AM1806 ARM Microprocessor Silicon Revisions 2.3, 2.1 and 2.0

Silicon Errata

Literature Number: SPRZ314H
March 2010–Revised September 2014

AM1806 ARM Microprocessor Silicon Revisions 2.3, 2.1 and 2.0

1 Introduction

This document describes the known exceptions to the functional specifications for the AM1806 ARM Microprocessor. For more detailed information, see the *AM1806 ARM Microprocessor* data manual (literature number SPRS658).

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 AM1xxx processors and support tools. Each commercial AM1xxx platform member has one of three prefixes: X, P, or null (no prefix). 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 (TMDX) through fully qualified production devices/tools (TMDS).

Device development evolutionary flow:

X	Experimental device that is not necessarily representative of the final device's electrical specifications
P	Final silicon die that conforms to the device's electrical specifications but has not completed quality and reliability verification
NULL	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

X and P devices and TMDX development-support tools are shipped against the following disclaimer:

"Developmental product is intended for internal evaluation purposes."

NULL 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 (X or P) 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.

Copyright © 2010–2014, Texas Instruments Incorporated

1.2 Revision Identification

The silicon revision is identified by a code marked on the package. Please see the figure Device nomenclature in the data manuals Packaging information section for coding details. Table 1 describes the relationship between device revision code and silicon revision

Table 1. Device Revision Codes

DEVICE REVISION CODE (xx)	SILICON REVISION	COMMENTS
E	2.3	-
B	2.1	Silicon Revision 2.1 can be differentiated from Revision 2.0 by a top side '21' marking on the package. See Figure 1 for details. Additionally the ROM ID code of may be used to distinguish Revision 2.1 from Revision 2.0
B	2.0	-

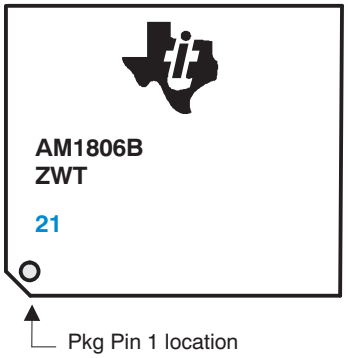


Figure 1. Device Markings

Copyright © 2010–2014, Texas Instruments Incorporated

2 Silicon Revision 2.3 Usage Notes and Known Design Exceptions to Functional Specifications

This section describes the usage notes and advisories that apply to silicon revision 2.3 of the device.

2.1 Usage Notes for Silicon Revision 2.3

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 usage 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.

2.1.1 USB0: Generic RNDIS Usage Note

On all silicon revisions, when using Generic RNDIS mode, the user should ensure that the DMA configuration has completed prior to the host starting a transfer. This condition is sometimes violated when performing a back-to-back data transfers (not transactions). If a new transfer is scheduled by a host while the device is working on the previous transfer and the data transfer size for the new transfer is different than the previous transfer data size, then there exists a contention between the two transfer sizes creating undesired behavior resulting with a DMA lock up. A case in point where this violation could happen is demonstrated by the example below.

A user configures the DMA in Generic RNDIS mode expecting a data size of 512 bytes or less from a host. The host sends 512 bytes or less of data to the device. While the device is in the process of working on the received data to figure out the size of the next data transfer, the host starts a new data transfer addressing the same endpoint. Since the endpoint FIFO is empty, the device accepts the data and the DMA starts to transfer the received data from the receive FIFO to memory. At the same time, the application on the device side finishes and figures out the next transfer data size (using the data received from previous transfer) and reconfigures the Generic DMA Size register for the second transfer. If the second transfer size is different from the first transfer size, the contention happens at this point. The host has already started the second transfer prior to the device re-configuring the DMA parameters. The application on the device side, updates the DMA size register content for the second transfer while the DMA is in the middle of the second transfer using the DMA size register content of the first transfer. This effectively results with altering the DMA size register content while the DMA is in the middle of a transfer. Changing DMA parameters while in the middle of a transfer is not allowed and when done it will create undesirable outcomes.

Workaround: This is not a bug and for this reason, there exists no workaround. This is a caution for the user to be aware of this issue and hence to ensure that this scenario is avoided. If there exists an idle time in between the two back-to-back transfers, this issue will not exist. When expecting a back-to-back transfer where RNDIS mode can not be used, the user needs to use TRANSPARENT mode. When using TRANSPARENT mode, the application will be receiving more interrupts, i.e., interrupt will be generated on each USB packets as opposed to receiving a single interrupt on the completion of a transfer.

2.1.2 USB0: Isochronous Interrupt Loading Usage Note

On all silicon revisions, when the USB Controller Endpoint is enabled to handle Isochronous type of transfer, the controller supports a single configuration for interrupt generation, which is for interrupt to be generated for every ISO packet received or sent, i.e., transfer size is equal to packet size. The option of generating interrupt on multiple ISO packets received or sent is not available. Since ISO transfer can be scheduled to happen on every micro-frame or frame, the number of interrupts generated could overwhelm the system. This is not a problem as long as there is enough CPU power available to handle all interrupts. However, some applications may be running low on available CPU time and may desire to service/process multiple ISO packets at a time. The option for handling ISO interrupts in a batch is not available. The user should ensure that enough CPU power is available to handle all ISO interrupts in order to avoid missing interrupts resulting with missing ISO packets.

Copyright © 2010–2014, Texas Instruments Incorporated

2.1.3 LCDC: Underflow During Initialization

On all silicon revisions during LCDC initialization, there is the potential for a FIFO underflow condition to occur. A FIFO underflow condition occurs when the input FIFO is completely empty and the LCDC raster controller logic that drives data to the output pins attempts to fetch data from the FIFO. When a FIFO underflow condition occurs, incorrect data will be driven out on the LCDC data pins.

An underflow condition will occur if the DDR2/mDDR Controller issues a refresh command to the SDRAM memory during LCDC start up/initialization. The error condition will be captured in the LCD Status register (LCD_STAT) in the FIFO underflow status (FUF) bit field. If the FUF_EN bit is enabled in Raster Control Register (RASTER_CTRL), the LCDC will send an interrupt to the CPU.

The FIFO underflow described above is not expected to be a common occurrence because of the unlikely alignment of events required to produce the underflow condition.

The DDR2/mDDR Controller hardware automatically schedules refresh commands to the SDRAM memory. Therefore, it is not possible for the user/application code to schedule DDR2/mDDR Controller refresh commands to prevent them from being initiated during LCDC start up. This means that it is not possible to prevent a DDR2/mDDR Controller refresh occurrence during the start up of the LCDC.

Software should poll the FUF bit field in the LCD_STAT register to check if an error condition has occurred or service the interrupt if FUF_EN is enabled when FUF occurs. If the FUF bit field has been set to 1, this will indicate an underflow condition has occurred and then the software should execute a reset of the LCDC via the LPSC.

This problem may occur if the LCDC FIFO threshold size (LCDDMA_CTRL[TH_FIFO_READY]) is left at its default value after reset. Increasing the FIFO threshold size will reduce or eliminate underflows. Setting the threshold size to 256 double words or larger is recommended.

2.1.4 Emulation Suspend Usage Note

Device peripherals often include an emulation suspend function that gracefully halts peripheral activity. This function is activated when the target CPU is halted through emulator debug. While halted, the control and status registers for the module can be viewed and manipulated for debug purposes.

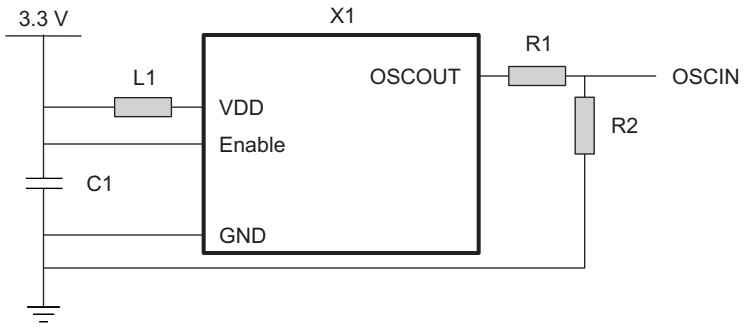
The peripheral emulation suspend functions on this device are disabled by default. In order to suspend a peripheral, the application must first enable the suspend function on an individual peripheral basis via the SYSCFG_SUSPSRC register.

2.1.5 System-Level ESD Immunity Usage Note

On all silicon revisions, certain design elements make this device susceptible to radiated noise during an ESD strike, as described in the standard IEC 61000-4-2. Exposure to the electrical noise caused by the ESD can cause soft device failures due to noise coupling on the system clock (OSCIN). ESD events within the IEC spec range do not cause permanent device damage and full functionality is recoverable with a device reset. The sensitivity to this noise issue is primarily due to the 1.2V oscillator/clock input implemented on this device. The low voltage range, coupled with slow rise and fall times, provides a lower noise margin than other TI devices with higher voltage internal oscillators (for example, 1.8V or 3.3V oscillators).

If ESD robustness is a concern, it is strongly recommended to avoid using the internal oscillator as a clock source. An external 3.3V clock source with a resistor voltage divider as in Figure 2 can be used to externally generate the required 1.2V input clock. By using an external clock input with fast rise/fall times (less than 5 ns), the noise margin improves significantly, increasing ESD noise resistance.

