

Intel[®] IXP2400 Network Processor / Vitesse GigaStream[®] Switch Fabric Solution White Paper

Revision 1.0

November, 2002

Legal Notices and Disclaimers

Disclaimers

Information in this document is provided in connection with Intel® and Vitesse products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's or Vitesse's Terms and Conditions of Sale for such products, Intel/Vitesse assumes no liability whatsoever, and Intel/Vitesse disclaims any express or implied warranty, relating to sale and/or use of Intel®/Vitesse products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel®/Vitesse products are not intended for use in medical, life saving, or life sustaining applications. Intel®/Vitesse may make changes to specifications and product descriptions at any time, without notice.

Legal Notices

Intel is a registered trademark of the Intel Corporation or its subsidiaries in the United States and other countries.

Vitesse is a registered trademark of the Vitesse Semiconductor Corporation or its subsidiaries in the United States and other countries.

Revision History

Version	Date	Description
0.0	8/23/02	Initial draft
0.1	10/8/02	Edited draft (cml)
1.0	11/20/02	Release

Table Of Contents

1	Introduction	9
1.1	Scope.....	9
1.2	Reference Documents.....	9
1.3	Acronyms	9
2	Intel IXP2400 Network Processor Overview	10
2.1	Product Description	10
2.2	Media and Switch Fabric Interface	10
2.2.1	Overview.....	10
2.2.2	CSIX-L1 Mode.....	11
3	Vitesse GigaStream Overview.....	12
3.1	Functional Description	12
3.2	Block diagram	12
3.2.1	GigaStream Queuing Engine.....	13
3.2.2	GigaStream Crossbar Device	13
3.2.3	Active Redundancy.....	13
3.2.4	CSIX-L1 Interface	13
4	System Block Diagram	14
4.1	Single 2.5 Gb/s Line Card	14
4.2	Dual 2.5 Gb/s Line Card	15
5	CSIX-L1 Electrical Interface	16
5.1	Intel IXP2400 / Vitesse VSC872 CSIX-L1 Signals Mapping.....	16
5.2	Clocking Scheme	17
5.3	PCB Layout Guideline	18
5.4	CBus Interface.....	18
6	CSIX-L1 Interface Functional Description.....	20
6.1	CSIX-L1 Common Spec.....	20
6.1.1	Summary of Supported CFrame Types.....	20
6.2	Base Header.....	21
6.2.1	Ready Field	21
6.2.2	Type Field.....	22

6.2.3	CSIX Reserved Bit (CR).....	22
6.2.4	Private Bit (P).....	22
6.2.5	Payload Length Field.....	22
6.3	Idle CFrames.....	22
6.4	Data CFrames	22
6.4.1	Ingress Unicast Extension Header	23
6.4.2	Egress Unicast Extension Header	23
6.4.3	Ingress and Egress Multicast Extension Header	23
6.4.4	Ingress and Egress Broadcast Extension Header	24
6.5	Flow Control CFrames.....	24
6.5.1	Fabric to NPE Flow Control Header.....	24
6.6	Parity	25
6.6.1	Horizontal Parity	26
6.6.2	Vertical Parity	26
6.7	Dead Cycle.....	27
6.8	Operation and Timing.....	27
6.8.1	Start-up	27
6.8.2	Transmission	27
6.8.3	Pause and Resume Operation	27
6.8.4	Parity Errors.....	27
6.8.5	Unexpected Start of Frame (SOF).....	27
6.9	Link Level Flow Control.....	28
6.10	VOQ Flow Control Frame Latency	28
6.10.1	Block Diagram Description.....	28
6.10.2	Flow Control Modeling.....	29
7	CPU Interface	30
7.1	Intel IXP2400 CPU Interface	30
7.2	Vitesse GigaStream VSC872 CPU Interface.....	30
7.3	Intel IXP2400 / Vitesse VSC872 CPU Interface Concept Drawing.....	30
8	Intel IXP2400 / Vitesse GigaStream Reference Design	32
8.1	Functional Description.....	32
8.2	Modes of Operation.....	33
8.3	Hardware Block Description.....	33
8.3.1	Mictor Connectors.....	33

8.3.2	CPLD	33
8.3.3	CSIX Clock Generation	34
8.3.4	VSC872 Transceiver.....	34
8.3.5	Power Generation	34
8.4	Thermal Considerations.....	35

Tables

<i>Table 5.1 Intel IXP2400 / Vitesse VSC872 CSIX Signal Mapping.....</i>	<i>16</i>
<i>Table 5.2 Intel IXP2400 / Vitesse VSC872 CSIX Interface PCB Layout Guideline.....</i>	<i>18</i>
<i>Table 6.1 CFrame Structure</i>	<i>20</i>
<i>Table 6.2 Overhead by Frame Type</i>	<i>20</i>
<i>Table 6.3 Base Header Layout.....</i>	<i>21</i>
<i>Table 6.4 Ready Bit Assignment by CFrame Type</i>	<i>21</i>
<i>Table 6.5 Type Field Values</i>	<i>22</i>
<i>Table 6.6 Ingress Unicast Extension Header.....</i>	<i>23</i>
<i>Table 6.7 Egress Unicast Extension Header.....</i>	<i>23</i>
<i>Table 6.8 Ingress and Egress Multicast Extension Header.....</i>	<i>23</i>
<i>Table 6.9 Ingress and Egress Broadcast Extension Header</i>	<i>24</i>
<i>Table 6.10 Fabric to NPE Flow-control Header Format (Payload Bytes).....</i>	<i>25</i>
<i>Table 6.11 Flow Control Entry Type Values.....</i>	<i>25</i>
<i>Table 6.12 Vertical Parity.....</i>	<i>26</i>

Figures

<i>Figure 2.1 Intel IXP2400 External Interface Block Diagram</i>	10
<i>Figure 3.1 Vitesse GigaStream Switch Fabric Chipset Block Diagram</i>	12
<i>Figure 4.1 Single 2.5 Gb/s Full-Duplex Line Card Block Diagram</i>	14
<i>Figure 4.2 Dual 2.5 Gb/s Full-Duplex Line Card Block Diagram</i>	15
<i>Figure 5.1 Intel IXP2400 / Vitesse VSC872 CSIX Interface</i>	16
<i>Figure 5.2 Intel IXP2400 / Vitesse VSC872 CSIX Interface Clocking Scheme</i>	17
<i>Figure 5.3 Intel IXP2400 / Vitesse VSC872 CSIX PCB Interface Topology</i>	18
<i>Figure 5.4 Full-Duplex CBus Connection Block Diagram</i>	19
<i>Figure 6.1 Intel IXP2400 / Vitesse VSC872 Flow Control Loop</i>	28
<i>Figure 7.1 Intel IXP2400 / Vitesse VSC872 CPU Interface Conceptual View</i>	31
<i>Figure 8.1 Anacapa Block Diagram</i>	32

1 Introduction

1.1 Scope

This white paper describes the interconnection between the Intel® IXP2400 Network Processor and the Vitesse GigaStream® Switch Fabric, and their interoperability through the CSIX-L1 interface to provide a complete network processor and switch fabric solution for high-performance communication systems.

