# IDT®
# RapidIO Packet Tracing Using the Tsi57x

## Application Note

August 5, 2009

6024 Silver Creek Valley Road San Jose, California 95138
Telephone: (408) 284-8200 • FAX: (408) 284-3572
Printed in U.S.A.

# 1. RapidIO Packet Tracing Using the Tsi57x

This document discusses the following:



## Revision History

### 80B8030_AN003_02, Formal, August 2009

There have been no technical changes to this document. The formatting has been updated to reflect IDT.

### 80B8030_AN003_01, Formal, August 2007

This was the first version of this document.

## 1.1 Overview

RapidIO system verification and defect correction requires the appropriate tools to ensure efficiency in development and verification groups. A *packet sniffer* is a tool which allows development and verification teams to view and analyze the flow of packets within a RapidIO system. This application note describes how the RapidIO standard enables packet sniffing, and how IDT products provide the best support for packet sniffing.

The RapidIO system that a packet sniffer is connected to is known as the *monitored network* in the rest of this document.

The packets captured by a packet sniffer are named *monitored packets* in the rest of this document.

For more information on RapidIO packet sniffers, refer to the following websites:

- www.fetcorp.com for RapidFET products
- www.tek.com for RapidIO Logic Analyzer probing solutions
- www.agilent.com for RapidIO Logic Analyzer probing solutions

## 1.2 RapidIO Packet Sniffing

A packet sniffer is a tool which captures, analyzes, filters, and displays monitored packets.

This section describes the assumptions around the behavior of the monitored network that supports packet sniffing. It also describes some of the characteristics and assumptions around the behavior of the RapidIO packet sniffer.

### 1.2.1 Monitored Network

A RapidIO packet sniffer must be connected to the monitored network. Therefore, a RapidIO packet sniffer must be compliant to the RapidIO physical layer in order to receive monitored packets.

Typically, a packet sniffer does not return monitored packets to the network. To support a packet sniffer, the RapidIO network must make copies of packets, and deliver the monitored packets to the packet sniffer.

A RapidIO packet sniffer connects to one RapidIO port on the monitored network. Packets anywhere in the monitored network can be monitored using one RapidIO packet sniffer.

### 1.2.2 Packet Speed

A RapidIO packet sniffer must accept monitored packets as fast as the switch can forward packets. For example, the Tsi57*x* can forward up to 10 Gbps.

Generally, a packet sniffer captures an amount of monitored traffic and then filters and analyzes the captured data offline. This allows comprehensive analysis and searching of traffic flows which cannot be performed in real time due to the high RapidIO data rate.

Note that some logic analyzers which support capture and display of RapidIO packets can perform minimal analysis and search of the captured flow in real time, but only for purposes of starting and stopping packet capture.

### 1.2.3 Packet Acceptance

The packet sniffer must accept monitored packets regardless of errors in the packet. This capability allows error conditions to be analyzed.

The packet sniffer must accept monitored packets for all destinationID values. This capability allows monitored packets for multiple different endpoints to be captured.

### 1.2.4 Network Management

A packet sniffer must be responsible for the control and configuration of the network routing tables and packet duplication functions for the monitored traffic.

Using a packet sniffer must not impact the application or control software of the monitored network.

### 1.2.5 Logical Layer Support

Some monitored packets may require logical layer responses. A packet sniffer must never issue responses to monitored packets, because the monitored packets are copies of packets in the monitored system. If a packet sniffer issued a response to a copy of the packet, the originator of the request would receive two responses for a single request, which breaks the logical level protocol.

Packet sniffers can be implemented using standard RapidIO endpoints, which issue logical level responses to monitored request packets. The Tsi57*x* can discard these response packets based on the port they are received on, using the *per-port routing tables*. Packets received on the port where the packet sniffer is connected can be discarded based on that port's routing table. Packets received on a port different from the packet sniffer port continue to be routed to their expected destination port.

## 1.3 Packet Sniffing Using the Tsi57*x*

The IDT Tsi578, Tsi576, and Tsi574 (Tsi57*x*) RapidIO switches use RapidIO-standard multicast functionality to enable packet sniffing.

Multicast is the RapidIO method of making copies of packets.

Packets are selected based on their destinationID. The selected packets are sent to a programmable set of ports on a switch. A packet is never sent out of the port that the packet was received on. This ability enables a group of endpoints to use the same destinationID to multicast data.