Copyright © 2010–2014, Texas Instruments Incorporated



A Legend: L1 = ferrite bead; C1 = 0.1 uF; R1 = 165 ohm / 5%; R2 = 100 ohm / 5%

Figure 2. External 3.3V Clock Source

In addition to using an external clock source, several other board and software recommendations specific to this device can improve system-level ESD immunity:

- The OSCIN and OSCVSS (and OSCOUT, if used) should be routed as short as possible to reduce their ability to pick up EMI noise.
- Route the OSCIN signal on inner board layers where it is shield by power and ground planes.
- Disable the DLL REFCLK signal in the DDR EMIF PHY. This prevents the DLL used by the DDR PHY from dynamically tracking glitches on the input clock. This can be done after normal DDR initialization by setting the following bit in the DDR PHY Control Register (0xB00000E4):

```
// Configure DDR PLL
   Set_DDRPLL_150MHz();
// Configure DDR timings
   DEVICE_DDR2Config(150);
// Minimum 600 MCLK cycle delay (allow master DLL to lock)
   Delay_600();
// Perform dummy DDR read
   volatile unsigned int k=0;
   ...
   k = *(volatile unsigned int*) (0xC0000000);
// Disable DLL REFCLK
   DRPYC1R |= 0x00002000;
```

- The processor should be provided as much power supply decoupling as is practical and placed as close to the processor as possible.
- Follow the entire DDR interface implementation requirements in the device datasheet.
- Implement the PLL filtering circuits shown in the device datasheet.

These recommendations are in addition to standard methods for increasing system ESD immunity, such as using shielding enclosures, proper grounding and PCB stackup, and ESD protection circuitry.

2.1.6 DDR2/mDDR Controller: mDDR Usage Note

On all silicon revisions, some mDDR memories designed with Status Register Read (SRR) support may be incompatible with the AM1806 device. These mDDR memories misinterpret certain initialization sequence commands sent by the DDR2/mDDR controller following a power-on-reset or a cold reset (by asserting the RESET pin) of the device. This results in the mDDR memory becoming unresponsive as it is stuck in the SRR state instead of returning to the IDLE state.

To ensure correct initialization of these types of mDDR memories, two consecutive mDDR initialization sequences have to be sent by the DDR2/mDDR controller to the mDDR memory after a power-on-reset or a cold reset.

If users rely on the ROM Bootloader (RBL) to perform DDR2/mDDR initializations, they can generate an Application Image Script (AIS) boot image with two consecutive DDR2/mDDR controller configuration sequences, with the help of the AIS GUI utility or the AIS command-line utility.

Copyright © 2010–2014, Texas Instruments Incorporated

Using the GUI utility (AISgen_d800k008.exe, version 1.13 or later):
Use the GUI tool to generate the AIS boot image with the required configuration and choose the 'mDDR with SRR support' option under the Memory Type section.

Using the command-line utility (HEXAIS_OMAP_L138.exe):
With the command-line utility, users should manually add a second set of PLL1 and DDR2/mDDR controller configuration sequence at the end of the first DDR2/mDDR controller configuration sequence in the .ini script.

Note: The *Using the OMAP-L132/L138 Bootloader* application report (literature number: SPRAB41) provides a link to the install package for the AIS tool which includes the following in the install directory: prebuilt patch files, the GUI AIS generation tool, command-line AIS generation tool, and example .ini files.

If the RBL is not used to initialize the DDR2/mDDR memory, two consecutive PLL1 and DDR2/mDDR controller configuration sequences should be incorporated into the user code that configures the DDR2/mDDR controller.

Boot images generated using the above methods can be used for all compatible mDDR memories. However, it is recommended to check with the memory vendor before doing so.

Note: Not all mDDR memories designed with Status Register Read (SRR) support exhibit the stuck-in-SRR-state behavior.

The set of two consecutive mDDR initialization sequence is only required during after a power-on-reset or a cold reset. It is not required after a soft reset (from Power and Sleep Controller or Watch-Dog Timer) of the DDR2/mDDR controller.

2.1.7 McASP: Inactive Slot Usage Note

On all silicon revisions, in any McASP serializer configured a transmit serializer with an n-slot TDM, data transfer can fail if both of the conditions below are true:

1. one or more time slots within the n-slot TDM are configured as inactive and
2. EDMA is used to transfer data to McASP

If the conditions mentioned above exists, either the transmit operation may fail to start with the XDATA bit in the Transmit Status Register (XSTAT) set or if the transmit operation has started, random underrun errors may occur breaking the data transfer operation.

To ensure correct McASP transmit operation with EDMA triggered data transfers, all time slots in an n-slot TDM should be configured as active slots. For example, if a serializer is configured for transmit operation with a 5-slot TDM frame in which it is only required to transmit data in slots 0 to 2, all five slots (0 to 4) should be configured as active in the Transmit TDM Time Slot Register (XTDM). In this example the remaining time slots (slot 5 onwards) can be configured as inactive. The EDMA configuration and user application should account for the transfer of extra data to the McASP for slots 3 and 4.

Copyright © 2010–2014, Texas Instruments Incorporated

2.2 Silicon Revision 2.3 Known Design Exceptions to Functional Specifications

The advisories may not always be enumerated in sequential order and hence some numbers may not appear in the document.

Table 2. Silicon Revision 2.3 Advisory List

Title	Page
Advisory 2.3.1 —USB0: Extraneous RESET Interrupt	9
Advisory 2.3.2 —EMIFA: Asynchronous Memory Timeout Error Persistence	10
Advisory 2.3.3 —Potential USB2.0 Soft Reset Timing Violation	11
Advisory 2.3.4 —Vil on Dual-Voltage LVCMOS Input Buffers Operated at 3.3V	12
Advisory 2.3.5 —ARM Interrupt Controller Vector Size Register (VSR) Initialization	15
Advisory 2.3.6 —A Single CHIPINTn Interrupt Event Can Register Multiple Times in the AINTC	16
Advisory 2.3.7 —DVDD18 Can Pull Up to 2.7V When Using Dual-Voltage IOs at 3.3V	18
Advisory 2.3.8 —USB 2.0 On-The-Go (OTG) Session Request Protocol (SRP) Is Not Supported	22
Advisory 2.3.12 —BOOT: Internal Pullup Resistors for BOOT[7:0] Pins Are Sometimes Enabled During Reset, Leading to Boot Failures	23
Advisory 2.3.13 —Boot: ECC Data Error in Spare Area Causes NAND Boot Failure	25
Advisory 2.3.25 —USB0: CPU gets Stale Receive Data from the Data Buffer located in External Memory	29
Advisory 2.3.26 —USB0: Early DMA Completion in DMA Receive Mode and More Than One Endpoint is Transferring Data	30
Advisory 2.3.27 —USB0: DMA Hung up in Frequent Teardowns	31

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.1 *USB0: Extraneous RESET Interrupt*

Revision(s) Affected 2.3 and earlier

Details When the USB controller is operating as a device and an attached host resets the device after the completion of the Device Attached state by driving both differential data lines low, the USB controller operating as a device could receive multiple RESET interrupts for the single RESET signaling invoked by the host. The multiple interrupt generation only happens for the duration of the RESET signaling on the bus. RESET Interrupt is not generated before or after the completion of RESET.

Workaround(s) Software must service every USB RESET interrupt received. Software should not proceed on performing any other task, like initialization, until RESET duration has come to completion. The POWER[RESET] bit field will be cleared by the USB Controller when RESET signaling on the bus is removed by the Host. The USB Controller clearing the POWER[RESET] bit field should be used by software as an indication for the completion of RESET signaling.

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.2 *EMIFA: Asynchronous Memory Timeout Error Persistence*

Revision(s) Affected

2.3 and earlier

Details

In Extended Wait mode, during a read access to an asynchronous memory, if the WAIT input does not go inactive within maximum extended wait cycles programmed in the Async Wait Cycle Config register, the EMIF will report a time-out error. The data returned for this access will be all zeros. If this access is followed by a read to the EMIFA's memory-mapped register (MMR) space, the EMIFA will still report a time-out error but with the correct data for the MMR read. The EMIF will hold the time-out error until another asynchronous access without a time-out error or an SDRAM access is performed.

This issue is only applicable if all of the following are true:

- The EMIF is used for asynchronous memory accesses in Extended Wait mode.
- There is a potential for a time-out error to occur, i.e., the asynchronous memory will not de-assert the WAIT input.
- If asynchronous memory read with time-out error is followed by an MMR read.

Workaround(s)

If a time-out occurs, perform any of the following:

- A dummy read to another asynchronous memory chip select that is not configured to be in Extended Wait mode.
- A dummy read to the same asynchronous memory chip select after disabling the Extended Wait mode on that chip select.
- A dummy read to SDRAM

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.3 *Potential USB2.0 Soft Reset Timing Violation*

Revision(s) Affected

2.3 and earlier

Details

When a soft reset is invoked by setting the RESET bit of the USB CTRLR register (CTRLR[RESET] = 1), the internal reset timing requirements may be violated. Although this timing violation has not been observed in practice, the potential for a timing violation exists.

USB resets initiated by system-reset and power-on-reset are immune from the timing violation.

There is no plan to fix this issue in future silicon revisions because:

1. No functional problems have been observed to date
2. A software workaround has been developed to avoid the problem

Workaround(s)

The reset timing violation can be avoided by providing the modified soft reset activation sequence outlined below:

1. Enable the USB controller module clock through the Power and Sleep Controller (PSC)
2. Perform a soft USB reset
3. Wait for the USB soft reset bit to clear
4. Disable the USB controller module clock through the PSC
5. Configure the USB PHY parameters
6. Enable the PHY
7. Enable the USB controller module clock through the PSC

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.4 *Vil on Dual-Voltage LVCMOS Input Buffers Operated at 3.3V*

Revision(s) Affected 2.3 and earlier

Details The input buffers on the device have shown timing sensitivity to the logic-low input voltage that can cause changes to the AC input timings. Due to this issue, input voltages must be driven to 0.5V or below on all dual-voltage LVCMOS input signals (signals associated with supplies DVDD1833_A, DVDD1833_B or DVDD1833_C).

