

TMS320TC1100
Digital Signal Processor
Silicon Errata

Silicon Revision 2.0, 1.0

SPRZ210I
October 2003
Revised April 2009



Copyright © 2009, Texas Instruments Incorporated

REVISION HISTORY

This silicon errata revision history highlights the technical changes made to the SPRZ210H revision to make it an SPRZ210I revision.

Scope: Applicable updates to the C64x device family, specifically relating to the TCI100 device, have been incorporated.

PAGE NO.	ADDITIONS/CHANGES/DELETIONS
7, 10	Added Power Sequencing – IO Before Core Usage Note

Contents

1	Introduction	4
1.1	Device and Development-Support Tool Nomenclature	4
1.2	Revision Identification	5
2	Silicon Revision 2.0 Known Design Exceptions to Functional Specifications and Usage Notes	6
2.1	Usage Notes for Silicon Revision 2.0	6
	HPI: HWOB Bit in HPIC Register is Writable by CPU	6
	TCP/VCP: TCP/VCP Memory Address Range Must be Accessed Using Only Doublewords (All C64x Devices)	6
	EMIFA: Dead CLKOUT4/6	6
	L1P Cache: Incorrect Update of the L1P Tag RAMs (All C64x Devices)	7
	Power Sequencing – IO Before Core	7
2.2	Silicon Revision 2.0 Known Design Exceptions to Functional Specifications	8
Advisory 2.0.1	PCI: Slave Reads With a Long Latency Can Return Bad Data	8
Advisory 2.0.2	PCI Reads May Get Ahead of PCI Writes When PCI Exceeds the Transfer Request Limit	9
3	Silicon Revision 1.0 Known Design Exceptions to Functional Specifications and Usage Notes	10
3.1	Usage Notes for Silicon Revision 1.0	10
	Electronic Discharge (ESD) Performance	10
	Power Sequencing – IO Before Core	10
3.2	Silicon Revision 1.0 Known Design Exceptions to Functional Specifications	11
Advisory 1.0.1	EMIF: PDT Write Transfers Fail When PDTWL Equals 3	11
Advisory 1.0.3	EMIF: Asynchronous RAM AC Timings Different From Data Sheet Specifications	11
Advisory 1.0.4	PCI: AC Timings Different From Data Sheet Specifications	12
Advisory 1.0.5	HPI: AC Timings Different From Data Sheet Specifications	12
Advisory 1.0.6	L2 Cache: Accesses to Mapped L2 RAM Update L2 LRU Information	13
Advisory 1.0.7	L2 Cache: L2 Controller Incorrectly Updates LRU for Accesses in L2 Cache	14
Advisory 1.0.8	JTAG: Boundary Scan IDCODE_REGISTER Incorrect	15
Advisory 1.0.9	EMIF: PDT Transfers Fail When Accessing the Same SDRAM Page as Non-PDT Transfers	15
Advisory 1.0.11	PCI/HPI: Data Corruption Can Occur When DV _{DD} Powers Up Before CV _{DD}	16

1 Introduction

This document describes the known exceptions to the functional specifications for the TMS320TCI100 digital signal processor. [See the *TMS320TCI100 Fixed-Point Digital Signal Processor* data sheet (literature number SPRS218).] Throughout this document, TMS320C64x, C64x and TCI100 refer to TMS320TCI100.

For additional information, see the latest version of the *TMS320C6000 DSP Peripherals Overview Reference Guide* (literature number SPRU190).

This document also contains “Usage Notes”. Usage Notes highlight and describe particular situations where the device’s behavior may not match presumed or documented behavior. This may include behaviors that affect device performance or functional correctness. These notes will be incorporated into future documentation updates for the device (such as the device-specific data sheet), and the behaviors they describe will not be altered in future silicon revisions.

1.1 Device and Development-Support Tool Nomenclature

To designate the stages in the product development cycle, TI assigns prefixes to the part numbers of all DSP devices and support tools. Each DSP commercial family member has one of three prefixes: TMX, TMP, or TMS (e.g., **TMS320TCI100GLZ7**). Texas Instruments recommends two of three possible prefix designators for its support tools: TMDX and TMDS. These prefixes represent evolutionary stages of product development from engineering prototypes (TMX/TMDX) through fully qualified production devices/tools (TMS/TMDS).