The Tsi57*x* supports RapidIO-standard multicast functionality, including a standard register interface. Multicast functionality duplicates packets based on their destinationID. The duplicated packets are sent to multiple output ports.

In all diagrams, the RapidIO destinationID of Endpoint XX is XX. For example, the destinationID of Endpoint CC is CC.

Figure 1 shows how the Tsi57*x* uses multicast to support packet sniffing. In this case, all packets sent from Endpoint CC must be monitored. Packets with destination ID AA and BB are duplicated in the Tsi57*x*'s Multicast Engine (MCE). Up to 10 Gbps of received packets can be duplicated by the Tsi57*x*'s MCE to any number of ports. Since 10 Gbps is the maximum speed that a packet sniffer can receive RapidIO packets, packet duplication for traffic monitoring is performed easily by the Tsi57*x*'s MCE.

**Figure 1: Monitoring all Traffic Sent from Endpoint CC**

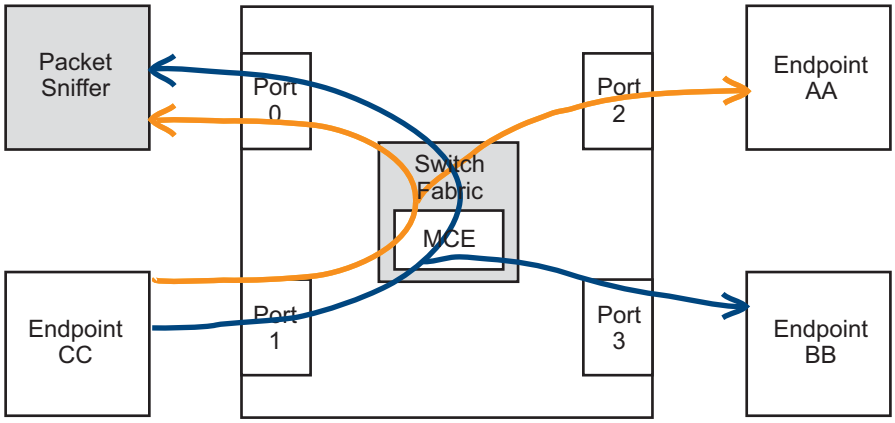


In Figure 2, all the packets sent to Endpoint CC are duplicated using the Tsi57*x*'s MCE. Monitored packets are sent to the packet sniffer, while the original packets are sent to Endpoint CC.

Packets are delivered in the same order to Endpoint CC when the Tsi57*x*'s MCE is used to monitor packets. The MCE accepts packets in the same order as port 1, and packets are duplicated in the order they are received.

**Figure 2: Monitoring all Traffic Sent to Endpoint CC**

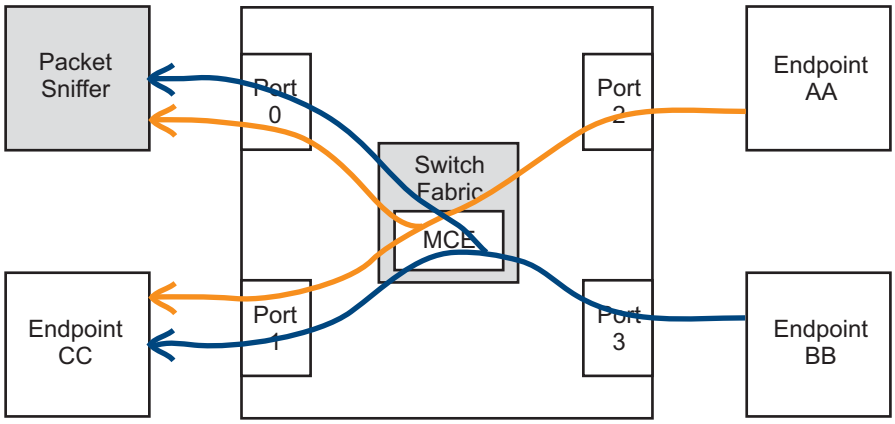

The obvious advantages of multicast for packet sniffing are made clear in Figure 3. Multicast allows all packets sent to and from a port to be monitored from a single switch port. Multiple traffic flows can be monitored in the same way.