Vil driven at or below 0.4V will cause no timing degradation. Vil driven up to 0.5V will cause up to 750 ps degradation in input timings.

The following datasheet parameters will be affected by Vil driven up to 0.5V. Their values adjusted for 0.75 ns degradation are shown.

Table 3. Timing Requirements for the EMIFA SDRAM Interface

NO.	PARAMETER		1.2V		UNIT
			MIN	MAX	
19	$t_{su(EMA_DV\text{-}EM_CLKH)}$	Input Setup time, read data valid on EMA_D[31:0]	2.75		ns

Table 4. Timing Requirements for the EMIFA Asynchronous Memory Interface

NO.	PARAMETER		1.2V		UNIT
			MIN	MAX	
12	$t_{su(EMDV\text{-}EMOEH)}$	Setup time, EMA_D[31:0] valid before $\overline{EMA_OE}$ high	3.75		ns

Table 5. Timing Requirements for McASP0

NO.	PARAMETER			1.2V		UNIT
				MIN	MAX	
6	$t_{h(ACLKRX\text{-}AFSRX)}$	Hold time, AFSR/X input after ACLKR/X (1)	AHCLKR/X ext input	1.15		ns
			AHCLKR/X ext output	1.15		
8	$t_{h(ACLKRX\text{-}AXR)}$	Hold time, AXR0[n] input after ACLKR/X (1) (2)	AHCLKR/X ext input	1.15		
			AHCLKR/X ext output	1.15		

(1) McASP0 ACLKXCTL.ASYNC=1: Receiver is clocked by its own ACLKR0
(2) McASP0 ACLKXCTL.ASYNC=0: Receiver is clocked by transmitter's ACLKX0

Table 6. Switching Characteristics for McASP0

NO.	PARAMETER			1.2V		UNIT
				MIN	MAX	
13	$t_{d(ACLKRX\text{-}AFSRX)}$	Delay time, ACLKR/X transmit edge to AFSX/R output valid (1)	ACLKR/X ext input	3.5		ns
			ACLKR/X ext output	3.5		
14	$t_{d(ACLKX\text{-}AXRV)}$	Delay time, ACLKX transmit edge to AXR output valid (1)	ACLKR/X ext input	3.5		
			ACLKR/X ext output	3.5		
15	$t_{dis(ACLKX\text{-}AXRHZ)}$	Disable time, ACLKR/X transmit edge to AXR high impedance following last data bit	ACLKR/X ext	3.5		

(1) McASP0 ACLKXCTL.ASYNC=1: Receiver is clocked by its own ACLKR0

Copyright © 2010–2014, Texas Instruments Incorporated

Table 7. Timing Requirements for McBSP0

NO.	PARAMETER			1.2V		UNIT
				MIN	MAX	
5	$t_{su(FRH-CKRL)}$	Setup time, external FSR high before CLKR low	CLKR ext	5.5		ns
7	$t_{su(DRV-CKRL)}$	Setup time, DR valid before CLKR low	CLKR ext	5.5		
10	$t_{su(FXH-CKXL)}$	Setup time, external FSX high before CLKX low	CLKR ext	5.5		

Table 8. Switching Characteristics for McBSP0

NO.	PARAMETER			1.2V		UNIT
				MIN	MAX	
4	$t_{d(CKRH-FRV)}$	Delay time, CLKR high to internal FSR valid	CLKR ext	2.75		ns
9	$t_{d(CKXH-FXV)}$	Delay time, CLKX high to internal FSX valid	CLKR ext	2.75		
13	$t_{d(CKXH-DXV)}$	Delay time, CLKX high to DX valid	CLKR ext	2.75 + D1 (1)		

(1) Extra delay from CLKX high to DX valid applies only to the first data bit of a device, if and only if DXENA = 1 in SPCR.
If DXENA = 0, then D1 = D2 = 0
If DXENA = 1, then D1 = 6P, D2 = 12P

Table 9. Switching Characteristics for McBSP1

NO.	PARAMETER			1.2V		UNIT
				MIN	MAX	
4	$t_{d(CKRH-FRV)}$	Delay time, CLKR high to internal FSR valid	CLKR ext	3.25		ns
9	$t_{d(CKXH-FXV)}$	Delay time, CLKX high to internal FSX valid	CLKR ext	3.25		
13	$t_{d(CKXH-DXV)}$	Delay time, CLKX high to DX valid	CLKR ext	3.25 + D1 (1)		

(1) Extra delay from CLKX high to DX valid applies only to the first data bit of a device, if and only if DXENA = 1 in SPCR.
If DXENA = 0, then D1 = D2 = 0
If DXENA = 1, then D1 = 6P, D2 = 12P

Table 10. Timing Requirements for Universal Parallel Port (uPP)

NO.	PARAMETER		1.2V		UNIT
			MIN	MAX	
4	$t_{su(STV-INCLKH)}$	Setup time, CHn_START valid before CHn_CLK high	4.75		ns
6	$t_{su(ENV-INCLKH)}$	Setup time, CHn_ENABLE valid before CHn_CLK high	4.75		
8	$t_{su(DV-INCLKH)}$	Setup time, CHn_DATA/XDATA valid before CHn_CLK high	4.75		
10	$t_{su(DV-INCLKL)}$	Setup time, CHn_DATA/XDATA valid before CHn_CLK low	4.75		
19	$t_{su(WTV-INCLKL)}$	Setup time, CHn_WAIT valid before CHn_CLK high	4.75		

Copyright © 2010–2014, Texas Instruments Incorporated

Table 11. Timing Requirements for Video Port Interface (VPIF)

NO.	PARAMETER		1.2V		UNIT
			MIN	MAX	
1	$t_{su(VDINV-VKIH)}$	Setup time, VP_DINx valid before VP_CLKIN0/1 high	4.75		ns

Workaround(s)

Although there is no specific workaround, the following recommendations can be used to help prevent this issue:

- Minimize loads as much as possible, especially DC loads that could cause the Vil to rise. **Point-to-point (single-load) connections are unlikely to be affected.**
- Falling edges should transition as rapidly as possible (so the signal passes through the 0.2V point as early as possible). Heavily loaded nodes resulting in degraded fall times may require drivers to provide rapid input edges.

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.5 *ARM Interrupt Controller Vector Size Register (VSR) Initialization*

Revision(s) Affected 2.3 and earlier

Details The VSR register in the ARM Interrupt Controller (AINTC) is not correctly initialized after reset. If this register is not explicitly configured, the AINTC will only allocate 1 byte per interrupt (instead of 4).

Workaround(s) The desired value (even if it is the default value) should be written to the VSR prior to using the interrupt controller.

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.6 *A Single CHIPINTn Interrupt Event Can Register Multiple Times in the AINTC*

Revision(s) Affected 2.3 and earlier

Details

Interrupts destined for the ARM CPU are managed by the ARM Interrupt Controller (AINTC). The AINTC detects, combines, and routes system interrupts to the two native ARM interrupt signals FIQ and IRQ. See the device System Reference Guide for additional information about the AINTC.

The AINTC module expects all incoming interrupts to be pulse interrupts, however the [SYSCFG_CHIPSIG_]CHIPINTn interrupts are level interrupts. This mismatch in interrupt types will cause a single CHIPINTn interrupt event to register as multiple interrupt pulses in the AINTC. However, the AINTC does not have the capacity to count the number of interrupt pulses received per system interrupt – it only maintains interrupt flags. A system interrupt is flagged as active until its status is cleared by the user through the AINTC, regardless of the number of interrupts detected.

If the status flag for AINTC CHIPINTn is cleared while the CHIPINTn interrupt is still active, the AINTC will continue to detect CHIPINTn interrupts and its status flag will be set again. This additional setting of the AINTC CHIPINTn status flag is false.

Workaround(s)

Method 1

Do not execute the intended interrupt service routine code if the associated CHIPSIGn status flag is not set in the SYSCFG_CHIPSIG register. A cleared CHIPSIGn status flag indicates that the device is responding to a false interrupt. This method is easy to implement, but does not eliminate false interrupts.

```
/** Pseudo code only **/

   void CHIPINT0_ISR(void) {
       /* Exit immediately if CHIPSIG0 is not set */
       if( (SYSCFG->CHIPSIG & 0x1) == 0 ) {
           return;
       }

       /* Intended service routine code */
       SYSCFG->CHIPSIG_CLR = 0x1;
       printf("Hello World!\n");
   }
```


Copyright © 2010–2014, Texas Instruments Incorporated

Method 2

Do not clear the AINTC CHIPINTn status flag until the CHIPSIGn status has been cleared. This method will eliminate false interrupts, but requires changes to the AINTC interrupt dispatch code. Changing the dispatch code may introduce undesired behavior in the application.

```
/** Pseudo code only **/

   /* Sequence that is susceptible to false CHIPINTn interrupts */
   void AINTC_ISR_DISPATCH_1(void) {
       Get_Interrupt_Information();

       /* CHIPINTn interrupts continue to be generated after */
       /* AINTC CHIPINTn flag is cleared.                    */
       Clear_AINTC_Interrupt_Flag();

       /* CHIPINTn interrupts are only stopped after ISR clears */
       /* the status flag.                                      */
       Branch_To_ISR();
   }

   /* Sequence that is not susceptible to false CHIPINTn interrupts */
   void AINTC_ISR_DISPATCH_2(void) {
       Get_Interrupt_Information();

       /* ISR will clear CHIPSIGn flag and discontinue CHIPINTn */
       /* interrupts to AINTC.                                  */
       Branch_To_ISR();

       /* Ok to clear AINTC CHIPINTn flag now.                  */
       Clear_AINTC_Interrupt_Flag();
   }
```


Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.7 *DVDD18 Can Pull Up to 2.7V When Using Dual-Voltage IOs at 3.3V*

Revision(s) Affected 2.3 and earlier

Details

A condition can occur during the device power supply ramp in which the DVDD18 supply, which should be at 1.8V nominal, can be pulled up on-chip by the DVDD3318_A, DVDD3318_B or DVDD3318_C supplies operated at 3.3V. For the sake of this description, DVDD3318_x means any combination of DVDD3318_A, DVDD3318_B or DVDD3318_C.