The paper covers the CSIX interface electrical and timing compatibilities, clock distribution, CPU register configurations, PCB layout guidelines, and a discussion of switch fabric interface flow control.

For more information about the Intel IXP2400 Network Processor and the Vitesse GigaStream chipset, including specifications and applications other than the switch fabric interface, refer to the reference documents listed below.

1.2 Reference Documents

CSIX-L1:Common Switch Interface Specification-L1	Rev 1.0	August 5, 2000
Intel IXP2400 Network Processor Advance Information Datasheet	Rev 002	March, 2002
Intel IXP2400 Network Processor Hardware Reference Manual	Rev 001	April, 2002
Intel IXP2400 Programmer's Reference Manual	Rev 0.4	June, 2002
Vitesse GigaStream Design Manual	Rev 2.2	April 2002
Vitesse GigaStream Datasheet	Rev 2.1	June 2002

1.3 Acronyms

CPLD	Complex Programmable Logic Device
FIFO	First In First Out Buffer
CSIX-L1	Common Switch Interface Specification - Level 1
Gb/s	Gigabit per second
HSSL	High Speed Serial Link
MSF	Media and Switch Fabric
PCB	Printed Circuit Board
PHY	Physical Layer
POS	Packet Over SONET
SerDes	Serializer/De-Serializer
SPI-3	System Packet Interface Level 3: OC-48 System Interface for Physical and Link Layer Devices.
UTOPIA	Universal Test and Operations PHY Interface for ATM
VOQ	Virtual Output Queue

2 Intel IXP2400 Network Processor Overview

2.1 Product Description

The IXP2400 Network Processor is the second-generation network processor from Intel Corporation. It enables fast deployment of intelligent network services by providing the maximum programming flexibility, code re-use, and high-performance parallel packet-processing power. It supports a wide variety of WAN and LAN applications with a broad range of speeds up to the OC-48 / 2.5Gb/s line rate.

The IXP2400 Network Processor has a store-and-forward architecture that combines a state of the art Xscale microprocessor core with eight multithreaded, independent, 32-bit programmable Microengines which include specific support for network packet processing.

The IXP2400 Network Processor's major external interfaces are shown as a block diagram in Figure 2.1.




Figure 2.1 Intel IXP2400 External Interface Block Diagram

2.2 Media and Switch Fabric Interface

2.2.1 Overview

The IXP2400 Network Processor Media and Switch Fabric (MSF) interface is used to connect the IXP2400 to a physical layer device (PHY) and/or to a Switch Fabric. The MSF consists of separate receive and transmit interfaces, which are unidirectional and independent of each other. See the diagram in Figure 2.1 for details. Each of the receive and transmit interfaces can be separately configured as either UTOPIA Level 1/2/3 or POS-PHY 2/3 for a PHY device interface, or CSIX-L1 protocol for a Switch Fabric

interface. This document focuses on the IXP2400 MSF in CSIX-L1 mode. For information on the IXP2400 MSF in UTOPIA or POS-PHY mode, see the “Intel IXP2400 Hardware Reference Manual”.

2.2.2 CSIX-L1 Mode

The IXP2400 MSF in CSIX-L1 mode is a 32-bit wide interface using 3.3V LVTTTL signaling and globally synchronous (common) clocking. Thus, it does not have electrical and clocking compatibility with the CSIX-L1 specification, which is 2.5V LVCMOS signaling and source synchronous clocking. The IXP2400 MSF CSIX-L1 interface operates at up to 125MHz, to transmit and receive packets at a peak rate of 4Gb/s.

The IXP2400 MSF in CSIX-L1 mode supports both Link Level and Virtual Output Queue (VOQ) flow-control as defined in the CSIX-L1 specification (see Section 6.9 and Section 6.10 for more details).

Per the CSIX-L1 specification, the network processor that handles traffic sent to the Switch Fabric is defined as the Ingress network processor; the network processor that handles traffic received from Switch Fabric is defined as the Egress network processor. The IXP2400 is implemented with a CBus, which is used for inter-processor communication of Link Level flow-control and VOQ flow-control information between the Egress processor and the Ingress processor. See Section 5.4 for more details about the IXP2400 CBus usage model.

For transmission from the Ingress IXP2400, CFrames are constructed under microengine software control and then written into the transmit buffer. Vertical and horizontal parity generation is performed by hardware. Also, the hardware will automatically handle the transmission of Idle CFrames with link level RDY bits constantly updated when there are no data or flow-control CFrames to transmit.

When the Egress IXP2400 is receiving data, Idle CFrames are recognized by hardware and discarded. VOQ flow-control CFrames and Link Level flow-control information are both handled by hardware (using the CBus). All other types of CFrames are buffered and passed to a Microengine to be parsed by software. However, Link Level flow-control information in the Base Header of all CFrames (including Idle CFrames) is handled by hardware.

3 Vitesse GigaStream Overview

3.1 Functional Description

The Vitesse GigaStream switch fabric chipset includes a queuing engine (VSC872) for line cards and a crossbar device (VSC882) for switch cards. The GigaStream chipset provides up to 32 CSIX interfaces, which can be seamlessly connected to the Intel IXP2400 network processor. This section gives an overview of the GigaStream architecture. More details can be found in the GigaStream Design Manual. GigaStream features include:

- Highly-Integrated, Two-Chip Set Solution
- Aggregate User Bandwidth of up to 80Gb/s
- Aggregate Backplane Bandwidth of up to 320Gb/s
- Maximum Port Configurations of up to 32 OC-48 or 64 Gigabit Ethernet
- Flexible N+1 and N+N *Active Redundancy* Schemes
- Integrated Queuing, Central Scheduling, and Crossbar Switching
- Sophisticated QoS handling with Eight Priorities
- Advanced Unicast and Multi/Broadcast Support
- Field-Proven, High-Speed 2.644Gb/s Serial Link Technology

3.2 Block diagram

The following figure is a block diagram of the GigaStream chipset in its maximum configuration of 32 CSIX layer 1 interfaces.




Figure 3.1 Vitesse GigaStream Switch Fabric Chipset Block Diagram

3.2.1 GigaStream Queuing Engine

The VSC872 queuing engine contains two 32-bit CSIX L1 interfaces and eight high-speed serial links (HSSLs). In a typical application, an ingress and an egress IXP2400 network processor can be connected to each of the two CSIX L1 interfaces. On the backplane side, up to eight VSC882 switch chips can be connected to a VSC872 engine using up to eight HSSLs. There is an ingress queuing structure for each of the two CSIX interfaces. Each queuing structure contains up to 16 virtual output queues (VOQs). Each VOQ plane contains up to eight priority levels. There is also a separate VOQ structure for multicast traffic.

3.2.2 GigaStream Crossbar Device

The VSC882 crossbar switch integrates 16 HSSLs, which operate at up to 2.5Gb/s. Up to eight VSC882 switches can be used in a GigaStream system. Each VSC882 switch contains a built in scheduler which accepts connection requests from, and issues grants to, the VSC872 engines on the line cards. Each scheduler operates independently. These devices support two levels of strict priority. Both high- and low-priority connection requests arrive in-band through the data channels. An iterative matching algorithm is used first on the high priority requests and next on the low-priority requests before sending the grant information in-band back to the line cards.