Device development evolutionary flow:

TMX	Experimental device that is not necessarily representative of the final device’s electrical specifications
TMP	Final silicon die that conforms to the device’s electrical specifications but has not completed quality and reliability verification
TMS	Fully qualified production device

Support tool development evolutionary flow:

TMDX	Development-support product that has not yet completed Texas Instruments internal qualification testing.
TMDS	Fully qualified development-support product

TMX and TMP devices and TMDX development-support tools are shipped against the following disclaimer:

“Developmental product is intended for internal evaluation purposes.”


TMS devices and TMDS development-support tools have been characterized fully, and the quality and reliability of the device have been demonstrated fully. TI’s standard warranty applies.

Predictions show that prototype devices (TMX or TMP) have a greater failure rate than the standard production devices. Texas Instruments recommends that these devices not be used in any production system because their expected end-use failure rate still is undefined. Only qualified production devices are to be used.

TMS320 is a trademark of Texas Instruments.
All other trademarks are the property of their respective owners.

1.2 Revision Identification

The device revision can be determined by the Die PG code marked on the top of the package. The location of the Die PG code for the GLZ package is shown in Figure 1.



NOTE: Qualified devices are marked with the letters “TMS” at the beginning of the device name, while nonqualified devices are marked with the letters “TMX” or “TMP” at the beginning of the device name.

Figure 1. Example, Die PG Codes for TMS320TCI100 (GLZ)

Silicon revision is identified by a code on the chip. The code is of the format Dxx-YMLLLLLS. If xx is 20, then the silicon is revision 2.0; if xx is 10, then the silicon is revision 1.0, etc.

Table 1. Die PG Codes

Die PG Code	Silicon Revision	Comments
20	2.0	TMX320TCI100, TMS320TCI100
10	1.0	TMX320TCI100, TMS320TCI100

2 Silicon Revision 2.0 Known Design Exceptions to Functional Specifications and Usage Notes

2.1 Usage Notes for Silicon Revision 2.0

Usage Notes highlight and describe particular situations where the device's behavior may not match presumed or documented behavior. This may include behaviors that affect device performance or functional correctness. These notes will be incorporated into future documentation updates for the device (such as the device-specific data sheet), and the behaviors they describe will not be altered in future silicon revisions.

HPI: HWOB Bit in HPIC Register is Writable by CPU

On silicon revision 2.0 and earlier, the HWOB bit in the HPIC register is writable by the CPU, even though it should be writable by the Host only. Writing the incorrect value to the HPIC register will corrupt any data passed through the HPI.

When performing HPIC writes via the CPU, be sure to write the correct value into the HWOB bit.

TCP/VCP: TCP/VCP Memory Address Range Must be Accessed Using Only Doublewords (All C64x Devices)

On silicon revision 2.0 and earlier, the TCP/VCP memory address range (0x5000 0000 to 0x5FFF FFFF) must always be accessed by doubleword requests. Even if the C64x device does not support the TCP/VCP peripherals or if the TCP/VCP coprocessors are not enabled, this memory address range must *always* be accessed via doublewords. A request of any other size will cause the chip to hang. Users should be especially careful when using Code Composer Studio™ Integrated Development Environment (IDE), because Code Composer Studio requests memory data using word-sized accesses. Do not attempt to view this memory map region (0x5000 0000 to 0x5FFF FFFF) in a Code Composer Studio window. (Note: this region can be disabled in the Code Composer Studio memory map to prevent accidental access.)

For all TMS320C64x™ DSP devices, always access the memory address range 0x5000 0000 to 0x5FFF FFFF via doublewords.

EMIFA: Dead CLKOUT4/6

On silicon revision 2.0 and earlier, there is a usage condition concerning EMIFA which can affect the functionality of CLKOUT4 and CLKOUT6.