Important Note: The 3.3V DVDD3318_x supplies must not be driven to 0V during zones A – E or high current capable of damaging the device may occur. The 3.3V DVDD3318_x supplies should not be driven during this time.

A normal example power-supply ramp is shown below:

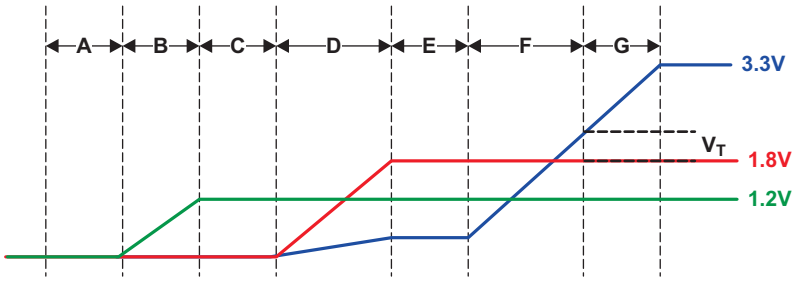


Figure 3. Normal power-supply ramp

- A. All device supplies are undriven
- B. The 1.2V supplies are ramped to their nominal levels
- C. Potential delay between supply ramps (not required)
- D. The 1.8V supplies (specifically DVDD18) are ramped to their nominal levels. During the DVDD18 ramp there will be some minor drift up in the undriven DVDD3318_x supplies. This is normal and is not a problem.
- E. Potential delay between supply ramps (not required)
- F. DVDD3318_x is ramping but is less than DVDD18 + V_T
- G. DVDD3318_x is ramping but is greater than DVDD18 + V_T

When the fault condition occurs, the supplies behave as shown below:

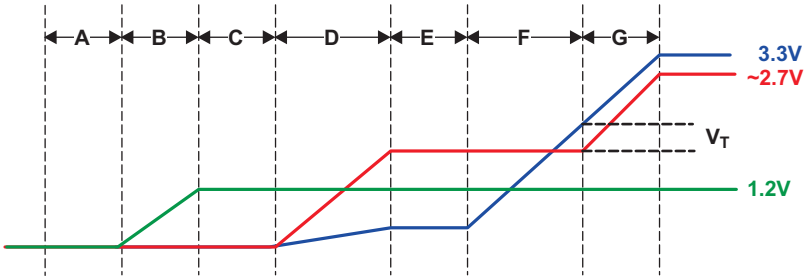
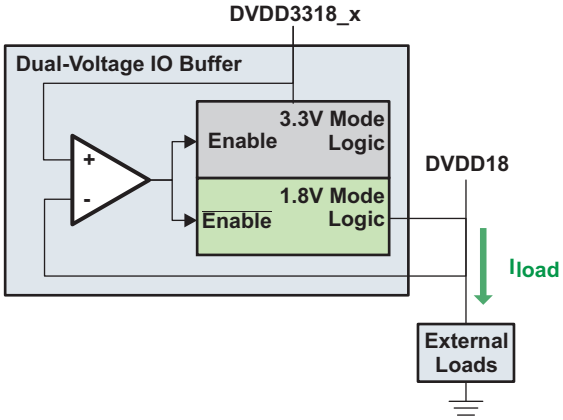


Figure 4. Faulty power-supply ramp

The behavior is the same until zone G. When the DVDD3318_x supply exceeds DVDD18 by a transistor V_T, the DVDD18 supply begins to be pulled up by the DVDD3318_x supply and follows the DVDD3318_x supply by DVDD3318_x- V_T. Since the DVDD18 supply is often connected to other 1.8V supplies in the design, these other supplies can be pulled up also.

Copyright © 2010–2014, Texas Instruments Incorporated

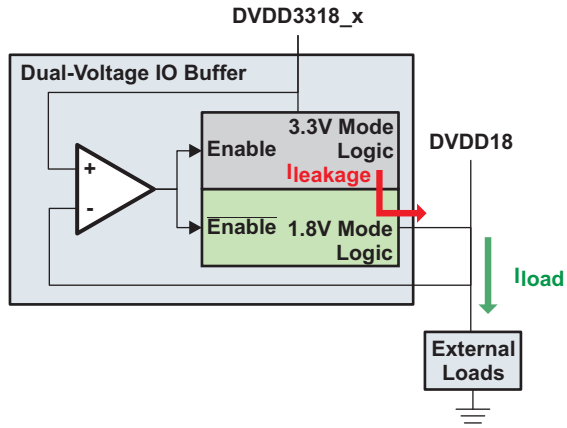
The condition occurs because the dual-voltage IO buffers have voltage detection circuitry that monitors DVDD3318_x during reset and determines whether the applied supply voltage is 1.8V or 3.3V. This detection circuitry then configures the IOs to operate at the appropriate voltage. As the DVDD3318_x supply ramps, there is a small range near DVDD18 + V_T where the voltage detection circuit finds the state indeterminate and briefly turns on circuitry associated with both voltage options creating a current path between them. This current path can cause the DVDD18 supply to be pulled up. The mechanism for this behavior is explained below:



Behavior of the IO buffer during Zone F:

As the DVDD3318_x supply ramps, but is not yet one V_T above the DVDD18 supply, the comparator enables the 1.8V mode logic section and disables the 3.3V mode logic section. No unexpected current flows between the two sections.

Figure 5. Behavior of the IO buffer during Zone F

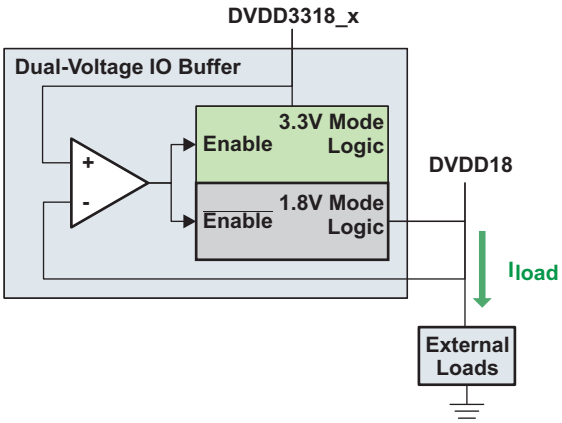


Error State of the IO buffer during Zone G:

As DVDD3318_x exceeds DVDD18+V_T, some of the circuitry across the boundary of the two voltage domains will turn on and causes a leakage current (in red) to flow between the two voltage domains. This current pulls up the DVDD18 supply unless the power source providing that supply can oppose it or unless the load current is strong enough to counteract it.

If DVDD18 continues to rise as DVDD3318_x ramps, the comparator never sees enough difference between the two supplies to switch to 3.3V mode.

Figure 6. Error State of the IO buffer during Zone G



Recovered State of the IO buffer during Zone G:

If the load current is high enough or the DVDD18 power supply can oppose the leakage current, then the DVDD18 voltage stays low enough for:

- **the comparator to recognize 3.3V mode**
- **the 1.8V mode logic is turned off**
- **the leakage stops**

Figure 7. Recovered State of the IO buffer during Zone G

Copyright © 2010–2014, Texas Instruments Incorporated

This error condition may occur when either of the following conditions are present:

- The regulator used to control the DVDD18 supply can only regulate voltage up and relies on the load to pull the output voltage down
- There is low load on the DVDD18 supply during the DVDD3318_x power supply ramp

This condition will not occur in designs where:

- All DVDD3318_x supplies are operated at 1.8V, or
- The regulator used to supply DVDD18 has the capability to actively regulate (drive) its output voltage up or down and doesn't rely on system load to pull the voltage down, or
- The load on the DVDD18 supply during the DVDD3318_x power supply ramp is sufficient to oppose the leakage current.

Workaround(s)

1. Design simulation has indicated that under the worst-case process/voltage/temperature conditions the maximum leakage current into the DVDD18 supply due to this phenomenon will be 140 mA (the 140 mA is a cumulative current generated by all of the dual-voltage IOs). So this error state can be prevented by ensuring that the load on the DVDD18 supply during the DVDD13318_x supply ramp exceeds 140 mA. This can be achieved by any combination of the following as long as they cumulatively produce >140 mA load on DVDD18 during the DVDD3318_x ramp period:

   (a) Maintain sufficient bulk capacitance on the DVDD18 supply such that the charging current for these capacitors provides all or part of the required >140 mA. Bulk capacitance in this context means the total capacitance seen by the DVDD18 supply (filter capacitors, bypass capacitors, etc.). Capacitor charging current is defined as I = C*(dV/dt). So the ramp rate of the DVDD3318_x supply and the total bulk capacitance on the DVDD18 supply can be used to calculate the current produced. This solution provides additional power supply filtering and little current leakage after the supplies are ramped (depending on the type of capacitors used).