3.2.3 Active Redundancy

The VSC872 provides a load balancing function that automatically distributes the load across up to eight VSC882 crossbar switches. The VSC872 engine determines the number of active VSC882 switches in the system, and only sends frames to working switches. This provides two benefits. First, a variety of redundant configurations can be supported, including 1+1 and N+1. Also, the system bandwidth can be upgraded by simply adding more VSC882 chips.

3.2.4 CSIX-L1 Interface

The VSC872 queuing engine contains two 32-bit wide CSIX L1 interfaces, using 3.3V LVTTTL signaling and source synchronous clocking. This interface operates at up to 166 MHz to transmit and receive packets at a peak rate of 5.3Gb/s. The VSC872 supports VOQs (Virtual Output Queues) and up to eight classes on the ingress side, and supports both Link Level and Virtual Output Queue (VOQ) flow control, as defined in the CSIX-L1 specification (see Section 6.9 and Section 6.10 for more details). On the egress side, the VSC872 engine supports two strict-priority levels, and only Link Level flow-control from the egress IXP2400 is acted upon.

For CSIX transmission from the VSC872 engine, CFrames are constructed and written into a transmit FIFO. Vertical and horizontal parity generation is performed by hardware. Also, the hardware will automatically handle the transmission of Idle Cframes, with link level RDY bits constantly updated, when there are no data CFrames to transmit.

On the VSC872 engine receive side, Idle CFrames are recognized by hardware and discarded. Frames are categorized as unicast or multicast and placed into the appropriate ingress VOQ and class queue. The frames are checked for both horizontal and vertical parity. CFrames with parity errors are discarded.

4 System Block Diagram

4.1 Single 2.5 Gb/s Line Card

The following figure shows two IXP2400 Network Processors in a typical single 2.5Gb/s port, full-duplex line card with a Vitesse VSC872 Queuing Engine.




Figure 4.1 Single 2.5 Gb/s Full-Duplex Line Card Block Diagram

As shown in the diagram, the Ingress IXP2400 MSF receive port is configured as SPI-3 for the Physical device interface, and the transmit port is configured as CSIX-L1 for the Switch Fabric interface to the Vitesse VSC872 Queuing Engine.

On the Egress side, the IXP2400 MSF receive port is configured as CSIX-L1 for the Switch Fabric interface to the VSC872 Queuing Engine, and the transmit port is configured as SPI-3 for the Physical device interface.

The IXP2400 CBus is configured as full-duplex mode, which is used for communicating Switch Fabric flow-control information from the Egress IXP2400 to the Ingress IXP2400. VSC872 Link Level flow control, as well as VOQ Flow Control Cframes, are received by the Egress IXP2400 from the VSC872 via the CSIX interface. The egress IXP2400 transmits Link Level flow-control information and Flow Control CFrames to the Ingress IXP2400 through the CBus. This allows the ingress IXP2400 to process flow-control information and respond accordingly.

The Vitesse VSC872 Queuing Engine supports two full-duplex CSIX-L1 ports. Only port 0 is used in this block diagram, so port 1 is unused and terminated.

The Vitesse VSC872 engine supports up to eight High Speed Serial Links (HSSL). For single 2.5 Gb/s full-duplex application, only four HSSL links are required. The SerDes of the unused four HSSL links are powered down through register settings.

Any one of the two IXP2400 Xscale cores can perform configuration access to the VSC872 internal registers. In this 2.5 Gb/s line card example, the Egress IXP2400 Xscale core is responsible for configuring the VSC872's internal configuration-and-status registers. A CPLD can be used as a bridge between the Egress IXP2400 SlowPort and the VSC872 CPU Interface for bus protocol conversion. See Section 7 for more details.

4.2 Dual 2.5 Gb/s Line Card

To utilize both CSIX-L1 channels on the Vitesse VSC872 Queuing Engine, a dual full-duplex 2.5 Gb/s line card can be implemented as shown in following figure.




Figure 4.2 Dual 2.5 Gb/s Full-Duplex Line Card Block Diagram

The two CSIX-L1 ports on the Vitesse VSC872 Queuing Engine operate independently of each other; each CSIX interface having its own ingress and egress queuing structure.

Any one of the four IXP2400 Xscale cores can access and configure the VSC872 internal registers. In the above, dual 2.5 Gb/s full-duplex line card, example the Egress IXP2400 Xscale core on CSIX channel 1 of the VSC872 is responsible for VSC872 configuration.

All eight of the 2.5 Gb/s HSSL links on the VSC872 engine are used to achieve dual 2.5 Gb/s full-duplex line card operation.

5 CSIX-L1 Electrical Interface

5.1 Intel IXP2400 / Vitesse VSC872 CSIX-L1 Signals Mapping

Both the Intel IXP2400 CSIX-L1 interface and the Vitesse GigaStream VSC872 CSIX-L1 interface are 32-bits wide and employ LVTTTL signaling. For a single 2.5 Gb/s full-duplex line card solution, the two IXP2400 Network Processors (Ingress processor and Egress processor) can be connected either to CSIX channel 0 or CSIX channel 1 of the VSC872, as shown in Figure 5.1.

However, the Intel IXP2400 CSIX-L1 interface uses a clocking scheme that is different from the VSC872 engine. The clocking distribution will be discussed in section 5.2.




Figure 5.1 Intel IXP2400 / Vitesse VSC872 CSIX Interface

The IXP2400 to VSC872 CSIX interface signal mapping is listed in following table.

IXP2400 Signal Name	I/O	VSC872 Signal Name	I/O	Description
RXCLK[0]	I			IXP2400 Receive Clock See Section 5.2 for more details
RXSOF[0]	I	C0_TxSOF Or C1_TxSOF	O	Transmit Start of Frame – asserted by the switch fabric for one clock to indicate the start of the CFrame.
RXDATA[31:0]	I	C0_TxData[31:0] Or C1_TxData[31:0]	O	Transmit Data – 32-bit output from the fabric to network processor.
RXPRTY[0]	I	C0_TxPar Or C1_TxPar	O	Transmit Bus Parity
		C0_TxCk Or C1_TxCk	O	VSC872 Transmit Clock – The switch fabric provides a data synchronization clock. See Section 5.2 for more details.
TXCLK[0]	I			IXP2400 Transmit Clock. See Section 5.2 for more details.
TXSOF[0]	O	C0_RxSOF Or C1_RxSOF	I	Receive Start of Frame – asserted by the network processor for one clock to indicate the start of the CFrame.
TXDATA[31:0]	O	C0_RxData[31:0] Or C1_RxData[31:0]	I	Receive Data – 32-bit input to fabric.
TXPRTY[0]	O	C0_RxPar Or C1_RxPar	I	Receive Bus Parity
		C0_RxCk Or C1_RxCk	I	VSC872 Receive Clock. See Section 5.2 for more details.