The EMIFA Global Control register (GBLCTL) controls the logic that outputs the internal CPU/4 or CPU/6 clocks to the CLKOUT4 and CLKOUT6 pins. The GBLCTL is registered with the AECLKIN clock; therefore, without a valid EMIF clock, it is possible to have unknown values in the GBLCTL register. Furthermore, without a valid EMIF clock, the EMIFA GBLCTL will not be reset to its default value.

To avoid a dead CLKOUT4 and/or CLKOUT6, a valid clock must be provided to EMIFA at all times. This can be implemented externally using AECLKIN or by setting the BEA[17:16] pins (EMIFA_CLK_SEL bits) to select the internal CPU/4 or CPU/6.

Code Composer Studio and TMS320C64x are trademarks of Texas Instruments.

L1P Cache: Incorrect Update of the L1P Tag RAMs (All C64x Devices)

On TCI100 silicon revision 2.0 and earlier, when the CPU is executing non-cacheable code from external memory and there is snoop activity from L2 to L1P occurring at the same time, an incorrect update to the L1P Tag RAM can occur.

Snoop activity from L2 to L1P can be generated two ways:

1. EDMA/QDMA activity to L2
2. Block cache invalidates initiated in L2

When there is a non-cacheable L1P fetch that is returned from L2 to L1P, **and** there is a snoop from L2 in the very next cycle, then the snoop tag read interferes with the tag/status RAM write for the non-cacheable data. This interference causes the tag RAMs to be incorrectly updated with the tag for that line, rather than discarding the write to the tags. When the NEXT access to that non-cacheable line in L2 occurs, the L1P incorrectly registers this as a hit and transfers data from the L1P rather than the desired external data.

To *avoid* an incorrect update of the L1P tag RAMs, do the following as best practice:

1. While executing code from non-cacheable space, **do not** perform either EDMA/QDMA transfers to L2 **or** block cache invalidates initiated in L2
2. Mark program code as cacheable as soon as possible.

Power Sequencing – IO Before Core

If customers are using IO before Core, which was initially supported, it is required to connect the Reserved pin W25 to ground due to insufficient pulldown strength. If IO comes up before Core, and W25 is not at ground level, this can result in the device coming up in an improper state. To ensure this does not occur, it is required to connect W25 to V_{SS} . If this is not done, then the power sequencing must be changed to Core before IO to ensure the internal state is set properly.

2.2 Silicon Revision 2.0 Known Design Exceptions to Functional Specifications

Advisory 2.0.1

PCI: Slave Reads With a Long Latency Can Return Bad Data

Revision(s) Affected: 2.0 and earlier

Details: When an external master attempts to read memory from the DSP, it is issued a “Retry”. The PCI will then go prefetch a FIFO’s worth (32 words) of data. If this data takes an exceptionally long time to fetch (approximately 32K PCI cycles, ~1 ms @ 33-MHz PCI), the PCI port can mishandle the return data and “delete” the first word of a PCI frame. The data returned to the master is address-shifted by one 32-bit word. For example, if {0,1,2,3,4, ...} is expected and {1,2,3,4,5 ...} is returned.

Data will only take that long to fetch if the access is to a very slow external memory, or there are large, slow DMAs using the same priority queue as PCI (which, by default, is medium). For performance reasons, the separation of DMA traffic is recommended.

For more detailed information on the EDMA peripheral, EDMA performance, and EDMA performance data, see the following reference guide and application notes:

- *TMS320C6000 DSP Enhanced Direct Memory Access (EDMA) Controller Reference Guide* (literature number SPRU234)
- *TMS320C64x EDMA Architecture* Application Report (literature number SPRA994)
- *TMS320C6000 EDMA IO Scheduling and Performance* Application Report (literature number SPRAA00)
- *TMS320C64x EDMA Performance Data* Application Report (literature number SPRAA02)

PCI slave reads that start in less than 32K PCI cycles will not return bad data.

Workaround:

Do not use PCI to directly read from exceptionally slow external memories.

Do not put any other DMA activity on the same priority level as PCI.

If system considerations force a user to put other DMA activity on PCI’s level, put only DMA traffic that will complete within the 32K PCI clock cycle limit.