   The table below shows some examples of the bulk capacitance that would be required to use this solution alone:

Table 12. Bulk Capacitance

DVDD3318_x ramp time for 3.3V	dV/dt (in V per second)	Required capacitance to generate 140 mA
100 µs	33000	4.2 µF
250 µs	13200	10.6 µF
500 µs	6600	21.2 µF
1 ms	3300	42.4 µF
10 ms	330	424.2 µF

   (b) Use an additional shunt regulator to control the voltage at DVDD18. The shunt regulator is placed between DVDD18 and Vss. As an example, the TLVH431 can provide up to 70 mA of additional load current to help maintain the DVDD18 voltage. When the voltage drops back to the normal 1.8V range, the current flow through the shunt regulator drops into the <100 uA range. This solution requires the shunt regulator and two additional resistors to set the desired regulation voltage.

   (c) Use a resistor to provide additional load between DVDD18 and Vss. This solution is less desirable since it continues to draw power even after the supply ramp is completed but would likely be the lowest cost. An improvement is to use a FET/switch in series with the resistor between DVDD18 and Vss that can later be turned off (by the RESET signal or a GPIO signal for example).

Copyright © 2010–2014, Texas Instruments Incorporated

2. Choose a regulator for the DVDD18 supply (or grouped 1.8V supplies) that is capable of actively regulating voltage up and down. Many high-efficiency switching regulators switch current into the load only when the voltage needs to be raised and rely on the load to pulldown the current. In this error condition, a regulator of this type will not be able to compensate for the leakage current described above. Synchronous buck regulators use external inductance to pulldown the regulated voltage when necessary.

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.8 *USB 2.0 On-The-Go (OTG) Session Request Protocol (SRP) Is Not Supported*

Revision(s) Affected 2.3 and earlier

Details The USB 2.0 On-The-Go (OTG) Session Request Protocol (SRP) allows a USB-peripheral to request the USB-host to enable Vbus and start a session. On this device, the SRP protocol is not supported.

The OTG Host Negotiation Protocol (HNP), which allows USB-devices to swap roles between host and peripheral, is supported.

Workaround(s) None

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.12 *BOOT: Internal Pullup Resistors for BOOT[7:0] Pins Are Sometimes Enabled During Reset, Leading to Boot Failures*

Revision(s) Affected 2.3 and earlier

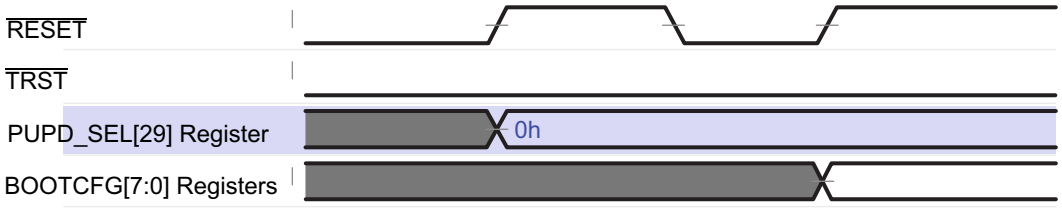
Details:

The PUPD_SEL[29] register does not get initialized when the device is first powered on and in Reset. This register controls the internal pullup and pulldown resistors for the BOOT[7:0] pins. The contents of the PUPD_SEL[29] register at this state are unpredictable and may contain random values. These random values can result in the internal pullups being enabled for some or all of the BOOT[7:0] pins during reset after every power-on, which conflicts with the datasheet claim that the internal pulldowns are enabled during reset.

Internal pullups being enabled on the BOOT[7:0] pins may result in boot failures. If weak external pulldown resistors are used on the PCB to select the boot mode, they may not be able to overpower the internal pullups. This can result in the wrong boot mode being latched in the BOOTCFG[7:0] register when coming out of Reset.

Once the device is out of Reset, the PUPD_SEL[29] register is initialized, and the internal pulldown resistors for the BOOT[7:0] pins are all enabled. Issuing a second Power-On Reset ($\overline{\text{POR}}$) results in the intended boot mode being latched, since the pins are internally pulled down as expected.

Initialization of PUPD_SEL[29] and BOOTCFG[7:0] Registers with Weak or No External Pull Resistors on BOOT[7:0] Pins shows the behavior of the PUPD_SEL[29] register before and after reset. Assuming weak or no external pull resistors are on the BOOT[7:0] pins, the BOOTCFG register will latch the unknown values after the device initially comes out of Reset. Note that, $\overline{\text{TRST}}$ ***must*** always be low in order to issue a $\overline{\text{POR}}$ and latch the boot pin values -- the boot pins are not latched after a Warm Reset.



Initialization of PUPD_SEL[29] and BOOTCFG[7:0] Registers with Weak or No External Pull Resistors on BOOT[7:0] Pins

Other device pins with configurable internal pullup or pulldown resistors are always internally pulled down during reset and are not affected by this advisory.

Workaround(s):

One of the following Workarounds must be implemented to ensure that the boot pins are always latched correctly:

Method 1

Use strong external pull resistors on BOOT[7:0] pins. Since either the internal pullup or pulldown resistors could be enabled after every power-on, the external resistors must be strong enough to oppose the internal pulls in either case.

Section 4, *Device Operating Conditions*, of the device datasheet shows the electrical characteristics information which can be used to calculate the maximum external pull resistance required. The value is dependent on the DVDD3318_C I/O supply level.

For BOOT pins which need to be sampled as logical low, the external pulldown resistance (R_{PDmax}) must be selected by assuming the internal pullup is enabled. The calculation is shown in Table 13.

Copyright © 2010–2014, Texas Instruments Incorporated

Table 13. Required Pulldown Resistance (R_{PDmax}) for Logical Low BOOT Pins

DVDD3318_C	I_I	V_{IL}	$R_{PDmax}\left(\frac{V_{IL}}{I_I}\right)$
3.3 V	310 µA	0.80 V	2.58 kΩ
1.8 V	310 µA	0.80 V	2.03 kΩ

For BOOT pins which need to be sampled as logical high, the external pullup resistance (R_{PUmax}) must be selected by assuming the internal pulldown is enabled. The calculation is shown in Table 2.

Table 14. Required Pullup Resistance (R_{PUmax}) for Logical High BOOT Pins

DVDD3318_C	I_I	V_{IH}	$R_{PUmax}\left(\frac{DVDD3318_C - V_{IH}}{I_I}\right)$
3.3 V	270 µA	2.00 V	4.81 kΩ
1.8 V	270 µA	1.17 V	2.33 kΩ

Method 2

For applications that have already implemented the Secondary Reset Workaround described in Advisory 2.0.9, no additional modifications are required as a fix for this Advisory.

As shown in Initialization of PUPD_SEL[29] and BOOTCFG[7:0] Registers with Weak or No External Pull Resistors on BOOT[7:0] Pins, issuing a second $\overline{\text{POR}}$ will always latch the BOOT pins while the internal pulldown resistors enabled.

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.3.13 *Boot: ECC Data Error in Spare Area Causes NAND Boot Failure*

Revision(s) Affected 2.3 and 2.1 (ROM Versions D800K008)

Details:

The ROM bootloader (RBL) reads a NAND page in segments of 512 bytes (user data) over the External Memory Interface A (EMIFA). The EMIF also reads the associated ECC data which is stored in the spare area of the flash (as shown in D800K008 ROM Behavior Before Application of Software Patch(1)). The ECC Correct function in the RBL can correct up to 4 bit errors in the user data and/or ECC data by using the syndrome generated from the ECC data and the parity of the user data calculated by the EMIFA module.

However, over the life span of the NAND flash, ECC data stored in the spare area can develop errors due to bit flips. When the calculated syndrome indicates an error in the ECC data, the ECC Correct and Read functions of the RBL abort the read process even though, it is possible to correct up to 4 bit errors combined in user and ECC data. Consequently the device fails to boot.

Explanation of Current ROM Bootloader Behavior:

The RBL implements a search mechanism to look for the boot image in NAND flash by using an Open function and a Read function. The Open function includes a bad block check where the RBL skips to the next block (shown as (A) in D800K008 ROM Behavior Before Application of Software Patch(1)) if the block under consideration is marked as "bad" (in the spare area). On finding a good block, the RBL attempts to read page0 (the first page) in that block.

If page0 of the good block has an ECC data error or an uncorrectable error (more than 4 bit errors combined), the RBL skips to the next block (shown as (B) in D800K008 ROM Behavior Before Application of Software Patch(1)). This bad block check mechanism enables the device to check up to the first 32 blocks in the NAND flash for booting (the boot process will fail if all 32 blocks have uncorrectable or ECC data errors in page0). Note that the RBL does not abort on detecting an error in ECC data on page0.

Once the RBL finds a good block with a good page0, it continues to read subsequent pages in that block. If an uncorrectable or ECC error is detected in subsequent pages, the RBL will abort with a boot error (shown as (C) in D800K008 ROM Behavior Before Application of Software Patch(1)). The workaround, described below, enables the boot process to continue for both types of errors (ECC data and uncorrectable errors).

Copyright © 2010–2014, Texas Instruments Incorporated



(1) All pages marked in **green** are good and all pages marked in **red** have "uncorrectable" errors.