Table 5.1 Intel IXP2400 / Vitesse VSC872 CSIX Signal Mapping

5.2 Clocking Scheme

The IXP2400 CSIX-L1 interface uses a global-synchronous clocking scheme, which is different from the CSIX-L1 specification issued by the Network Processor Forum. The VSC872 engine uses the source-synchronous clocking scheme that is defined in the CSIX-L1 specification. To implement a glueless CSIX interface between the IXP2400 and the VSC872, a global synchronous clocking scheme with delay lines can be used to compensate for the clock skews, as shown in the following figure. In this example, CSIX port 0 is used on the VSC872 chip. CSIX port 1 uses the same connection scheme, but shares the TxRefClk input.

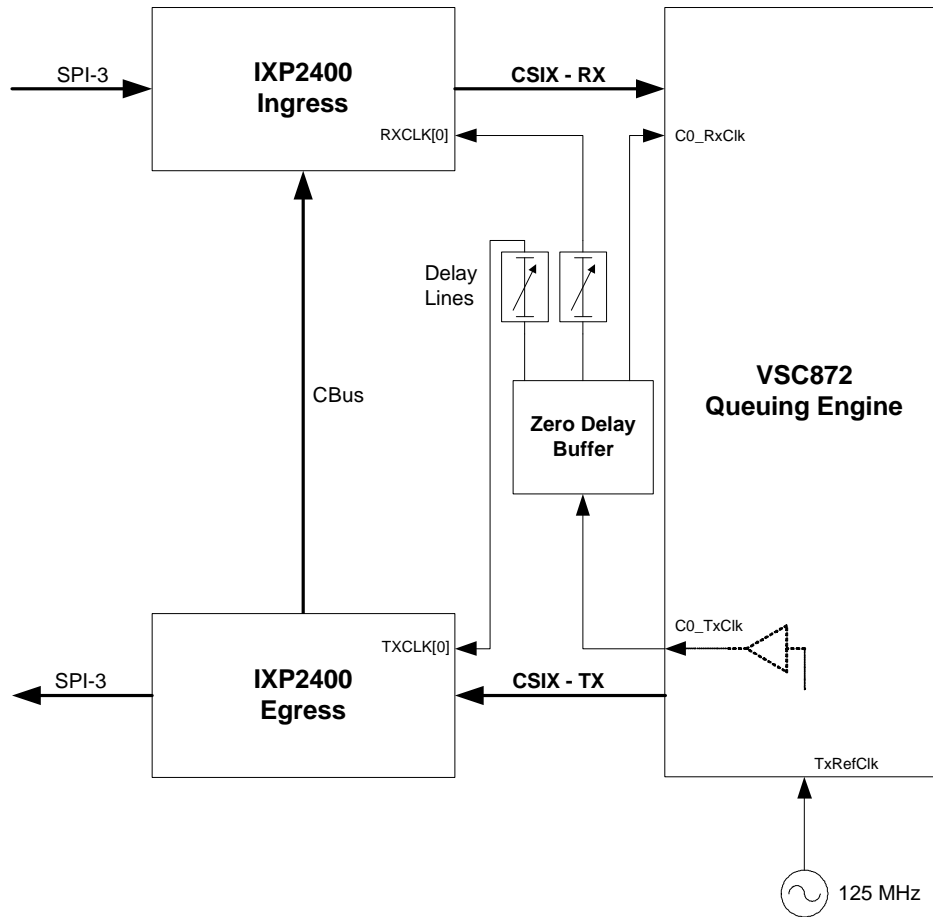


Figure 5.2 Intel IXP2400 / Vitesse VSC872 CSIX Interface Clocking Scheme

A 125 MHz oscillator is used to provide a clock source to the VSC872 TxRefClk input. A clock buffer inside the VSC872 provides a 125MHz C0_TxClk output from this TxRefClk input. This C0_TxClk output is used as an input to a very-low-skew Zero Delay buffer. One output of the Zero Delay buffer is used to provide a receive clock to the C0RxClk input of the VSC872, which latches the CSIX receive data.

The other two copies of the Zero Delay buffer output will be used through delay lines as the clock sources to the IXP2400 CSIX transmit clock and receive clock. By adjusting delay lines to change the clock skews, the IXP2400 CSIX data will match the VSC872 engine's source-synchronous clocks for both transmit and receive CSIX interfaces.

The LDH54 series chip delay lines from muRata can be used; these can provide delay in 100ps steps.

5.3 PCB Layout Guideline

The IXP2400 and the VSC872 CSIX 32-bit interface can operate together at 125 MHz. Certain PCB layout rules must be followed to ensure good signal integrity. The CSIX PCB interface topology is shown in following figure.




Figure 5.3 Intel IXP2400 / Vitesse VSC872 CSIX PCB Interface Topology

It is highly recommended that system designers perform signal integrity simulations for the PCB layout. The IXP2400 and the VSC872 CSIX interface IBIS models are available from Intel and Vitesse. The following table provides a PCB layout guideline for reference.

Parameter	Routing Guidelines
Signal Group	Data /Control / Parity
Topology	Point to point
Reference Plane	Ground
Characteristic Trace Impedance	50-60 $\Omega \pm 10\%$
Nominal Trace Width	5 mils
Nominal Trace Separation	9 mils minimum, 50 mils recommended
Group Spacing	Isolation from non-media-related signals = 50 mils
IXP2400 Breakout Guideline	4 mil with 4 mil space for a max of 300 mils
Trace Length P^1+A	Max = 6.0" The trace length from ball to ball should be less than 20 mils.
Max Via Count per Signal	4 vias

Note: P^1 refers to the package trace length.

Table 5.2 Intel IXP2400 / Vitesse VSC872 CSIX Interface PCB Layout Guideline

5.4 CBus Interface

The IXP2400 Network Processor CBus interface has two modes of operation:

- Full-Duplex mode - Used to communicate Link Level flow-control and VOQ flow-control information from the Egress IXP2400 to the Ingress IXP2400 directly.
- Simplex Mode - Used to communicate Link Level flow-control and VOQ flow-control information directly to the switch fabric.

When using the IXP2400 Network Processor and the GigaStream VSC872 Switch Fabric in a glueless interface solution, the CBus must be configured in full-duplex mode.

When the IXP2400 Network Processor MSF is used in CSIX mode, the CBus can be configured as either four-bit mode or eight-bit mode by register settings. For eight-bit mode, the eight unused signals will be used as TXCDATA[7:4] and RXCDATA[7:4], see the IXP2400 Advanced Information Datasheet for eight-bit mode signal mappings.

Figure 5.4 shows an example of the CBus in four-bit full-duplex mode including the connections between the Ingress IXP2400 and the Egress IXP2400.




Figure 5.4 Full-Duplex CBus Connection Block Diagram

The Link level flow-control information is transmitted serially on the TXCSR output pin on the Egress IXP2400 and received serially on the RXCSR input pin on the Ingress IXP2400.

The VOQ flow-control CFrame is transmitted from the Egress IXP2400 using the TXCSOF, TXCDATA[3:0], and TXCPAR signals to the Ingress IXP2400 through the RXCSOF, RXCDATA[3:0], and RXCPAR signals.

The CBus operates using the CSIX interface clocks with the global-synchronous clocking scheme.

For more information about the IXP2400 CBus signals definition and operation, see the Intel IXP2400 Network Processor Hardware Reference Manual.