Advisory 2.0.2*PCI Reads May Get Ahead of PCI Writes When PCI Exceeds the Transfer Request Limit*

Revision(s) Affected: 2.0 and earlier

Details: When the PCI port exceeds its allocation of EDMA Transfer Requests (TRs), it can re-order read and write memory transactions. This can cause PCI slave reads to return “stale” data.

When the PCI port reaches its limit for the number of outstanding TRs, it must wait for previously issued TRs to complete. While the PCI port is waiting, any pending write data and/or read requests will be held pending in internal buffers. When a TR allocation is available, any pending read requests are serviced before any pending write data, regardless of the order in which the write data and read request were received. This can potentially allow the read request to get ahead of the write data, and return “stale” results.

Workaround: TI recommends that good EDMA resource allocation be used to prevent the problem. For detailed information on EDMA resource allocation, refer to the *TMS320C64x EDMA Architecture Application Report* (literature number SPRA994).

The following guidelines can help prevent PCI from running out of available TRs:

- Do not place large or slow transfers on EDMA priority levels at or above the PCI's priority level.
- Increase PCI's TR allocation limit. For detailed information, refer to the *TMS320C6000 DSP Peripheral Component Interconnect (PCI) Reference Guide* (literature number SPRU581).

The following guidelines can ensure correct read data:

- Do not read any of the previous 32 words written
- Write 32 words of “dummy” data after writing “real” data

3 Silicon Revision 1.0 Known Design Exceptions to Functional Specifications and Usage Notes

3.1 Usage Notes for Silicon Revision 1.0

Usage Notes highlight and describe particular situations where the device's behavior may not match presumed or documented behavior. This may include behaviors that affect device performance or functional correctness. These notes will be incorporated into future documentation updates for the device (such as the device-specific data sheet), and the behaviors they describe will not be altered in future silicon revisions.

Electronic Discharge (ESD) Performance

The TM320TCI100 devices were subjected to Human Body Model (HBM) ESD testing per EIA/JESD22-A114. Tests were passed at ± 2 KV with the exception of a single pin combination. When the RSV pin (G2) is stressed negative to the PLLV pin (J6) the passing voltage is 750 V. These two pins are not adjacent on the package and no more than normal ESD handling precautions are required.

The device has also been subjected to Charged Device Model (CDM) ESD testing per EIA/JESD22-C101, passing at the ± 500 V level.

Power Sequencing – IO Before Core

If customers are using IO before Core, which was initially supported, it is required to connect the Reserved pin W25 to ground due to insufficient pulldown strength. If IO comes up before Core, and W25 is not at ground level, this can result in the device coming up in an improper state. To ensure this does not occur, it is required to connect W25 to V_{SS} . If this is not done, then the power sequencing must be changed to Core before IO to ensure the internal state is set properly.

3.2 Silicon Revision 1.0 Known Design Exceptions to Functional Specifications

Advisory 1.0.1

EMIF: PDT Write Transfers Fail When PDTWL Equals 3

Revision(s) Affected: 1.0

Details: During a PDT write transfer, the $\overline{\text{PDT}}$, $\overline{\text{PDTA}}$, $\overline{\text{PDTDIR}}$, $\overline{\text{SDWE}}$, and $\overline{\text{SDCAS}}$ signals will not be driven to their appropriate states when a non-PDT write is followed by a PDT write to a different bank. The incorrect behavior of $\overline{\text{PDT}}$, $\overline{\text{PDTA}}$, $\overline{\text{PDTDIR}}$, $\overline{\text{SDWE}}$, and $\overline{\text{SDCAS}}$ can result in data corruption, as well as bus contention. This only occurs when PDTWL is set equal to 3 in the PDTCTL register.

Workaround: When performing both non-PDT writes and PDT writes to the same CE space, do not set PDTWL equal to 3.