Up to 10 Gbps of traffic, from any number of ports, can be monitored using the Tsi57*x* MCE. The Tsi57*x*'s MCE ensures that monitored traffic received by the packet sniffer is received in the same order by all other ports on the switch. This ability makes sure that the sequence of events in the system is preserved by the packet sniffer.

**Figure 3: Multicast: Monitoring Traffic Sent FROM and TO Endpoint CC**

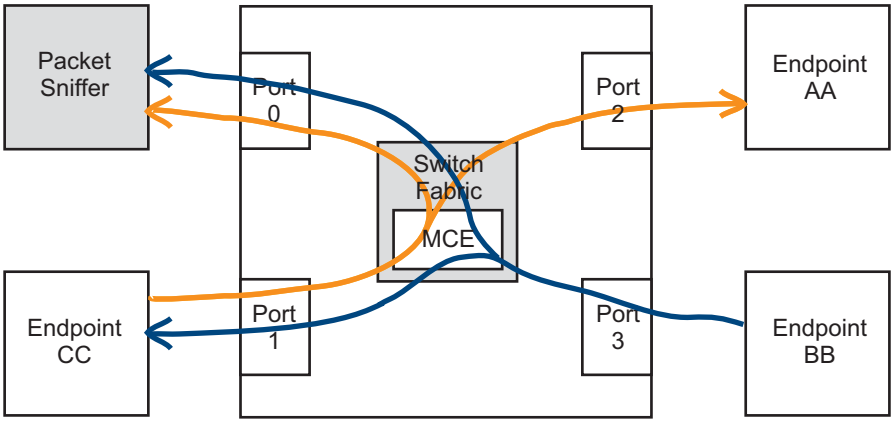

The Tsi57*x* enables packet sniffing in a multi-switch system by using a separate routing table in each port. As shown in Figure 4, monitored packets are routed to one switch where the packets are duplicated and then routed back into the system for delivery to their original destination.

In Figure 4, the blue packets have two different routes through the switch on the right. This is made possible by the use of the separate routing tables in port 3 and port 0 in the switch on the right. Port 3 routes packets with destination ID AA to port 1, while port 0 routes packets with destination ID AA to port 2. The separate routing tables in each port, in combination with the Tsi57*x*'s MCE, allows traffic to be routed to and from a Tsi57*x* MCE for duplication anywhere in the system.

**Figure 4: Multicast: Monitoring Traffic in a Multi Switch System**

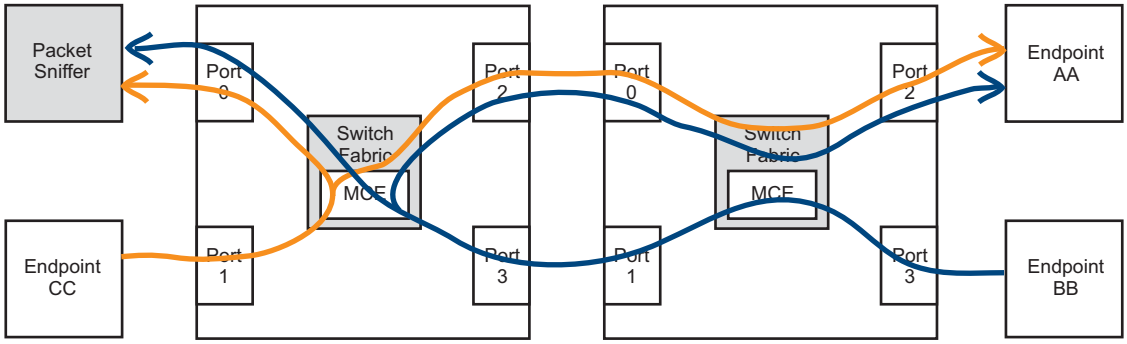


Multicast packets are never sent back out the port they were received on.

The traffic from Endpoint BB to Endpoint AA must be sent to the switch on the left using the bottom path, and received from the switch on the left using the top path. For an alternative approach, see "The Tsi57x Used in both the System and the Packet Sniffer" on page 9.

Switches in the monitored network must have per-port routing tables to support packet sniffing.

Switches that have a single routing table cannot support packet sniffing in a multi-switch system.