6 CSIX-L1 Interface Functional Description

This section contains specification details of the common CSIX Layer 1 interface between the Intel IXP2400 Network Processor and the Vitesse GigaStream Switch Fabric. Although each product implements a super-set of this specification, the following must be used to insure inter-operability between the two devices.

6.1 CSIX-L1 Common Spec

A CFrame consists of a base header, an optional (determined by type) extension header, an optional payload, and a 16-bit vertical parity field. CFrame headers are based on a layered approach to minimize total overhead for varying types of frames. Every CFrame begins with the base header, which is two bytes long and contains the payload-length, frame-type, and ready bits (for Link level flow control). Type-specific extension headers, if needed, follow the base header. The base header frame type determines the extension header format. Extension headers contain additional information needed to handle the frame, such as the destination fabric egress CSIX interface for unicast frames. The header formats and field definitions are described in subsequent sections.

A two-byte vertical parity field follows the payload. The parity bytes are always the highest-numbered bytes of the last word (recall that the most-significant bytes will be the least-significant bits of the last word). If the payload would result in 0 or 1 bytes unused in the last CWord, a new CWord must be added and the parity bytes will appear in the least-significant bits of the added word. In addition, padding, added to the CFrame to match CWord boundaries, is included in the vertical parity calculations. See the Padding definition in CSIX specification.

Although the CSIX specification allows a maximum payload length of 256 bytes, the GigaStream chipset supports a maximum payload length of 40 to 120 bytes, and the IXP2400 supports a maximum payload length of 64 to 256 bytes. Therefore, the maximum payload value when using both of these devices can range from 64 to 120 bytes. The CFrame structure is described in Table 6.1.

CFrame Component	Base Header	Extension Header	Payload	Vertical Parity
Length	2 bytes	0-7 bytes	Maximum payload value: 64-120 bytes	2 bytes

Table 6.1 CFrame Structure

6.1.1 Summary of Supported CFrame Types

Table 6.2 summarizes the number of overhead bytes for each supported CFrame type.

Type Encoding	Frame Type	Total CSIX Overhead (Base Header + Extension Header + Parity)
0x0	Idle	2+0+2=4 bytes
0x1	Unicast	2+4+2=8 bytes
0x3	Multicast ID	2+4+2=8 bytes
0x5	Broadcast	2+4+2=8 bytes
0x6	Flow Control Frame*	2+4+2=8 bytes
All others	N/A	N/A

* Flow control does not use an extension header, instead it uses the first 2 bytes of the payload

Table 6.2 Overhead by Frame Type

6.2 Base Header

Table 6.3 shows the layout of the base header.

Byte Number	Bit Position															
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
0/1	Ready		Type				CR	P	Payload Length							

Table 6.3 Base Header Layout

6.2.1 Ready Field

The ready field is used to indicate when the transmitting entity is ready to receive data. This is sometimes referred to as Link level flow control. A low (0) ready bit indicates that the entity is not ready to receive the specified traffic type; a high (1) ready bit indicates that the entity is ready to receive the specified traffic type. There are 2 ready bits representing 2 Link level queues. The two groups and the CFrame types assigned to these groups are listed below. Note that the command and status frame ready bits are unsupported.

Ready[0] (bit 6 of byte 0) is for Flow Control Traffic

- Flow-control Frames

Ready[1] (bit 7 of byte 0) is for Data Traffic

- Unicast
- Multicast
- Broadcast

When no data is being transmitted, the ready field is kept alive by regular transmission of Idle CFrames. Idle frames are neither control or data frames, and can be sent at any time.

When an error is detected (such as a parity error) on the CFrame, ready bits from the CFrame are ignored and interpreted as “not” ready. In the GigaStream switch fabric, the CSIX ingress interface detecting such an error causes the corresponding CSIX egress interface to suspend transmission of CFrames until the ingress interface receives a CFrame without error. The IXP2400 can be programmed to perform various operations based on the reception of an errored CFrame. Table 6.4 shows which ready bit controls the transmission of each CFrame type. The Ready field is valid in both the ingress and the egress directions.

Encoding	Type	Ready Bit CRdy=Ready[0] Drdy=Ready[1]
0x0	Idle	None
0x1	Unicast	Drdy
0x3	Multicast ID	Drdy
0x5	Broadcast	Drdy
0x6	Flow Control	Crdy
Others	N/A	N/A

Table 6.4 Ready Bit Assignment by CFrame Type

6.2.2 Type Field

The type field encodes the type of CFrame being transferred. The payloads of all the frame types, except flow-control frames, are unspecified by CSIX. The flow-control CFrame payload is fully defined in Section 6.5. For the unicast, broadcast and multicast frames, there is an extension header defined to follow the base header. The formats for these extension headers are described in the following sections. The Type field is valid in both the ingress and egress directions.

Encoding	Type
0x0	Idle
0x1	Unicast
0x3	Multicast ID
0x5	Broadcast
0x6	Flow Control Frame
Others	N/A

Table 6.5 Type Field Values

6.2.3 CSIX Reserved Bit (CR)

The CSIX reserved bit is not used.

6.2.4 Private Bit (P)

The CSIX private bit is not used.

6.2.5 Payload Length Field

Payload length is the number of bytes in the payload of the message. For data CFrames, the payload length can have values from 1 to N, where N is the maximum payload size, and can be programmed to have a value from 64 to 120 bytes. For outgoing idle CFrames, the payload length field is set to 0. The Payload Length is valid in both the ingress and egress directions.

6.3 Idle CFrames

Idle Cframes will be transmitted during start up in standard CSIX mode. Idle Cframes will also be used during periods of inactivity to maintain ready bits and data synchronization.

If more than one Idle CFrame is needed, a dead cycle will be inserted between Idle frames to allow the SOF line to toggle by both the IXP2400 and the VSC872 transmitter. An idle CFrame consists of 32 bits; a base header and a two-byte vertical parity. For idle frames, the payload length field is always set to 0.

6.4 Data CFrames

Data Cframes are used to send data payloads through the switch fabric. This section describes unicast, multicast, and broadcast CFrame types.

Throughout the following tables, bit locations designated with shaded areas are CSIX reserved bits. Bit locations designated with cross hatching are unused portions of a defined CSIX field. In both cases, these bits can be set to zero by the transmitter and ignored by the receiver.

6.4.1 Ingress Unicast Extension Header

The ingress unicast extension header is shown in Table 6.6.

Byte Number	Bit Position															
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
2/3	Class[2:0]															
4/5							DA[11:0]									

Table 6.6 Ingress Unicast Extension Header

6.4.1.1 Class Field (Class[2:0])

A total of eight classes are supported. Classes 0x0 to 0x7 are used to select the VSC872 queue for unicast data traffic. The highest priority class is defined as 111 (binary). The lowest priority class is defined as 000 (binary).

6.4.1.2 Destination Address Field (DA[11:0])

The destination address is a 12-bit field. The destination address field is only valid in the ingress direction.

6.4.2 Egress Unicast Extension Header

The egress unicast extension header is shown in Table 6.7.

Byte Number	Bit Position															
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
2/3	Class[2:0]															
4/5																

Table 6.7 Egress Unicast Extension Header

6.4.2.1 Class Field (Class[2:0])

A total of 8 classes are supported. Classes 0x0 to 0x7 are used to select the IXP2400 egress queue for unicast data traffic. The highest priority class is defined as 111 (binary). The lowest priority class is defined as 000 (binary).