Advisory 1.0.3

EMIF: Asynchronous RAM AC Timings Different From Data Sheet Specifications

Revision(s) Affected: 1.0

Details: The following “timing requirements for asynchronous memory cycles for EMIFA module” table reflects the differences between the actual TCI100 silicon revision 1.0 values and the values specified in the *TMS320TCI100 Fixed-Point Digital Signal Processor Data Sheet* (literature number SPRS218B or higher) for the EMIF peripheral – Asynchronous RAM.

timing requirements for asynchronous memory cycles for EMIFA and EMIFB modules†‡§ (see Figure 21 and Figure 22)

NO.		DATA SHEET VALUES		REVISION 1.0 ACTUAL VALUES		UNIT
		MIN	MAX	MIN	MAX	
7	$t_{h(\text{EKO1H-ARDY})}$ Hold time, ARDY valid after ECLKOUTx high	1		1.5		ns

Advisory 1.0.4*PCI: AC Timings Different From Data Sheet Specifications***Revision(s) Affected:** 1.0

Details: The following “timing requirements for PCI inputs” table reflects the differences between the actual TCI100 silicon revision 1.0 values and the values specified in the *TMS320TCI100 Fixed-Point Digital Signal Processor Data Sheet* (literature number SPRS218B or higher) for the PCI peripheral.

*PCI: AC Timings Different From Data Sheet Specifications (Continued)***timing requirements for PCI inputs (see Figure 48)**

NO.		DATA SHEET VALUES		REVISION 1.0 ACTUAL VALUES		UNIT
		MIN	MAX	MIN	MAX	
6	$t_{h(IV-PCLKH)}$ Hold time, input valid after PCLK high	0		0.9		ns

Advisory 1.0.5*HPI: AC Timings Different From Data Sheet Specifications***Revision(s) Affected:** 1.0

Details: The following “switching characteristics over recommended operating conditions during host-port interface cycles” table reflects the differences between the actual TCI100 silicon revision 1.0 values and the values specified in the *TMS320TCI100 Fixed-Point Digital Signal Processor Data Sheet* (literature number SPRS218B or higher) for the HPI peripheral.

switching characteristics over recommended operating conditions during host-port interface cycles†‡ (see Figure 39 through Figure 46)

NO.	PARAMETER	DATA SHEET VALUES		REVISION 1.0 ACTUAL VALUES		UNIT
		MIN	MAX	MIN	MAX	
8	$t_{d(HDV-HRDYL)}$ Delay time, HD valid to \overline{HRDY} low	-3		-3.5		ns

Advisory 1.0.6*L2 Cache: Accesses to Mapped L2 RAM Update L2 LRU Information*

Revision(s) Affected: 1.0

Details: CPU accesses to L2 RAM addresses incorrectly cause updates to the “Least-Recently Used” (LRU) state information in the L2 cache. This may cause an increased number of L2 cache misses in some systems.

The L2 cache implements a 4-way set-associative cache. The cache uses the (LRU) information to determine what “way” within each set is least recently used. When the CPU accesses data in L2 cache, the L2 controller determines what “way” holds the data in that set, and marks that “way” as most-recent. When allocating a new line in the cache, the L2 controller evicts the line in the set from the least-recently used “way” under the assumption that more-recently accessed data is more relevant.

When the CPU accesses data in L2 SRAM, either via program fetches or data accesses, the L2 cache is incorrectly updated by the L2 controller. The L2 controller updates the LRU as if the access was to “way 0” in the cache. This causes the LRU history to *not* reflect the actual sequence of accesses to L2 cache. As a result, the L2 may not choose the actual least-recently used line during an eviction.

Only CPU accesses to L2 RAM cause this update to LRU information in L2 cache. DMA accesses to L2 RAM *do not* trigger updates to the L2 LRU information.

Workaround: Perform one of the following three workarounds:

- Choose an L2 cache size that fits your cached working set. The advisory primarily impacts programs that are significantly larger than the L2 cache size.
- Explicitly remove cached contents from L2 when finished with them. The L2 cache allocates “invalid” lines within each set before consulting the LRU. Programs may do this using “block invalidate” or “block writeback-invalidate” commands in the L2 cache.
- Lay out buffers in L2 RAM so they *do not* conflict with buffers or code held in L2 cache. L2 RAM addresses map onto L2 sets in the same manner as external memory addresses.

For more detailed information on the organization and manipulation of the L2 cache, see the *TMS320C64x DSP Two-Level Internal Memory Reference Guide* (literature number SPRU610).