D800K008 ROM Behavior Before Application of Software Patch(1)

Workaround(s):

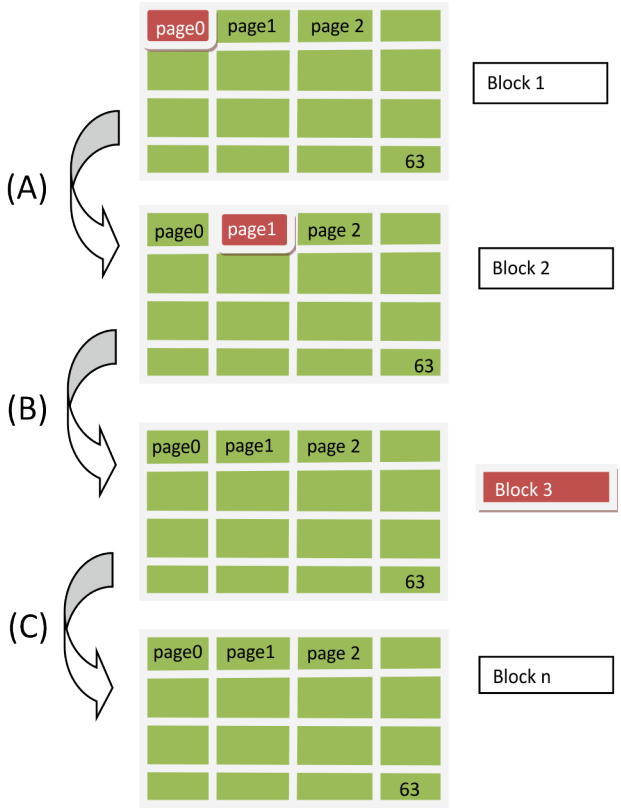
The workaround ignores errors in ECC data (these errors do not need to be corrected) so that the boot process can continue, correcting errors in the user data (up to 4 bit errors total), as necessary. The workaround also includes a mechanism to restart the NAND read process from the next good block if an uncorrectable error (more than 4 bit errors) is detected on any page of a good block.

This workaround is implemented using a software patch that is loaded in device RAM at boot time and is designed to change the default behavior of the ECC Correct and Read functions in the RBL. The patch binary replaces function pointers to the ECC Correct and Read functions in the ROM function table, (defined in device internal memory), during boot. The patch binary will reside in page0 of the NAND block and so will be applied only after the page0 of a good block has been read.

Once the patch is read, all subsequent page reads in the block will use the new ECC Correct and Read functions. ECC data errors on any page (other than page0) are ignored, any user data error (up to 4 bit errors total) is corrected and the boot process continues. If an uncorrectable error is detected on any page, the block is skipped and the boot process is restarted on the next good block. The behavior of the device boot from the NAND after application of the patch is shown in D800K008 ROM Behavior With Software Patch Applied(1).

Copyright © 2010–2014, Texas Instruments Incorporated

D800K008 ROM Behavior Before Application of Software Patch(1) (continued)



(1) All pages marked in **green** are good and all pages marked in **red** have "uncorrectable" errors.

D800K008 ROM Behavior With Software Patch Applied(1)

The software patch is available as a pre-built file with the latest version of the AIS tool that is used to generate the NAND flash boot image. The *Using the TMS320C6748/C6746/C6742 Bootloader Application Report* (Literature number: SPRAAT2) provides a link to the install package for the AIS tool which includes the following in the install directory: prebuilt patch files, the GUI AIS generation tool (AISGEN.exe, version 1.11 or later), command-line AIS generation tool and an example INI file.

Application of the Software Patch to NAND Boot Images:

- For the GUI tool, AISGEN.exe version 1.11 or later (found in the install directory), the patch integrates the modified ECC Correct function into the user application file to generate one binary AIS file.
- For the Command line AIS tool, HEXAIS_OMAP_L138.exe users (found in the install directory), the patch can be integrated into the user application file to generate one binary AIS file by inserting lines below to the end of the INI configuration file.

The patch name is: **ARM_nand_ecc_patch_OMAP-L138.out**

Copyright © 2010–2014, Texas Instruments Incorporated

D800K008 ROM Behavior With Software Patch Applied(1) (continued)

```
[INPUTFILE] ; get the NAND ECC patch file
FILENAME=Patch_name.out

; patch the NAND ECC handling routine
[AIS_Jump]
LOCATION=_NAND_ECC_patchApply
```

- The patch will be burnt on page0, as it is in the head of AIS file. Once page0 has been read successfully, the AIS set command will overwrite the function pointer in the RBL with the modified function pointer and the modified function will be applied to the later pages read. Memory usage of the patch at boot time is as follows:
  - 500 bytes at location 0xFFFF 0B00

***Recommendations to Improve Robustness*:**

- Page0 should be stored on multiple blocks as backup to take advantage of the safety mechanism built into Silicon Revision 2.1 to cycle to the next block when a page0 read fails in a good block.
- Maintain backup boot images in multiple blocks. The patch to the Abort function reinitializes the boot process and forces the boot to cycle to the next good block and restart the boot process by reading from it.
- Silicon Revision 2.1 supports booting from Block 0 of the flash which many NAND manufacturers guarantee as a “more reliable” block than all other blocks. Hence, setting up boot to start from that Block 0 could help improve the reliability of boot. This is a hardware change, requiring the bootmode pins BOOT[6:5] = 1x.

  For more details on the bootmode pins, see the *Using the AM18xx Bootloader Application Report* (Literature number: SPRABA5), *NAND-Boot Modes*.

Copyright © 2010–2014, Texas Instruments Incorporated

D800K008 ROM Behavior With Software Patch Applied[1] (continued)

Advisory 2.3.25 *USB0: CPU gets Stale Receive Data from the Data Buffer located in External Memory*

Revision(s) Affected 2.3 and earlier

Details When CPPI DMA completes a receive data transaction it posts a write to the Rx data buffer located in external memory, posts a write to update the descriptor located in external memory, and raises an interrupt to CPU. When the system load is high, the posted writes to DDR may not be complete before the CPU receives the interrupt. In this case, the CPU would fetch stale receive data from the Rx data buffer located in external memory.

Workaround(s) Initialize the datalength descriptor field to zero. CPPI DMA updates this field after the completion of an RX DMA operation with the actual number of bytes received. In the ISR (actually in a deferred call context), poll this field until it becomes a non-zero value to ensure data buffer has been updated with actual data. The descriptor buffer write is posted after the data buffer write, so waiting for the descriptor field to be updated ensures the data buffer has been updated. Since this workaround involves deferred procedure calls (whose schedule can be delayed depending on OS load), the latency sensitive application (like ISO Audio) might be affected by delay in notification to the application.

Copyright © 2010–2014, Texas Instruments Incorporated

D800K008 ROM Behavior With Software Patch Applied(1) (continued)

Advisory 2.3.26 *USB0: Early DMA Completion in DMA Receive Mode and More Than One Endpoint is Transferring Data*

Revision(s) Affected 2.3 and earlier

Details The erroneous short packet status can be detected on current endpoint and XDMA closes the Rx transfer in current endpoint. When more than one endpoint have been processed, if one of the endpoints has a short packet, then the short packet status is broadcasting to all endpoints.

This results in premature completion of a Rx descriptor in generic RNDIS CPPI DMA mode.

Workaround(s) The workaround involves monitoring transfer data size before and after transferring and reconfiguring data transfer size by software if the before and after size is different. Software must keep tracking every endpoint data transferring size. When DMA completion interrupt is received, software checks size difference. If the size is not equal, software requests the remaining data.

Copyright © 2010–2014, Texas Instruments Incorporated

D800K008 ROM Behavior With Software Patch Applied(1) (continued)

Advisory 2.3.27 *USB0: DMA Hung up in Frequent Teardowns*

Revision(s) Affected 2.3 and earlier

Details Teardown receive DMA is not working perfectly. This happens when a teardown is initiated by software during the endpoint is still active. Frequent teardown results in XDMA hung up situation.

Workaround(s) Software should make sure that DMA does not get to an unknown state during teardown by disabling the DMAEN bit in the RXCSR register. After this the teardown procedure can be initiated. Software should also add 250 ms delay during teardown.

Copyright © 2010–2014, Texas Instruments Incorporated

3 Silicon Revision 2.1 Usage Notes and Known Design Exceptions to Functional Specifications

This section describes the usage notes and advisories that apply to silicon revision 2.1 of the device.

3.1 Usage Notes for Silicon Revision 2.1

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 usage 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.

The usage notes detailed in Section 2.1 also apply for revision 2.1.

3.2 Silicon Revision 2.1 Known Design Exceptions to Functional Specifications

Silicon revision 2.1 applicable advisories have been found on a later silicon revision. For more details, see Section 2, *Silicon Revision 2.3 Known Design Exceptions to Functional Specifications*.

Table 15. Silicon Revision 2.1 Advisory List

Title	Page
Advisory 2.1.10 —USB0 PLL Mean Frequency Can Drift Across Large Temperature Swings	33

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.1.10 *USB0 PLL Mean Frequency Can Drift Across Large Temperature Swings*

Revision(s) Affected 2.1 and earlier

Details Under conditions in which the device is subjected to large variations in operating temperatures, the USB0 PLL temperature compensation circuitry does not have enough margin to guarantee compensation for PLL drift across all temperature ranges.

As a result, the mean frequency generated by the USB0 (USB 2.0 OTG) PHY PLL will begin to drift (relative to the expected 480 Mbps) when the temperature of the device is subjected to large swing from the original temperature in which the USB0 PHY was most recently calibrated (initialized).

Once the onset of PLL drift occurs, the mean frequency will continue to drift outside the expected frequency and will eventually cause the PLL to lose lock resulting in failure of USB packet reception and/or transmission. This break in transmission will continue until the USB0 PHY is recalibrated during a USB0 PHY Reset.

