characteristics

- *Core is parameterisable to support a programmable number of MAC addresses.*
- *Cells : Internal tables require memory cells.*
- *Architected to synthesise easily into ASIC or FPGAs – designed for minimum number of logic layers.*
- *Clock Frequency: 200 MHz. Single clock domain (outside MAC's).*
- *The core is architected to support up to 8 ports.*

Test suite

- *Comprehensive OVM compliant System Verilog Testbench.*

Specifications

L2 MAC address database

- The L2 MAC database stores a configurable number of MAC addresses. Independent VLAN learning is supported. Hardware assist is provided for adding and deleting MAC addresses to the database. An 'age' bit is provided per MAC address entry to support ageing.
- The database is searched by hardware via a hash table followed by a linked list of entries.

VLANs

- Support is provided for 32 VLANs. Tables are provided to map between 802.1q tags and an internal VLAN ID. On Tx a port can be configured (per VLAN) to append or strip 802.1q VLAN tags.
- A forwarding port bit mask per VLAN is provided for unknown MAC addresses.

Spanning tree

- Per port/ per VLAN spanning tree state is maintained. BPDU's are forwarded to the host.

Priority

- 2 priority queues are supported on transmit. The priority of frames can be determined via multiple sources. 3 priority queues are supported on the host interface.

Flow Control

- Xon/Xoff flow control is supported.

Snooping

- Support is provided to snoop on the Ethertype or IPv4 protocol field.

Phy Interfaces

- RGMII and MDIO per port.

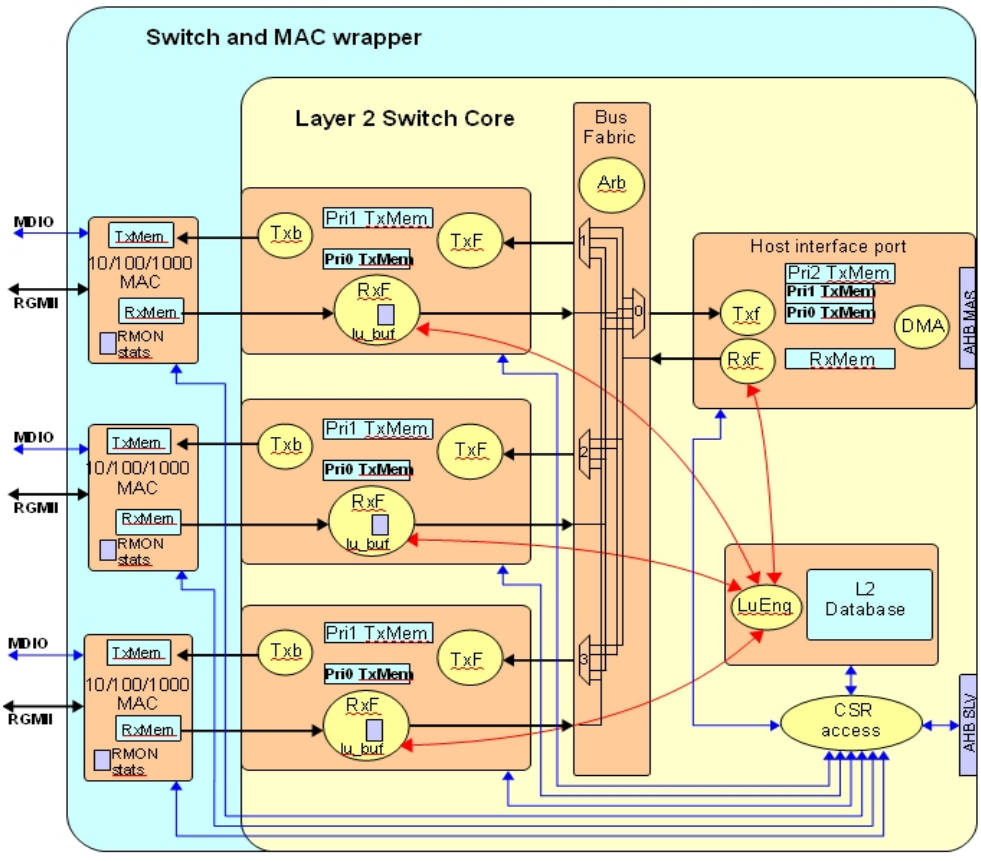

Specifications

L2MAC address database

An entry per MAC address is maintained in the L2 database. Each entry stores a MAC address, a destination port mask, VLAN ID and other associated data. Independent VLAN learning is supported.

