

PRELIMINARY

DEVICE ERRATA

PMC-2001247



PMC-Sierra, Inc. PM931x ETT1™ TT1™ CHIP SETS

ISSUE 4

SCHEDULER

PM9311-UC

ETT1 and TT1 Chip Sets

Scheduler

DEVICE ERRATA

RELEASED

Issue 4: July 2001



This document is the proprietary and confidential information of PMC-Sierra Inc. Access to this information does not transfer or grant any right or license to use this intellectual property. PMC-Sierra will grant such rights only under a separate written license agreement.

Any product, process or technology described in this document is subject to intellectual property rights reserved by PMC-Sierra, Inc. and are not licensed hereunder. Nothing contained herein shall be construed as conferring by implication, estoppel or otherwise any license or right under any patent or trademark of PMC-Sierra, Inc. or any third party. Except as expressly provided for herein, nothing contained herein shall be construed as conferring any license or right under any PMC-Sierra, Inc. copyright.

Each individual document published by PMC-Sierra, Inc. may contain additional or other proprietary notices and/or copyright information relating to that individual document.

THE DOCUMENT MAY CONTAIN TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE REGULARLY MADE TO THE INFORMATION CONTAINED IN THE DOCUMENTS. CHANGES MAY OR MAY NOT BE INCLUDED IN FUTURE EDITIONS OF THE DOCUMENT. PMC-SIERRA, INC. OR ITS SUPPLIERS MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCTS(S), PROCESS(ES), TECHNOLOGY, DESCRIPTION(S), AND/OR PROGRAM(S) DESCRIBED IN THE DOCUMENT AT ANY TIME.

THE DOCUMENT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTY OR MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

ETT1, Enhanced TT1, TT1, and LCS are trademarks of PMC-Sierra, Inc.

PRELIMINARY

DEVICE ERRATA

PMC-2001247



PMC-Sierra, Inc. PM931x ETT1™ TT1™ CHIP SETS

ISSUE 4

SCHEDULER

REVISION HISTORY

Issue Number	Issue Date	Details Of Change
1	August 2000	Creation of document
2	September 2000	Added Scheduler TDM Frame Select Defect
3	March 2001	Added ESD Protection
4	July 2001	Added TT1 Port Redundancy Backpressure Defect

PRELIMINARY

DEVICE ERRATA

PMC-2001247



PMC-Sierra, Inc. PM931x ETT1™ TT1™ CHIP SETS

ISSUE 4

SCHEDULER

CONTENTS

1 INTRODUCTION 1

2 FUNCTIONAL DEFICIENCIES 2

 2.1 Go-Schedulers Command After Refresh 2

 2.2 Scheduler TDM Frame Select Defect 3

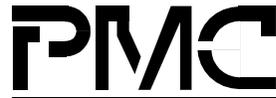
 2.3 ESD Protection Absolute Maximum Rating 5