6.4.3 Ingress and Egress Multicast Extension Header

The ingress and egress multicast ID extension header is shown in Table 6.8.

Byte Number	Bit Position															
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
2/3	Class[2:0]															
4/5							Multicast ID[11:0]									

Table 6.8 Ingress and Egress Multicast Extension Header

6.4.3.1 Class Field (Class[2:0])

A total of eight classes are supported. Classes 0x0 to 0x7 are used to select the queue for multicast data traffic. The highest priority class is defined as 111 (binary). The lowest priority class is defined as 000 (binary).

6.4.3.2 Multicast ID Field (ID[11:0])

The 12-bit Multicast ID communicates a lookup tag to the switch fabric. The switch fabric uses this lookup tag to determine which set of egress queues should receive copies of the frame. This field can also be used by the egress IXP2400 for further multicast fan-out to multiple physical layer devices.

6.4.4 Ingress and Egress Broadcast Extension Header

The ingress and egress broadcast extension header is shown in Table 6.9.

Byte Number	Bit Position															
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
2/3	Class[2:0]															
4/5																

Table 6.9 Ingress and Egress Broadcast Extension Header

6.4.4.1 Class Field (Class[2:0])

A total of eight classes are supported. Classes 0x0 to 0x7 are used to select the queue for broadcast data traffic. The highest priority class is defined as 111 (binary). The lowest priority class is defined as 000 (binary).

6.5 Flow Control CFrames

Flow-control CFrames provide a finer level of flow-control than Link level flow-control by specifying the specific destination and Class that are oversubscribed. Flow-control frames are used to transmit fabric ingress queue status information to the ingress IXP2400. These flow-control frames in the egress direction are referred to as fabric-to-NPE flow control. Only Link level flow-control is used from the egress IXP2400 to the fabric egress queuing system.

A flow-control frame consists of a base header (as previously defined) with the type set to flow-control frame, a control payload, any necessary padding, and the two-byte vertical parity field. A flow-control CFrame does not contain an extension header. Instead, flow-control information is located within the payload. The basic flow-control element is the four-byte Flow Control Entry. The CSIX standard allows for a variable number of flow-control elements within the control payload.

6.5.1 Fabric to NPE Flow Control Header

The Fabric to NPE ingress flow-control header format is shown in Table 6.10.

Byte number	Bit Position															
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
1 of 4 / 2 of 4	Class[2:0]							Type[1:0]		C	P	S				
3 of 4 / 4 of 4				DA[11:0]												

Table 6.10 Fabric to NPE Flow-control Header Format (Payload Bytes)

6.5.1.1 Class Field (Class[2:0])

The class field is the same class field specified previously for unicast, multicast or broadcast extension headers.

6.5.1.2 Entry Type Field (Type[1:0])

The entry type indicates what resource is being reported. The entry type determines how the Entry is processed. The encoding of the entry type field is shown in Table 6.11.

Encoding	Name
00	Unicast
01	Multicast/Broadcast
10	N/A
11	N/A

Table 6.11 Flow Control Entry Type Values

6.5.1.3 Class Wildcard Field (C)

This is a proprietary field definition. If this bit is set to a '1', flow-control is set for all classes which overrides the Class[2:0] field.

6.5.1.4 Port Wildcard Field (P)

This is a proprietary field definition. If this bit is set to 1, flow-control is set for all ports which overrides the DA[11:0] field. For multicast and broadcast frames, this field is ignored.

6.5.1.5 Speed Bit Field (S)

In the NPE to Fabric flow control, only the most-significant bit of the speed field is used. It indicates the ability of the fabric to receive data. 0 indicates the fabric is not ready to receive data; 1 indicates the fabric is ready to receive data.

6.5.1.6 Destination Address Field (DA[11:0])

The destination address for unicast queues has the same field definition as in the unicast extension header. For multicast and broadcast flow-control frames, this field is ignored.

6.6 Parity

Both horizontal and vertical parity is supported. Checking of vertical parity is user programmable.

6.6.1 Horizontal Parity

Horizontal parity is an odd parity across 32 bits of data. The horizontal-parity bit is carried across the CSIX interface by the TxPar[1:0] and RxPar[1:0] busses. Each bit of parity covers a 32-bit data word, see CSIX Specification. For example, the transmit horizontal parity, TxPar, is calculated as follows:

$$\text{TxPar} = \text{!(TxData[31])} \wedge \text{TxData[30]} \wedge \dots \wedge \text{TxData[0]}.$$

TxPar[0] and RxPar[0] are associated with CSIX port 0.

TxPar[1] and RxPar[1] are associated with CSIX port 1.

6.6.2 Vertical Parity

Vertical parity provides an optional second set of parity bits that provide a substantial improvement (over horizontal parity alone) in the probability of error detection. To calculate vertical parity bits, the CFrame is treated as a series of 16-bit words, organized in a two-dimensional block as shown in Table 6.12. A Vertical parity bit is generated for each of the 16 bit positions (columns) in the block across all rows required to contain the CFrame payload. The resulting 16-bit Error Detecting Code (VPar) is appended to the payload. All padding will be included in the vertical parity calculation.

16-bit Words	Byte 0 Bit 7	Byte 0 Bit 6		Byte 1 Bit 0
Word 0	b(15,0)	b(14,0)		b(0,0)
Word 1	b(15,1)	b(14,1)		b(0,1)
Word m	b(15,m)	b(14,m)		b(0,m)
Vpar	Vpar15	Vpar14		VPar0

Table 6.12 Vertical Parity

$$\text{VPar}[i] = \text{!(b(i,0) } \wedge \text{ b(i,1) } \wedge \dots \wedge \text{ b(i,m))};$$

Where,

Vpar[i] = ith bit of Vertical Parity word.

b(i,m) = ith bit of data word m

m = number of 16-bit words (Excluding Vpar word) in CFrame

The 16-bit word ordering within the payload is big-endian, where two 16-bit words comprise the 32-bit wide data input. For example, the first 32 bits of data would be broken down into two 16-bit words as follows:

$$\text{Data}[31:0] = \text{b(15,0) } \dots \text{ b(0,0), b(15,1) } \dots \text{ b(0,1)}$$

Where b(i,m) = ith bit of 16-bit data word m

6.7 Dead Cycle

The purpose of a dead cycle is to prevent consecutive SOFs. A dead cycle is invalid data and a low on the SOF (i.e., de-asserted). For a dead cycle, horizontal parity is irrelevant, but is generated correctly on the VSC872 egress anyway. The VSC872 and the IXP2400 ingress have no expectation of correct horizontal parity for incoming dead cycles. In fact, on the ingress side, both horizontal and vertical parity are ignored for dead cycles by both the VSC872 and the IXP2400.

There can be more than one consecutive dead cycle. Dead cycles are inserted between all single word CFrames including idle frames. Dead cycles may or may not appear between multi-word CFrames, depending upon the internal timing between the multi-word CFrames. At least one dead cycle is always inserted following an idle frame. A dead cycle may also be inserted between two CFrames, depending on where those frames are stored in RAM.