In summary, multicast is standardized, interoperable RapidIO functionality that provides superior packet sniffing capabilities. IDT Tsi57*x* switches Multicast capability, coupled with the Tsi57*x*'s per-port routing tables, enables packet sniffing of up to 10 Gbps of traffic anywhere in a multi-switch system.

### 1.3.1 The Tsi57*x* Used in both the System and the Packet Sniffer

There are multiple advantages when both the system and the packet sniffer use Tsi57*x* switches. Figure 5 on page 10 uses the per-port routing of the Tsi57*x*, as well as the Tsi57*x*'s MCE, to enable and implement packet sniffing anywhere in a multi-port system. Also, only the Tsi57*x*'s MCE on the packet sniffer is used. All Tsi57*x* MCEs in the monitored system can be used by the monitored system, instead of packet sniffing.

It is possible to combine the use of the Tsi57*x* MCEs in the system, and Tsi57*x*'s MCE in the packet sniffer, to monitor many different traffic flows.

Multicast packets are never sent from the port that they are received on. This is an issue if only one RapidIO port is connected to the packet sniffer. However, looping back the traffic to be monitored on a separate switch port solves this problem (see Figure 5 on page 10). This allows the MCE to duplicate packets to both the packet sniffer and to the switch port the monitored traffic was originally received on.


In Figure 5 on page 10, the traffic to be monitored is sent out port 1 and received on port 0. This is done for clarity in the diagram. It is also possible to loop port 1 back to itself, so that the packet sniffer requires the use of three ports on the Tsi57*x*.