The L2 MAC database stores a configurable number of MAC addresses. 32k bytes of RAM are required to support 1k MAC addresses.

Hardware assist is provided for adding and deleting MAC addresses to the database. Software indicates the MAC address (and VLAN ID) to be added/deleted, the hardware generates the hash and adds/deletes the MAC address to/from the database.

An 'age' bit is provided per MAC address entry to support ageing. A ‘permanent’ bit is provided per MAC address to identify MAC addresses that should not be aged out. Up to 4 unknown source MAC address will be stored in a FIFO for learning by software.

Software also has direct access to the lookup database.

The database is searched by hardware via a hash table followed by a linked list of entries. The depth of a linked list can be limited to ensure lookups are always wire rate. The hash table size is configurable to reduce hash collisions and is typically 4 to 8 times the number of MAC addresses.

For frames from the host, a method is provided to bypass the layer 2 lookup and forward the frame as indicated by the host. For example this is required for forwarding BPDU’s from the host.

VLANs

Support is provided for 32 VLANs. There is a default VLAN ID per port for frames that are not 802.1q tagged or frames that are tagged with a null VLAN. 802.1q tagged frames are mapped to an internal 5 bit VLAN ID via a programmable table. On the transmit side a port can be configured (per VLAN) to append or strip 802.1q VLAN tags. A VLAN id to VLAN tag table is provided to support this.

A forwarding port bit mask per VLAN is provided for unknown MAC addresses (unicast, multicast or broadcast).

Spanning Tree / Fast spanning Tree

To support spanning tree a port can be placed in one of 3 states (discarding, learning or forwarding). The state of each port can be defined per VLAN.

BPDU's received on the Ethernet ports are always forwarded to the host. BPDU’s from the host are forwarded as requested by the host.

Multicasts/Broadcasts

Multicasts and Broadcast MAC addresses may also be entered in the L2 database (per VLAN). A destination bit mask per VLAN is provided for unknown multicasts

Broadcast storm

If a broadcast storm is detected by the host software (via RMON stats counters) then the broadcast MAC address entry in the database can be temporarily modified to discard broadcasts.

An optional extra is to maintain a ‘broadcast rate count’ per Ethernet port on receive. When the rate exceeds a programmable value broadcasts frames are dropped. Typically broadcast shouldn’t exceed 10M bit/sec.

Priority

2 priority queues are supported on transmit. The priority of frames can be determined from 4 possible sources listed in order of precedence (1 being highest) these are:

1. A priority programmed on a source port.
2. The 'DSCP' field in the IPV4 TCP header or IPV6 COS.
3. A priority associated with a MAC address.
4. The 802.1p header

Buffering

32k bytes of fixed buffering is provided on transmit. This is split into 24k bytes for low priority traffic (pri1) and 8k bytes for high priority low latency traffic (pri0).

Flow Control

The switch can be configured to either ‘discard frames on transmit’ if the transmit FIFO is full or ‘store frames at receive’ until the transmit FIFO has space.

Xon/Xoff flow control is supported via the MAC. If the Rx FIFO in the MAC reaches a programmable threshold then flow control is triggered. When flow control is enabled the switch should be setup to ‘store frames at receive’.

L2/L3 Snooping

Support is provided to snoop on the Ethertype or IPv4 protocol field. 4 different Ethertypes and 4 different protocol fields can be configured for forwarding frames ‘only to the host’ or ‘copying to the host’.

For example it is possible to snoop on IGMP frames, which are forwarded as normal and also copied to the host interface.

HOST DMA

A descriptor based DMA engine is provided to transmit and receive frames from the host. 3 priority queues are provided on the interface to the host interface. A single queue is supported on the interface from the host. The DMA engine is an AHB master.

CSR access

An AHB slave interface is provided to access control registers, the lookup database and MAC control registers and statistics.

MAC functions

IEEE1588 is supported via the MAC.

RMON1 statistics are provided via the MAC.

The receive side of the MAC needs at least 6k bytes. The transmit side needs 2k bytes.

An RGMII and MDIO interface is provided by each MAC.

Miscellaneous
The max frame size supported is 2048bytes.
A software reset is provided for each Ethernet port.
Flush per queue is provided on the Ethernet ports and the host port.