Advisory 1.0.7*L2 Cache: L2 Controller Incorrectly Updates LRU for Accesses in L2 Cache*

Revision(s) Affected: 1.0

Details: The L2 cache implements a 4-way set-associative cache. The cache uses the “Least-Recently Used” (LRU) information to determine what “way” within each set is least recently used. When the CPU accesses data in L2 cache, the L2 controller determines what “way” holds the data in that set, and marks that “way” as most-recent. When allocating a new line in the cache, the L2 controller evicts the line in the set from the least-recently used “way” under the assumption that more-recently accessed data is more relevant.

For this advisory, CPU accesses which hit L2 cache *do not* correctly update the LRU information for the set accessed. Instead of storing the LRU information back to the set being accessed, the L2 controller stores the information to 3 adjacent sets.

LRU information is stored in groups of 4 sets. The 3 adjacent sets affected by the current set are defined as follows:

- Group 1 contains sets 0, 1, 2, 3
- Group 2 contains sets 4, 5, 6, 7
- Etc.

For example, during an access to set 5, the L2 controller incorrectly stores the LRU information to sets 4, 6, 7.

As a result of this issue, repeated misses to the same set with no intervening accesses to adjacent sets will allocate from the same “way”. This can make the L2 cache appear to “thrash”. A series of misses to consecutive sets in L2 cache may appear to allocate with reduced associativity; that is, L2 could appear to behave as a 2-way or direct-mapped cache.

Workaround: Perform one of the following three workarounds:

- Choose an L2 cache size that fits your cached working set. The advisory primarily impacts programs that are significantly larger than the L2 cache size.
- Explicitly remove cached contents from L2 when finished with them. The L2 cache allocates “invalid” lines within each set before consulting the LRU. Programs may do this using “block invalidate” or “block writeback-invalidate” commands in the L2 cache.
- Offset external buffers that are accessed as part of the same working set so that accesses to the buffers are at least 4 L2 sets apart (512 bytes). This will prevent the buffers from “thrashing” each other in L2 cache.

For more detailed information on the organization and manipulation of the L2 cache, see the *TMS320C64x DSP Two-Level Internal Memory Reference Guide* (literature number SPRU610).

Advisory 1.0.8*JTAG: Boundary Scan IDCODE_REGISTER Incorrect***Revision(s) Affected:** 1.0

Details: When reading the Boundary Scan IDCODE_REGISTER, an incorrect value is returned. The upper 16 bits of the Boundary Scan IDCODE_REGISTER, which contain the variant field (**silicon revision** information) and a portion of the **device id** field, are incorrect. This problem is due to data corruption in the reset configuration logic.

Workaround: Disregard the upper 16 bits of data contained in the Boundary Scan IDCODE_REGISTER. The lower 16 bits of the Boundary Scan IDCODE_REGISTER, which include the lower 4 bits of the **device id** field, can still be read.

This advisory issue will be fixed in the next **major** silicon revision.

Advisory 1.0.9*EMIF: PDT Transfers Fail When Accessing the Same SDRAM Page as Non-PDT Transfers***Revision(s) Affected:** 1.0

Details: When PDT and non-PDT transfers occur to the same SDRAM page, $\overline{\text{PDTA}}$, $\overline{\text{PDT}}$, and PDTDIR may not be driven to their appropriate state. The incorrect behavior of these signals can result in PDT data corruption.

Workaround: Place all PDT transfers, whether reads or writes, in a memory range that is an aliased version of the physical SDRAM. For example, if SDRAM is in CE0 and is 128 Mbytes (MB) in depth, then the functional addressable space is 0x8000 0000 through 0x87FF FFFF and all normal CPU and non-PDT DMA transfers should access this memory range. The “aliased” view of the SDRAM is at address 0x8800 0000 through 0x8FFF FFFF and must be used for all PDT transfers. Similarly, if SDRAM is 64 MB in depth, the functional addressable view is 0x8000 0000 through 0x83FF FFFF and the “aliased” view is 0x8400 0000 through 0x87FF FFFF. The aliased view accesses the same underlying physical address as the functional view.