**Figure 5: The Tsi57*x* in a System Packet Sniffing Solution**

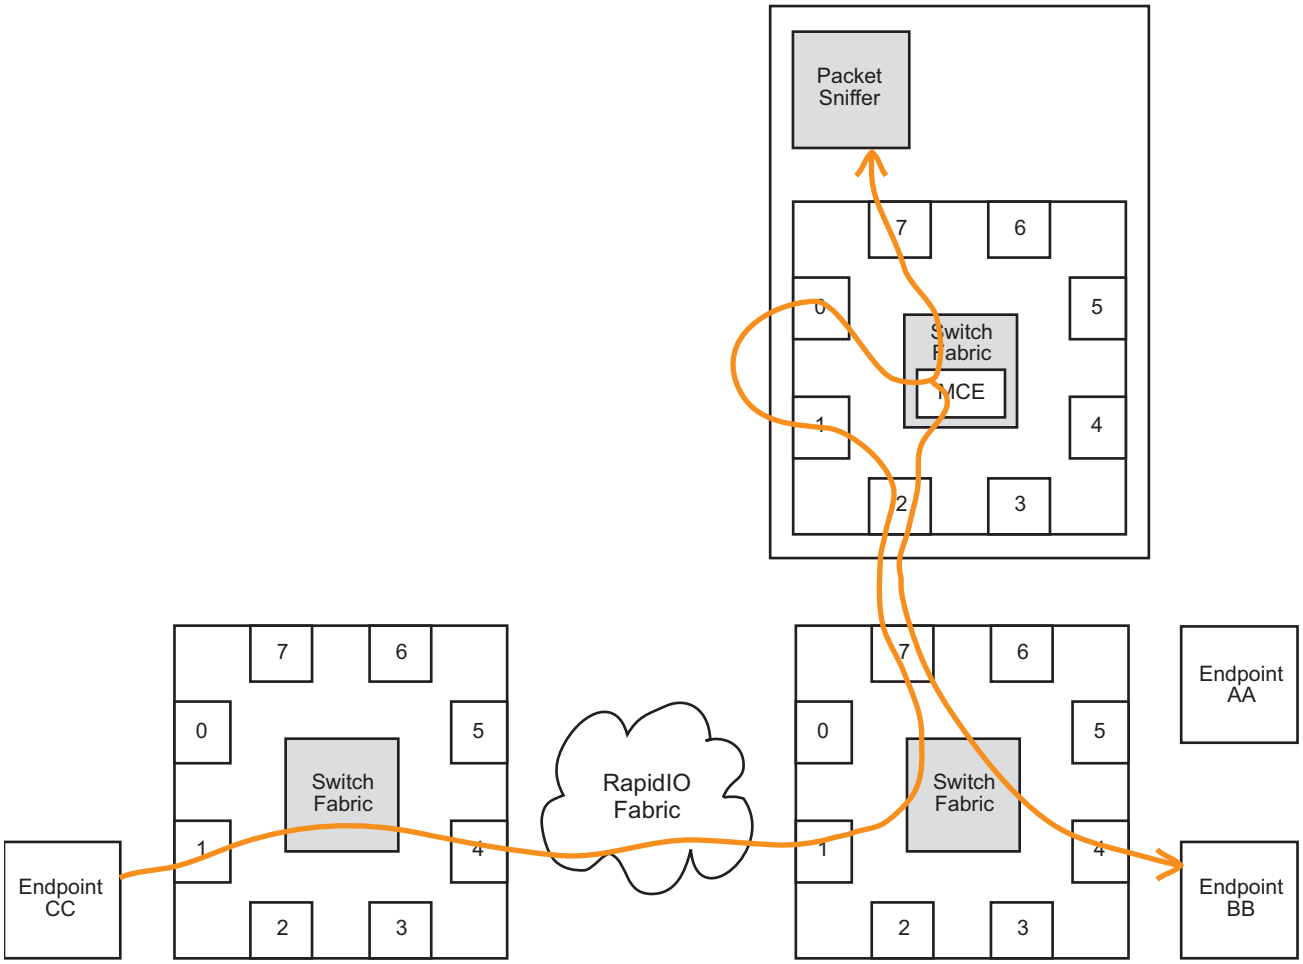

The Tsi57*x* has one copy of the multicast configuration registers (*RapidIO Specification* requirement). This implies every switch port sends packets with multicast destinationIDs to the MCE. However, the Tsi57*x* allows users to select which Multicast destination IDs are supported by individual ports, through an implementation specific programming procedure.

This procedure is implemented in the packet sniffer only - the monitored network is unaffected.

Switches in the monitored network must have per-port routing tables to support packet sniffing.

Switches that have a single routing table cannot support packet sniffing in a multi-switch system.