If the device is not exposed to large variations in temperature relative to the temperature at which the USB0 PHY was most recently initialized, the temperature compensation circuitry is expected to provide the proper compensation to prevent the mean PLL frequency from losing lock and beginning to drift.

More specifically, this advisory is most applicable in applications where the device is expected to operate outside the commercial temperature space (0˚C-90˚C). TI has identified a point-to-point device temperature range of 0˚C-65˚C in which there is very high confidence in which the compensation circuitry will properly compensate for all variations in temperature provided that the USB0 PHY was most recently initialized (calibrated) within this same temperature range.

Operating outside the 0˚C-65˚C temperature range increases the susceptibility of the device to experience PLL drift, but does not mean that the application will always experience a failure in USB transmission.

Root Cause The Voltage Controlled Oscillator (VCO) Compensation circuitry local to the USB0 PHY was not designed with a large enough range to compensate for all variations in temperature across the specified operating range of the device.

How to Most Easily Reproduce the Issue: Reproduction of this issue can most easily be accomplished by the following steps:

1. Allowing the unit to soak in an ambient temperature of -35˚C until the device temperature reaches approximately the same temperature.
2. Power up the device and provide the necessarily software programming in order to invoke the USB Signal Quality Test Pattern.
3. Using a USB 2.0 Certified Test Platform, execute the USB signal quality test procedure across the following temperature set points. -35˚C, 0˚C, +35˚C, +70˚C. Record the measured mean frequency by the compliance software.

NOTE: The set points can be varied to obtain finer temperature resolution of when the PLL begins to drift a per platform basis. The above temperature profile is provided for reference.

Copyright © 2010–2014, Texas Instruments Incorporated

Workaround(s)

When a break in transmission is detected, USB0 traffic can be recovered by a software reset of the USB0 PHY. A PHY reset implies recalibration of the PHY PLL at the reset temperature. The system has not been observed to reliably recover on its own. A PHY reset also implies re-enumeration of all devices. There is no way to recalibrate the USB0 PHY without a re-enumeration.

In order to invoke the recovery mechanism (i.e. a USB0 PHY reset) one needs to determine when the issue is present. One such approach is to look for an absence of USB0 Core interrupts over a specified time window. This window should be optimized for the expected USB traffic based upon the application.

As an additional safeguard, an application can also intentionally schedule pre-determined USB PHY resets at specific temperature points if operation over a broad range is expected.

Here is an example of one way to power cycle the USB0 PHY via the Chip Configuration 2 Register in the System Configuration (SYSCFG) Module:

```
#define CFGCHIP2  *((volatile unsigned int *) 0x01C14184)
    #define USBPHY_PHYPDWN 0x00000200

Void phy_reset(void) {
        CFGCHIP2 |= USBPHY_PHYPDWN;        /* Power down the USB PHY */

        mdelay(1);                         /* Wait 500ms */

        CFGCHIP2 &= ~USBPHY_PHYPDWN;       /* Power up the USB PHY */
}
```


Copyright © 2010–2014, Texas Instruments Incorporated

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

This section describes the usage notes and advisories that apply to silicon revision 2.0 of the device.

4.1 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 usage 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.

The usage notes detailed in Section 2.1 also apply for revision 2.0.

4.2 Silicon Revision 2.0 Known Design Exceptions to Functional Specifications

Table 16. Silicon Revision 2.0 Advisory List

Title	Page
Advisory 2.0.9 — Intermittent Boot Failures	36

Copyright © 2010–2014, Texas Instruments Incorporated

Advisory 2.0.9 *Intermittent Boot Failures*

Revision(s) Affected 2.0 and earlier

Details

For affected silicon revisions, a proprietary Boot Controller (which is not accessible by the user) initiates the system boot sequence when the device is released from reset. Before the ARM can take control of the user boot mode, the Boot Controller must first initialize the ARM reset vector table so that the ARM CPU will execute from its boot ROM.

The ARM reset vector table is located in the ARM's local RAM, however the ARM local RAM can only be accessed by two bus masters: ARM and PRU0. Therefore, the Boot Controller must program PRU0 to copy the desired reset vector table into the ARM's local RAM.

The PRU instructions are located inside of an instruction RAM (IRAM) which is initialized by the Boot Controller (see Figure 8). After the instructions are stored to IRAM, the PRU is reset and enabled to execute its instructions. In this case, the PRU is instructed to initialize the ARM reset vector table.

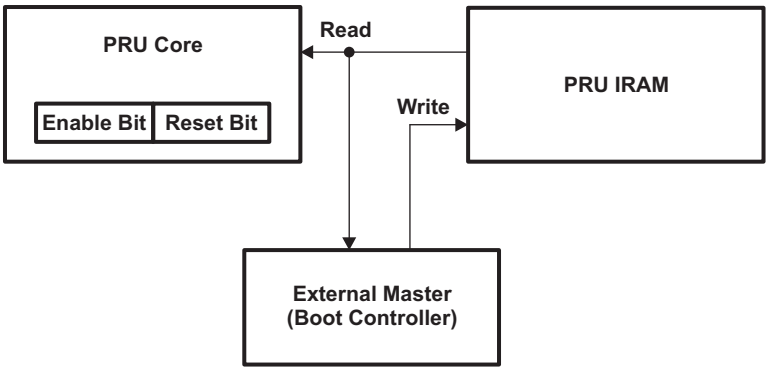


Figure 8. PRU and Boot Controller Block Diagram

When the device is first powered-on, the read bus from the PRU IRAM is not initialized and will contain random values (see Figure 9). Under unpredictable circumstances, the random value on the read bus may resemble a reserved instruction which can be interpreted by the PRU when the core is reset and not enabled.

If the PRU core executes this reserved instruction, it will not be able to properly execute the first functional op-code in the PRU IRAM when the core is later enabled. In this fail state, the PRU will never acknowledge to the Boot Controller that the reset vector table was successfully initialized and the Boot Controller will be stuck in a polling loop waiting for the PRU to complete its task.

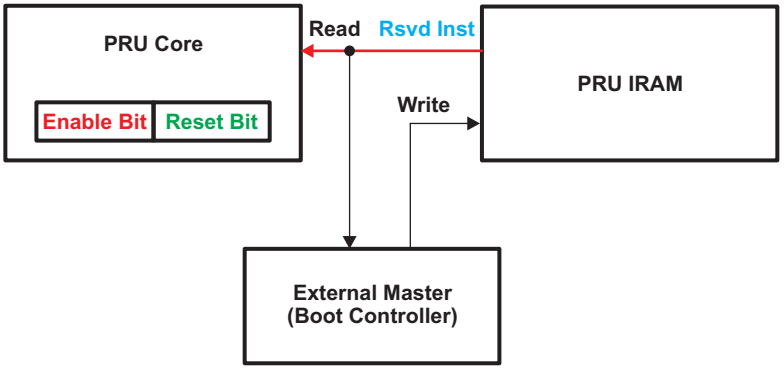


Figure 9. Boot Failure on Power-On

Copyright © 2010–2014, Texas Instruments Incorporated

Although the PRU core execution is stuck, the PRU IRAM read bus is now initialized with a non-reserved instruction that was fetched from the IRAM by the PRU core (see Figure 10). If a secondary reset is provided to the device (either POR or WARM), the PRU will be able to execute its functional instructions as expected.

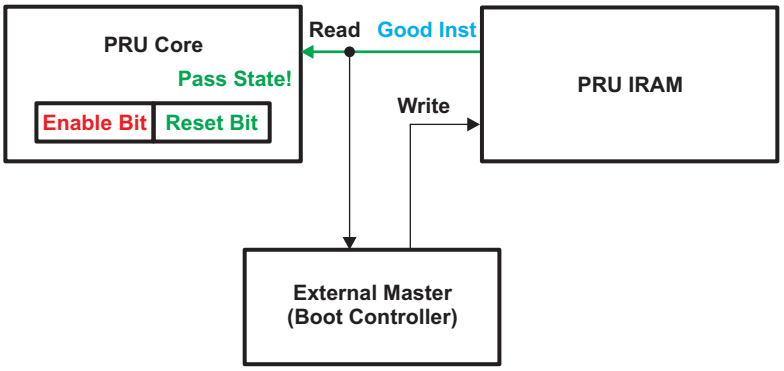


Figure 10. Secondary Reset

Note that in order to recover from this fail state with a secondary reset, the Boot Controller must be allowed to execute up to the point where the PRU has fetched a known instruction from the PRU IRAM. The approximate count of 15k clock cycles is sufficient.

The 15k clock cycle count does not include the 6192 clock cycles required to complete a device POR reset (see Figure 11). With a 24MHz crystal, the first RESET signal must be asserted high for at least 883us (or approximately 1ms).

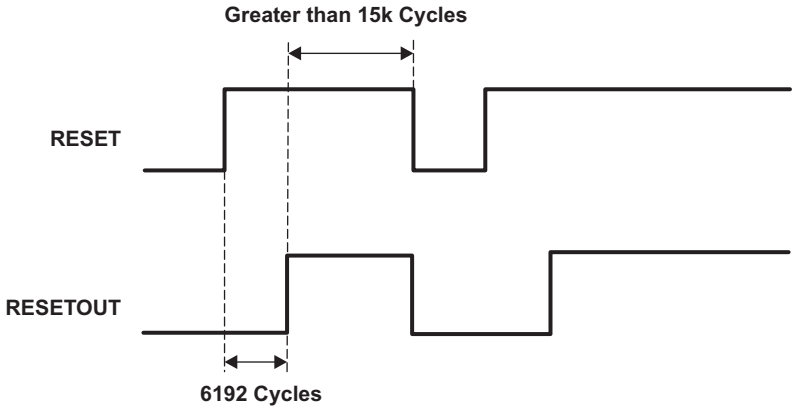


Figure 11. First POR Reset Timing

Workaround(s)

Modify the target board so that the affected device is given a secondary reset on power-up as shown in Figure 11. Two example methods are described in the sections that follow.