 2.4 TT1 Port Redundancy Backpressure Defect 5

PRELIMINARY

DEVICE ERRATA

PMC-2001247



PMC-Sierra, Inc. PM931x ETT1™ TT1™ CHIP SETS

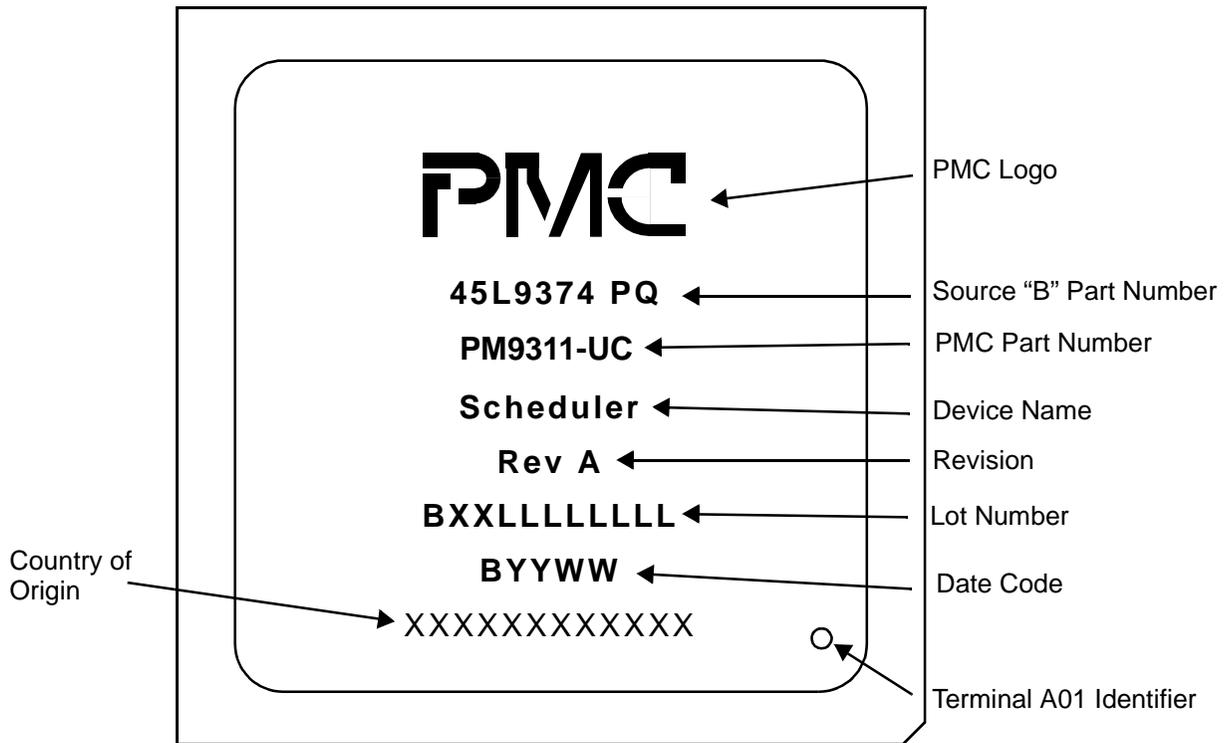
ISSUE 4

SCHEDULER

1 INTRODUCTION

The information contained in this document applies to PM9311-UC, ETT1/TT1 Scheduler, only.

Figure 1. Device Marking



TOP VIEW
SCALE: 3:1 (APPROX.)

2 FUNCTIONAL DEFICIENCIES

This section lists the known functional deficiencies of the Scheduler device as of the publication date of this document. For each deficiency, the known workaround and operating constraints, if any, are described.

Please report any functional deficiencies observed to PMC-Sierra at:

PMC-Sierra Inc.
105-8555 Baxter Place
Burnaby, BC
Canada V5A 4V7
Tel: (604) 415-6000
App Support: (604) 415-4533
Fax: (604) 415-6001

Product Information:	info@pmc-sierra.com
Applications Information:	apps@pmc-sierra.com
Web Site:	http://www.pmc-sierra.com

2.1 Go-Schedulers Command to EPP/PP After Refresh

When a Go-Schedulers command is issued to the EPP/PP to restart the Scheduler after a refresh process, the Scheduler may temporarily turn off some of its AIB links to the EPP/PP. This causes the EPP/PP to receive errors from the Scheduler. A software workaround can avoid the problem.

Description

The ETT1/TT1 switch fabric uses a refresh mechanism to restore Scheduler state after a fault has occurred or a new board has been added to a dual-Scheduler ETT1/TT1 system. The refresh process is initiated by the ETT1/TT1 CPU. The CPU sends a Refresh command to all ETT1/TT1 ports that need to have their Scheduler state refreshed.

The refresh process has two main phases:

1. Send data from EPP/PP to Scheduler.
2. The CPU writes to the Go Scheduler register in the EPP/PP, which has the effect of synchronizing the two Schedulers and restarting the arbitration of non-TDM traffic.

The defect occurs at phase 2. The Go-Schedulers command may sometimes cause some of the Scheduler AIB links to turn themselves off for up to 4 cell times, after which they turn back on. This may cause data loss, and will certainly cause unwanted interrupts.

What happens is that, internally to the Scheduler, the AIB logic associated with a given port can mis-read its configuration state and will use the configuration state of a different port. So, for example, the AIB for port 8 (inside the Scheduler) may temporarily adopt the configuration of the AIB for port 0. If port 0 is not

connected, then its AIB is turned off, and so port 8 will turn off its own AIB. The port will restore its correct configuration within four cell times (160ns), but that is too late to prevent the problem.

The deficiency is in the Scheduler logic that manages the state information for each port and AIB, and has been demonstrated in our simulations of the system and on real TT1 and ETT1 systems.

Workaround

This error only occurs for certain configurations, and never occurs if port 0 is occupied. Most of our recent testing has been conducted on a 32-port system, and that configuration also does not exhibit the problem.

The workaround to this problem is very straightforward:

Immediately before issuing a Go-Schedulers command, the software should set the following Scheduler registers:

AIB Reset	(SAIBRS) :	Set to 32'h0000_0000
Enable Port	(SENBPRT) :	Set to 32'hffff_ffff
AIB Tx Enable	(SPORTEN) :	Set to 32'hffff_ffff
Reset Port	(SPORTRS) :	Set to 32'h0000_0000
Enable non-TDM Traffic	(SNTDMEN) :	Set to 32'hffff_ffff

Immediately after the Go-Schedulers command has been issued, the software should restore the five registers to their previous value.

NOTE: The Reset Port register (SPORTRS) should never be left asserted, even for ports that are not present in the system.

2.2 Scheduler TDM Frame Select Defect

The TDM Frame Select that the Scheduler sends to the EPP/PP may change while the Scheduler is asserting TDM Sync to the EPP/PP. In some cases, this may cause the EPP/PP to store a different TDM Frame Select from that which it sends to the linecard. This causes cell loss/misdirection for up to one TDM sync period. No software workaround has been found which completely avoids the defect.