6.8 Operation and Timing

CSIX operation and timing is defined in the same way for both the receive path and the transmit path.

6.8.1 Start-up

The IXP2400/GigaStream implementation uses the CSIX Specification startup mechanism based on the ready bits and the idle CFrame. At Power-Up or RESET, each CSIX port entity (NPU and fabric) begins transmitting idle CFrames while holding the ready bit low. When an entity detects receipt of idle CFrames, it raises the ready bit to high and continues to send idle CFrames. When an entity detects receipt of idle CFrames and a high ready bit from its connected entity, it assumes normal operation.

6.8.2 Transmission

During times when no data is being transmitted, a pattern of alternating IDLE CFrames and dead cycles is transmitted to maintain synchronization and keep the ready bits alive. See the CSIX documentation for a description of the startup and transmission state machines.

6.8.3 Pause and Resume Operation

Pause and Resume information, as represented by the ready bits, provides Link level control across the CSIX interface. De-assertion of a ready bit indicates that the destination device cannot accept another frame after the frame currently being transferred for that Link level queue. When a pause is received, the CSIX interface will complete the current frame transmission and not resume transmitting until the ready bit is reasserted.

In both devices, the queue structures will respond to a pause request within 32 clock ticks. The 32 ticks do not include the time required to finish the last transmission. The last transmission may require an additional 35 tick delay, depending on the mode of operation.

6.8.4 Parity Errors

When a parity error is detected on a CFrame, the ready bits from that CFrame will be ignored. The CSIX interface will then stop sending further data or control until it receives another CFrame that is free of errors.

6.8.5 Unexpected Start of Frame (SOF)

If an unexpected SOF is detected, the current CFrame will be discarded and the CSIX ingress will begin processing the new frame.

6.9 Link Level Flow Control

The CSIX layer 1 specification requires that the transmitting port must stop transmitting within 32 CSIX clock periods after receiving a Link level flow-control signal on the Ready Bits. Both the IXP2400 and the GigaStream VSC872 line card transceiver meet that requirement.

6.10 VOQ Flow Control Frame Latency

When one or several of the queues in the GigaStream ingress queuing structure fill up past a programmed threshold level, flow-control CFrames are generated and transmitted back to the Ingress IXP2400. The Egress IXP2400 will first receive these flow-control CFrames, and must relay them over to the Ingress IXP2400 through the CBus. This flow-control loop has been studied to determine optimal VSC872 and IXP2400 register configurations. The block diagram of the flow-control loop is shown in the figure below.

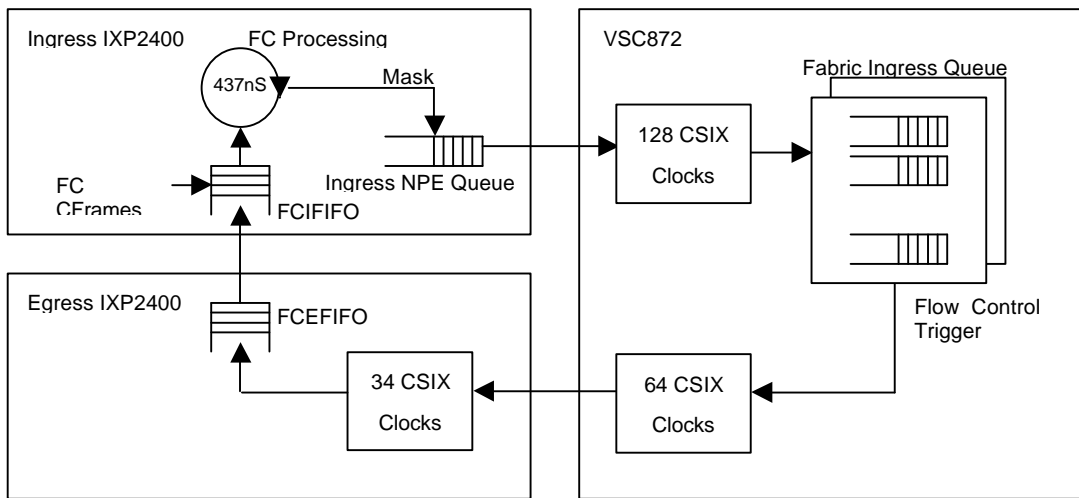


Figure 6.1 Intel IXP2400 / Vitesse VSC872 Flow Control Loop

6.10.1 Block Diagram Description

When flow-control is triggered from the fabric ingress queuing system, There is a delay of up to 64 CSIX clocks until a flow-control CFrame is generated and presented at the VSC872 egress CSIX interface. When the egress IXP2400 receives the flow-control CFrame, it will extract it from the data flow and place it in a flow-control FIFO (FCEFIFO) to send to the ingress IXP2400. The total delay from when the Egress IXP2400 receives a control CFrame until it transmits it to the ingress IXP2400 through the CBus is 34 CSIX clocks; this is based on four-bit CBus interface for the worst-case scenario.

The ingress IXP2400 maintains a FIFO (FCIFIFO) of incoming flow-control messages. A Microengine thread is assigned to process the flow-control CFrames. Since the Microengine is fully programmable, the flow-control CFrames processing latency is dependant on the application usage model and the software. For a typical OC-48 IPv4 forwarding with DiffServ application, one Microengine is assigned to process flow-control CFrames stored in the FCIFIFO and perform three other tasks of processing data packets from the normal transmit data flow. In this scenario, the Microengine processes one flow-control CFrame, if there is any, after processing every data packet, and the total latency to process one flow-control CFrame is 437ns.

6.10.2 Flow Control Modeling

The goal of the flow-control modeling effort is to provide system integrators with the optimal flow-control register configurations for various traffic conditions in both the IXP2400 and GigaStream. The model uses the IXP2400 software development kit which contains a C model of the ingress and egress NPU. A C model of the VSC872 ingress queuing structure and flow-control feedback mechanism is included with the SDK. Traffic patterns can then be run while adjusting configuration values in order to optimize performance. These results will be published.

7 CPU Interface

7.1 Intel IXP2400 CPU Interface

The IXP2400 Network Processor SlowPort bus is used for Flash EEPROM access and as a microprocessor interface to other devices, such as the VSC872 Queuing Engine. The Xscale core inside the IXP2400 is a master that can perform read/write data transfers through its SlowPort to the VSC872 CPU interface.

The IXP2400 SlowPort is an 8-bit address and data multiplexed bus, which supports multiple protocols to communicate with 8-bit, 16-bit, and 32-bit microprocessor interfaces. See the Intel IXP2400 Hardware Reference Manual for more details about the IXP2400 SlowPort signal definitions and operation.

The IXP2400 SlowPort can operate in a synchronous mode at speeds up to 50 MHz.

7.2 Vitesse GigaStream VSC872 CPU Interface

The VSC872 engine provides a CPU Processor Interface, which allows an external CPU to do the following:

- Configure the VSC872 engine.
- Read the performance counters in the VSC872 engine.

The CPU Interface is a 32-bit multiplexed address/data bus functioning as a 16-bit address bus and a 32-bit data bus. It provides the host CPU with access to the VSC872 counters and registers.

See the GigaStream Design Manual for more details on the VSC872 CPU interface.