**Figure 6: IDT Tsi57x: System Packet Sniffing Solution**

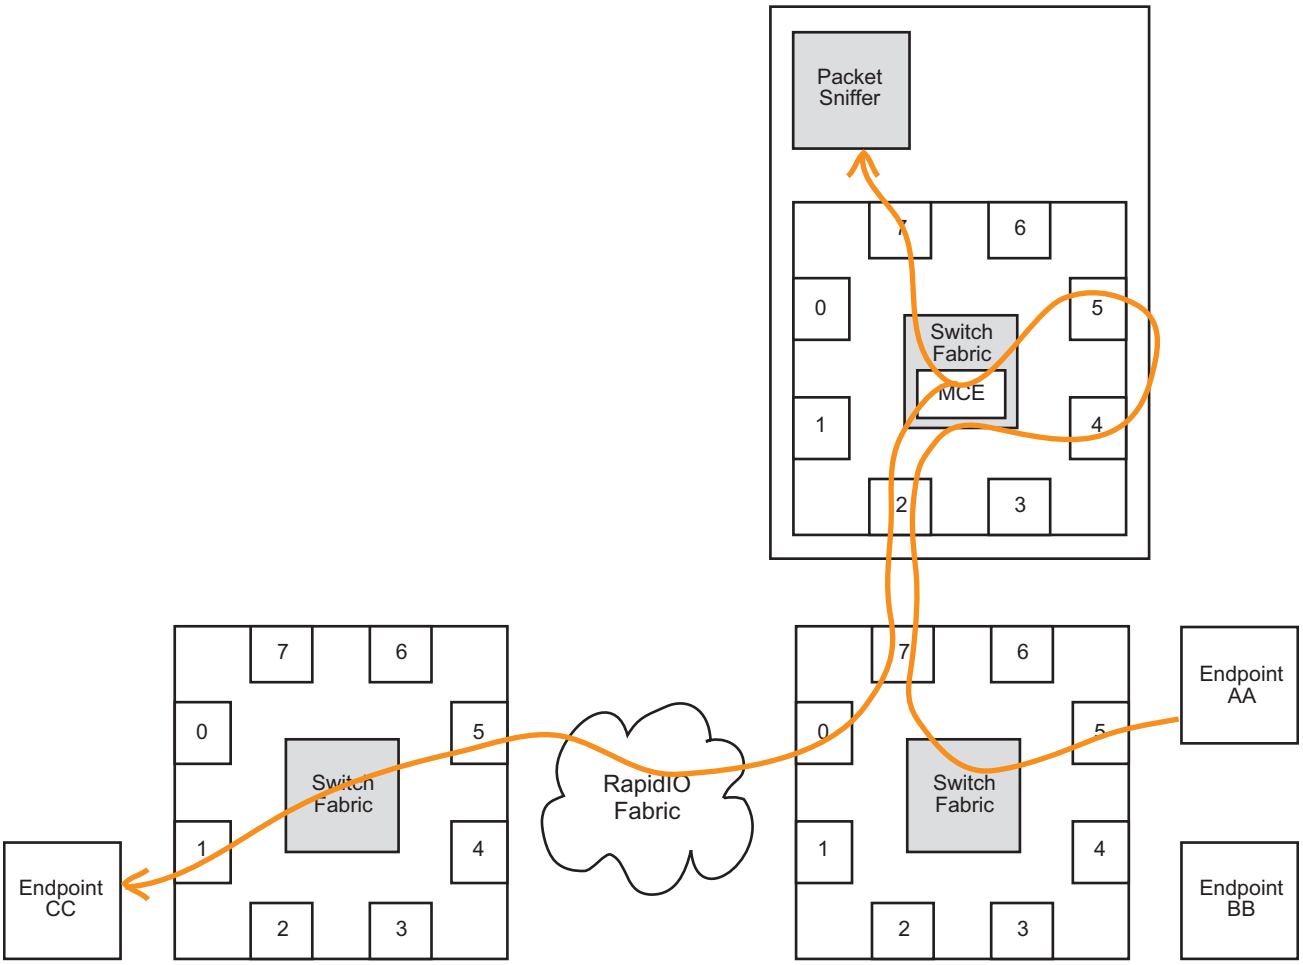


The packets shown in Figure 5 can be monitored at the same time as the packets shown in Figure 6. Up to 10 Gbps of packets from anywhere in a system can be monitored.

#### 1.3.1.1 Per-port Routing Procedure

The following procedure allows a packet with destinationID Y received on port Z to be routed to port Q instead of being multicast:

1. Remove the multicast destinationID Y from the multicast tables (if it previously existed)
2. Power down port Z, by setting the PWDN_X4 or PWDN_X1 bit as appropriate in the SMACx_DLOOP_CLK_SEL register for port Z
3. Reset port Z, by setting the SOFT_RST_X1 or SOFT_RST_X4 bit as appropriate in the SMACx_DLOOP_CLK_SEL register for port Z
4. Configure the multicast table for multicast destination ID Y
5. Remove reset from port Z, by clearing the SOFT_RST_X1 or SOFT_RST_X4 bit (which was set in step three)
6. Power-up port Z, by clearing the PWDN_X4 or PWDN_X1 bit (which was set in step two)
7. Re-initialize registers in port Z
   — Pay special attention to those registers reference in the “Default Configurations on Power Down” section of the *Tsi57x User Manual*
8. In port Z, set the routing table entry for destinationID Y to be Port Q.
9. Configure the Look Up Table (LUT) in port Z to route multicast destination ID Y to the required output port

When these steps are complete, port Z routes packets with destination ID Y to port Q without multicast. Packets with destinationID Y received on any other port are multicast.


The registers reference in this procedure are described in detail in the *Tsi57x User Manual*.

## 1.4 Packet Tracing using Port Mirroring

Another switch function that can be used to support packet tracing is called *port mirroring*. Port mirroring operates by selecting a port, known as the *mirror port*, to listen for packets sent to a port, or packets sent from a port. When port mirroring is used, the packet sniffer must be connected to the mirror port.

Port mirroring can be used to monitor all traffic sent from a switch port, or all traffic sent to a switch port. Figure 7 illustrates how port mirroring can monitor all traffic sent from a switch port. In the figure, port 1 is the mirrored port. All traffic sent from port 1 is sent to port 0 as well as its usual destination. Port 0 is the debug endpoint, while port 1 is the mirrored port.