Description

When the Scheduler sends TDM Sync to the EPP/PP, it sends it for two consecutive celltimes. When switching frames, the Scheduler may send a TDM Frame Select of 0 along with the first celltime's TDM Sync, and 1 along with the second, or vice versa.

The PP always stores the TDM Frame Select from the second celltime, and sends both celltimes' TDM Syncs and TDM Frame Selects to the linecard. An idle cell from the switch may preempt the second celltime's TDM Sync/Frame Select. In that case, the linecard will have received only the first TDM Sync/Frame Select, which may have a different value from the second Frame Select that the PP will use internally for the next TDM sync period. Depending on the timing of idle cells from the switch, the linecard may see two consecutive TDM Syncs, or only one. If the linecard sees two, it should use the TDM Frame Select from the second. This reduces the probability of the problem, but does not prevent it entirely.

The EPP also stores the TDM Frame Select from the second celltime; however, unlike the PP, the EPP sends only one TDM Sync Control Packet to the linecard, with the TDM Frame Select from the first celltime. An idle cell from the switch may delay the TDM Sync Control Packet by one celltime, but will not affect the Frame Select used in the Control Packet. Therefore, if the TDM Frame Select changes mid-TDM Sync, the EPP and linecard will use different Frame Selects for the next TDM sync period.

When the EPP/PP and linecard have different Frame Selects, each uses a different TDM Frame Table. Depending on how tags are assigned in each table, there may be cell misdirection, and almost certainly will be head-of-line blocking which causes cell loss.

When not switching frames, the Scheduler will send the same TDM Frame Select in both celltimes of the two-celltime TDM Sync, so the problem will not occur.

The frequency of the problem depends on several factors:

1. How often TDM Frames are switched (since this problem can occur only in the first TDM sync period when frames are switched).
2. What method is used to drive TDM Syncs (since this determines the probability of the Scheduler changing the TDM Frame Select mid-TDM Sync).
3. PP only: the probability of an idle cell preempting the second TDM Sync/Frame Select to the linecard.

Number 1 depends on the customer's TDM programming.

Number 2 depends on whether TDM Syncs are driven by the Scheduler's TDM Control register, or by Control Packets from the linecard.

If TDM Syncs are generated by the Scheduler, frame-switching is accomplished by writing a new frame select to the TDM Control register. Since the timing of OOB writes cannot be precisely controlled, the OOB write with the new frame select may take effect mid-TDM Sync. Lab tests have shown this probability to be approximately $1/\text{TDMSyncPeriod}$ frame switches.

If TDM Syncs are driven by Control Packets from the linecard, then TDM sync period jitter will expose TDM Frame Select to changing mid-TDM Sync, even after the Scheduler has locked on to the TDM Sync frequency from the linecard. Factors that contribute to TDM sync period jitter are clock domain crossings, idle cell rates, OC-48 muxing function, and any jitter in the source of TDM Sync Control Packets. Lab tests in which the jitter is ± 3 (due to Enhanced Bithose implementation) show the probability of mid-TDM Sync change to be about $1/5$ frame switches.

Number 3 depends on the rate at which the PP is programmed to send idles to the linecard, and the TDM sync period.

The deficiency is in the Scheduler logic that sends the TDM Sync/Frame Select to the EPP/PP. and has been demonstrated on the ETT1 system.

The probability of this error is fairly low, so frame-switching tests in our simulation environment did not uncover it.

Workaround

A software workaround that completely avoids the problem has not been found. Modification of the EPP to Revision B is being implemented as a solution to the problem.

2.3 ESD Protection Absolute Maximum Rating

The ESD protection of the PM9311 Revision A Scheduler Device is rated at 500 V for production released product. Section 6.2 of the TT1 and ETT1 Switch Fabric Datasheets will be changed in future issues to reflect this rating.

2.4 TT1 Port Redundancy Backpressure Defect

The Scheduler does not correctly assert backpressure when the redundant port of a TT1 port pair is made active. This causes potential cell loss for as long as backpressure is being asserted by the redundant port. No software workaround has been found which completely avoids the defect. This defect only affects a TT1 system and not an ETT1 system.

Description

The defect is that although the rtag port is copied across between the primary and redundant ports, the priority information is not. So when the redundant port becomes the active port and the Scheduler gets a backpressure indication from the PP, it will always set backpressure for unicast Priority 0, regardless of the actual priority and traffic type that caused backpressure to be asserted. This means the Scheduler is potentially backpressuring the wrong QID. It will continue to issue grants to the QID that is backpressured by the egress PP and the cells will get dropped at the egress PP. For example, if the redundant port's PP sets backpressure for a multicast Priority 2 from port 12 then it will assert backpressure to unicast Priority 0 from port 12 and multicast traffic will not be backpressured.

Workaround

A software workaround that completely avoids the problem has not been found.