The address space for the “aliased” view can be created by bit-wise ORing the “logical address” (functional address) in use as follows.

- For 128 MB, OR with 0x0800 0000
- For 64 MB, OR with 0x0400 0000
- For 32 MB, OR with 0x0200 0000
- For 16 MB, OR with 0x0100 0000

This workaround is **ONLY** applicable if the CE space has less than or equal to 128 MB of SDRAM connected to it. If a CE space is full (maximum addressable space is 256 MB), then that CE space cannot support PDT transfers.

Advisory 1.0.11	<i>PCI/HPI: Data Corruption Can Occur When DV_{DD} Powers Up Before CV_{DD}</i>
------------------------	---

Revision(s) Affected: 1.0

Details: **Advisory 1.0.11- PCI/HPI: Data Corruption Can Occur When DV_{DD} Powers Up Before CV_{DD} (for all device speeds) supersedes Advisory 1.0.2 – PCI/HPI: Data Corruption Can Occur When CV_{DD} = 1.1 V.**

When DV_{DD} powers up before CV_{DD}, the PCI/HPI I/O buffers can be initialized incorrectly – causing them to be stuck in a Hi-Z condition. This can result in incorrect data on the bus, deadlocked interface or other undesirable operation.

Workaround: One of the following options is available as a Workaround:

- Power up the CV_{DD} supply before the DV_{DD} supply [see timing requirements for power-on sequence table and Figure 2] *or*
- Do not use the PCI and/or HPI peripherals. Do not use GP[15:9] as outputs.

timing requirements for power-on sequence† (see Figure 2)

NO.	Description	-600 , -720		UNIT
		MIN	MAX	
1	t _d (CV _{DD} -DV _{DD}) Delay time, CV _{DD} supply ready to DV _{DD} supply ramp start	0.5	200	ms

† Delay time from CV_{DD} is referenced to the supply reaching its minimum operating voltage.




Figure 2. Power-On Sequence Timing

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

TI products are not authorized for use in safety-critical applications (such as life support) where a failure of the TI product would reasonably be expected to cause severe personal injury or death, unless officers of the parties have executed an agreement specifically governing such use. Buyers represent that they have all necessary expertise in the safety and regulatory ramifications of their applications, and acknowledge and agree that they are solely responsible for all legal, regulatory and safety-related requirements concerning their products and any use of TI products in such safety-critical applications, notwithstanding any applications-related information or support that may be provided by TI. Further, Buyers must fully indemnify TI and its representatives against any damages arising out of the use of TI products in such safety-critical applications.

TI products are neither designed nor intended for use in military/aerospace applications or environments unless the TI products are specifically designated by TI as military-grade or "enhanced plastic." Only products designated by TI as military-grade meet military specifications. Buyers acknowledge and agree that any such use of TI products which TI has not designated as military-grade is solely at the Buyer's risk, and that they are solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI products are neither designed nor intended for use in automotive applications or environments unless the specific TI products are designated by TI as compliant with ISO/TS 16949 requirements. Buyers acknowledge and agree that, if they use any non-designated products in automotive applications, TI will not be responsible for any failure to meet such requirements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products

Amplifiers	amplifier.ti.com
Data Converters	dataconverter.ti.com
DLP® Products	www.dlp.com
DSP	dsp.ti.com
Clocks and Timers	www.ti.com/clocks
Interface	interface.ti.com
Logic	logic.ti.com
Power Mgmt	power.ti.com
Microcontrollers	microcontroller.ti.com
RFID	www.ti-rfid.com
RF/IF and ZigBee® Solutions	www.ti.com/lprf

Applications

Audio	www.ti.com/audio
Automotive	www.ti.com/automotive
Broadband	www.ti.com/broadband
Digital Control	www.ti.com/digitalcontrol
Medical	www.ti.com/medical
Military	www.ti.com/military
Optical Networking	www.ti.com/opticalnetwork
Security	www.ti.com/security
Telephony	www.ti.com/telephony
Video & Imaging	www.ti.com/video
Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2009, Texas Instruments Incorporated