**Figure 7: Monitoring all Traffic Sent FROM A Port Using Port Mirroring**

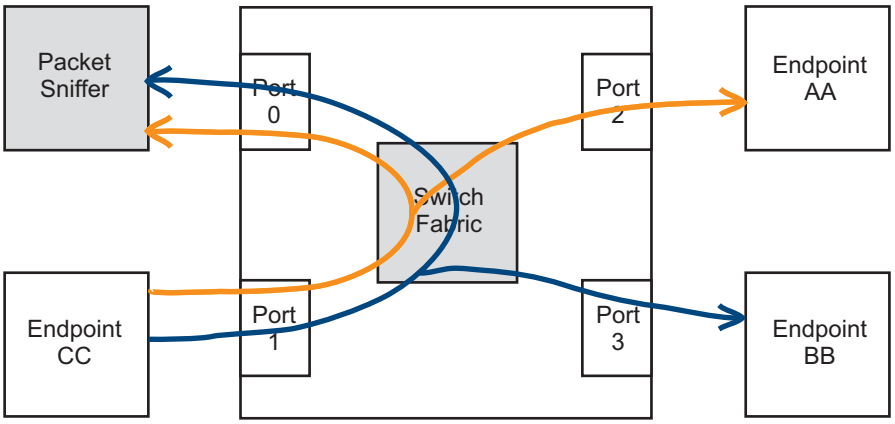


Port mirroring can also be used to monitor all traffic sent to a switch port. As shown in Figure 8, port 0 (the debug endpoint) receives copies of all packets destined for port 1. However, because port mirroring is implementation-specific functionality, it is not certain that the packets are delivered in the same order to the mirror port and to port 1.

**Figure 8: Monitoring all Traffic sent TO a Port Using Port Mirroring**

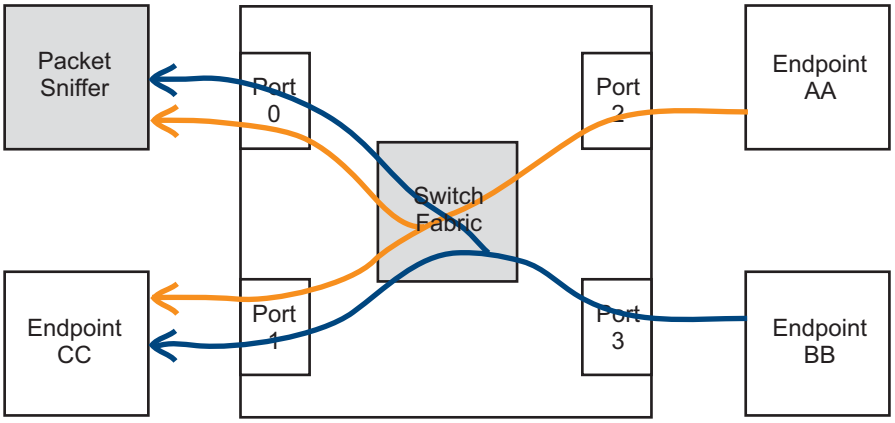

### 1.4.1 Limitations

Many packet sniffing scenarios require configurations that port mirroring functionality does not support. For example, to mirror traffic sent from port 1 and to port 1, two mirror ports are required, and two packet sniffers. Assuming that the system designer could make two mirror ports available, and two packet sniffers are purchased, the order of the packets sent and received by Endpoint CC is uncertain (see Figure 8).

The disadvantages of port mirroring are:

- Port mirroring needs a dedicated switch port to receive the monitored traffic on every board where monitoring is performed.
- Port mirroring is implementation specific functionality, so software invested in port mirroring cannot be ported to other RapidIO devices.
- A mirror port can only listen to one other port in one direction. When monitored traffic is transferred between more than two ports, more than one mirror port is required.
- Port mirroring increases the cost of debug in multi-switch customer products, because each switch requires a mirror port.
- It is not certain that the packets delivered to the debug port are in the same order as those delivered to the mirrored port.

## 1.5 Conclusion

The two types of packet tracing, packet sniffing and port mirroring, can be used in different systems with different requirements. Packet sniffing in multi-switch systems requires that the product switches support per-port routing tables. Switches with a single routing table cannot support packet sniffing in large, multi-switch systems. Packet sniffing using the Tsi57*x*'s multicast functionality and per-port routing tables supports a wide-range of applications and large, multi-switch systems.

Port mirroring is sufficient for limited debugging in small, single switch systems. It is expensive to enable packet sniffing in multi-switch systems using port mirroring.