7.3 Intel IXP2400 / Vitesse VSC872 CPU Interface Concept Drawing

A glue logic circuit implemented in a CPLD is required between the IXP2400 SlowPort and the VSC872 CPU interface. This function will convert an 8-bit address bus into a 16-bit address bus and convert an 8-bit data bus into a 32-bit data bus.

The IXP2400 SlowPort will be configured as mode 1, which is 32-bit data bus operation. The glue logic is used for bus protocol conversion between the IXP2400 and the VSC872 as well.




Figure 7.1 Intel IXP2400 / Vitesse VSC872 CPU Interface Conceptual View

8 Intel IXP2400 / Vitesse GigaStream Reference Design

The Vitesse switch fabric card (code name Anacapa) is a daughter card designed to work with the Intel IXP2400 Advanced Development Platform (ADP). Anacapa interfaces the Intel IXP2400 Network Processors to the Vitesse GigaStream switch fabric chipset. With the daughter card attached to the IXP2400 ADP, Intel and Vitesse are able to show compatibility of chipsets as well as functional demonstration of a complete OC-48 switching reference design.

8.1 Functional Description

Anacapa provides the necessary functionality to connect the Intel IXP2400 Network Processor to the Vitesse GigaStream switch fabric chipset. A functional block diagram is shown in Figure 8.1.




Figure 8.1 Anacapa Block Diagram

8.2 Modes of Operation

The primary mode of operation of Anacapa is to accept 32-bit CSIX data from the Ingress IXP2400 Network Processor, loop the traffic through the VSC872 GigaStream transceiver, and return the 32-bit CSIX data to the Egress IXP2400 Network Processor. The VSC872 GigaStream transceiver is configured, via the SlowPort interface of the Egress IXP2400 Network Processor, such that the VSC872 high-speed serial links are internally looped back, so that no external cables or connections are required. In addition, the Anacapa daughter card supports a number of other modes of operation. These modes are:

- CSIX loopback mode
- External high-speed serial link loop back mode using the SMA connects and coaxial cables.
- Single-mode packet switching with the Anacapa card connected to a VSC882 switch card, using the SMA connects and coaxial cables.

The IXP2400 Network Processor, using the VSC872 configuration registers, sets each of these modes of operation.

8.3 Hardware Block Description

8.3.1 Mictor Connectors

All signal connections between the IXP2400 ADP and the Anacapa card are made via two 144-pin Mictor connectors. Signals attached to the Mictor connectors include:

- 3.3V power (up to 6A)
- Ground
- SlowPort interface
- 32-bit CSIX interface
- JTAG port (connected, but not used by Anacapa)
- Utopia/SPI-3/POS-2 interface (not used by Anacapa)
- PCI interface (not used by Anacapa)
- I²C interface (not used by Anacapa)

All signals on the Anacapa card will be length matched to the Mictor connectors.

8.3.2 CPLD

The CPLD provides the necessary logic to translate the IXP2400 SlowPort interface to the VSC872 CPU interface. This is necessary since the IXP2400 SlowPort interface is an 8-bit multiplexed address/data bus, while the VSC872 CPU interface is a 32-bit multiplexed address/data bus. Both read and write accesses to the VSC872 from the IXP2400 Network Processor are supported.

8.3.3 CSIX Clock Generation

The clock buffer logic creates CSIX Tx and Rx clocks to the IXP2400 Network Processor, as well as a CSIX Rx clock to the VSC872 transceiver. This functionality is required because the IXP2400 Network Processor does not source the CSIX clock with the CSIX data. The master CSIX clock is generated from a 125 MHz LVTTTL oscillator and fed into VSC872 transceiver. Oscillator frequencies up to 160 MHz are possible, which allows supporting the maximum CSIX clock rate of 133 MHz for the IXP2400 Network Processor. A socket is used to allow easy swapping of oscillators. The VSC872 transceiver generates a CSIX Tx clock, which drives the zero-delay clock-buffer logic. Passive delay elements are placed in the CSIX Tx and Rx clock paths to the IXP2400 Network Processor. This allows skewing these clocks relative to the CSIX data, to meet the CSIX setup time specification of 1.5 ns. The passive delay elements can be used to adjust the clock from 0.1 ns to 2.5 ns relative to the CSIX data.

In addition, a copy of the 125 MHz CSIX clock is sent to the CPLD.

8.3.4 VSC872 Transceiver

The VSC872 transceiver provides the CSIX to high-speed serial-link conversion functionality that is required to connect the IXP2400 Network Processor to the VSC882 switch fabric device. The VSC872 has two independent CSIX ports, however, only a single port, CSIX port 0, is used in the Anacapa design. All CSIX inputs for CSIX port 1 are tied high to effectively disable the port. All CSIX port 0 outputs have source termination to reduce reflections.

Of the sixteen high-speed, differential, serial links (eight Tx and eight Rx links), only two Tx links and two Rx links are employed in the Anacapa design, since only a single CSIX port is supported. The Tx and Rx links are connected to PCB SMA connectors to support external SMA loopback mode and packet switch mode, when connected to a VSC882 switch fabric card via SMA cables. Strip-line PCB layout techniques are employed on the high-speed serial links for impedance matching to 100 Ohm differential values. In addition, trace-length matching to 25 ps is used for each serial link pair.

A low-jitter, LVTTTL oscillator is used to provide the serial link reference clock. This clock can range in frequency from 125 MHz to 155.52 MHz, and is internally multiplied by the VSC872 transceiver to provide a bit clock of from 2.125 GHz to 2.644 GHz. A socket is used to allow easy swapping of oscillators.

8.3.5 Power Generation

Six Amps at 3.3 Volts is available from IXDP2400 ADP to Anacapa via the Mictor connectors. Supply requirements for Anacapa are:

- 100 mA at 3.3 Volts
- 700 mA at 2.5 Volts
- 2000 mA at 1.8 Volts

Linear, low dropout (< 700 mV at 3A), voltage regulators are used to provide the 2.5 and 1.8 Volt supply requirements. Each linear regulator is able to supply up to 3A of current for each of the supply voltages. Bypass capacitors are used at all device supply pins to provide charge storage as well as AC noise filtering. This is especially critical for the power supplied to the VSC872 transceiver's PLL circuitry. Each supply voltage has a separate power plane as well as a ground plane for the entire board.

8.4 Thermal Considerations

In the worst case, approximately 6W must be dissipated through the VSC872 transceiver. The IXDP2400 ADP has approximately 250LFPM of airflow over the location of the Anacapa card. With a maximum case temperature of 110 °C, the VSC872 transceiver does not require a heat sink, provided that the ambient air temperature is below 58 °C and the airflow is at least 200LFPM.

$$T_{ca} = 8.66^{\circ}\text{C/W for 200 LFPM}$$

$$\text{Max Power} = 6\text{W}$$

$$\text{Max Case Temperature} = 110^{\circ}\text{C}$$

$$\text{Max_ambient} + 8.66^{\circ}\text{C/W} * 6\text{W} = 110^{\circ}\text{C}$$

$$\text{Max_ambient} = 58^{\circ}\text{C (with 200LFPM)}$$

Intel Part Number: 252142-001