Although secondary resets are compatible with future silicon revisions, they are not required for devices where the root cause has been fixed via an updated Boot Controller. In order to reduce BOM costs, board designers may want to route a reset signal bypass path so that the workaround circuit can be depopulated on future PCB builds.

1. Use a reset supervisor device that includes a watchdog timeout function so that the reset supervisor will issue a secondary reset if the device fails to boot. The watchdog should be serviced with a device signal that is controlled by software. Options for servicing the watchdog timeout include GPIO, unused clock sources such as CLKOUT, or a periodic output peripheral like TIMER and ePWM. Potential reset supervisors are TPS382x for 3.3V IOs (see Figure 12) and TPS312x for 1.8V IOs.

Copyright © 2010–2014, Texas Instruments Incorporated

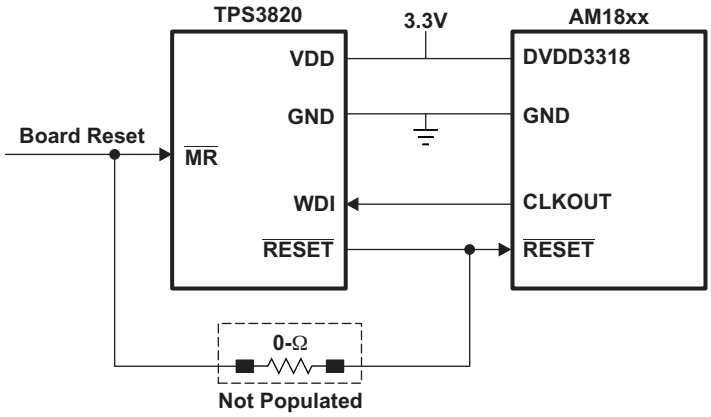


Figure 12. Reset Supervisor with Watchdog Function

The watchdog supervisor workaround is easy to implement, however the watchdog timeout period may exceed application boot-up time requirements. For example, the TPS3820 has a typical watchdog timeout period of 200ms. The second workaround can speed up the reset process.

2. Implement a logic-based secondary reset circuit which is timed using RC components. For the circuit shown in Figure 13, a single board reset control signal can trigger three logic transitions in a dual XOR gate device.

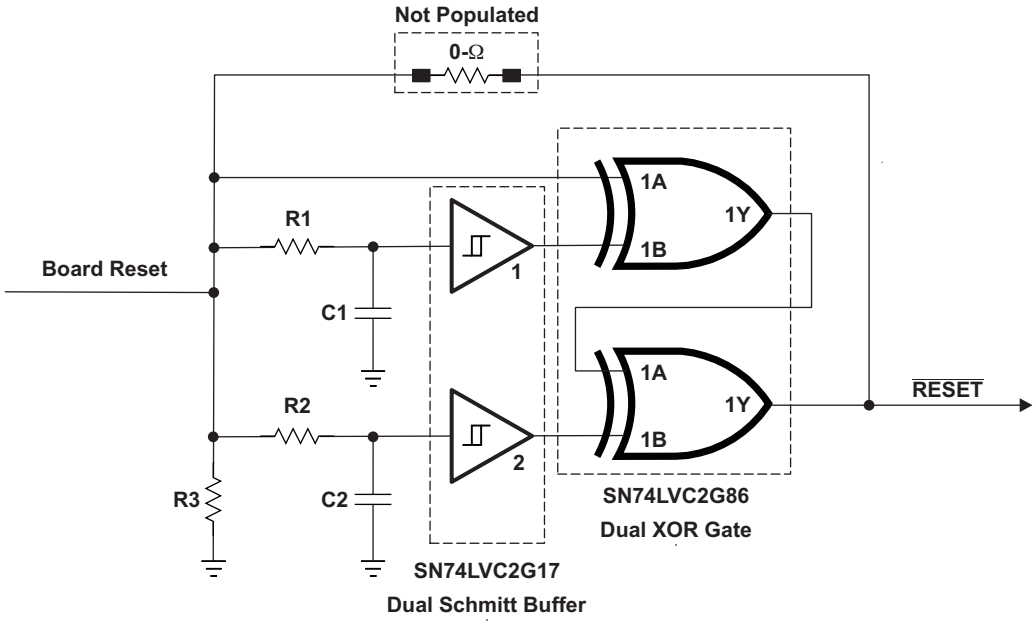


Figure 13. RC-Timed Secondary Reset

This is possible because each RC load connected to the board reset control signal can output a different rising-edge waveform. With increasing RC load, the resulting control signal will reach the Schmitt buffers' Vih level at a later point in time. Figure 14 shows the relationship between the board reset signal and the RESET signal produced by the circuit. The blue and green lines represent the voltage as seen by the Schmitt buffers. The output voltage of a charging RC circuit is defined as: Vo = Vi * (1 - e^[-t / RC])

Copyright © 2010–2014, Texas Instruments Incorporated

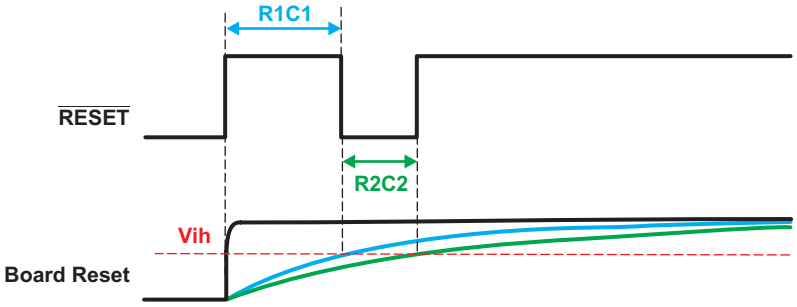


Figure 14. RESET Signal vs Board Reset

Given ideal conditions, a 3.3V board reset signal, and an input buffer Vih of 1.4V, the following set of component values would generate an initial RESET high period (R1C1 region) of approximately 2ms and a RESET low period (R2C2 region) of approximately 0.5ms:

R1 = 36k, C1 = 100nF,

R2 = 45k, C2 = 100nF,

R3 = 450k

When implementing this workaround, some important aspects should be kept in mind:

(a) The dual Schmitt buffer is included because the dual XOR gate has an input rise-time requirement that is violated by the RC circuits,

(b) The Board Reset signal must meet the XOR gate input rise-time requirement and must provide enough output current to charge the RC circuits to the target Vih level,

(c) It is critical for the Vih level of the two input buffers to be very close together so only single-device buffers should be considered for this circuit (such as the 2-in-1 dual Schmitt buffer device used in this example),

(d) Variations in the electrical characteristics of the circuit components may produce waveforms that deviate from ideal calculations, and

(e) The sole purpose of the R3 pulldown resistor is to discharge the RC components before the board reset signal is driven high. Therefore, the value selected for R3 should be sufficiently large enough to not interfere with the RC circuits as they are charging.

Copyright © 2010–2014, Texas Instruments Incorporated

Revision History

NOTE: Page numbers for previous revisions may differ from page numbers in the current version.

This silicon errata revision history highlights the technical changes made to this document.

Revision History

ADDITIONS/MODIFICATIONS/DELETIONS
Changes that were made from SPRZ314 to SPRZ314A • Advisory 2.0.6 - Potential USB 2.0 Soft Reset Timing Violation • Advisory 2.0.18 - DVDD18 Can Pull Up to 2.7V When Using Dual -Voltage IOs at 3.3V • Advisory 2.0.19 - USB2.0 On-The-Go (OTG) Session Request Protocol (SRP) is Not Supported • Added Usage Note - SATA and USB digital supplies must be powered for proper device operation
Changes that were made from SPRZ314A to SPRZ314B • Advisory 2.0.20 - Intermittent Boot Failures
Changes that were made from SPRZ314B to SPRZ314C • Advisory 2.x.18 Numbering Corrected to 2.x.7 • Advisory 2.x.19 Numbering Corrected to 2.x.8 • Advisory 2.0.20 Numbering Corrected to 2.0.9 - Fixed in Silicon Revision 2.1 • Revision Identification Table Updated for Silicon Revision 2.1 • Device Marking Figure Added
Changes that were made from SPRZ314C to SPRZ314D • Usage Note: System-Level ESD Immunity • Advisory 2.1.10 - USB0 PLL Mean Frequency Can Drift Across Large Temperature Swings
Changes that were made from SPRZ314D to SPRZ314E • Added Advisory 2.1.12 - BOOT: Internal Pullup Resistors for BOOT[7:0] Pins Are Sometimes Enabled During Reset, Leading to Boot Failures
Changes that were made from SPRZ314E to SPRZ314F Section 2.1.5, System-Level ESD Immunity Usage Note: • Added additional code to the "Disable the DLL REFCLK signal in the DDR EMIF PHY ..." bullet Section 2.1.6, DDR2/mDDR Controller: mDDR Usage Note: • Added *new* Section 3.2, *Silicon Revision 2.1 Known Design Exceptions to Functional Specifications*: • Advisory 2.1.10, *USB0 PLL Mean Frequency Can Drift Across Large Temperature Swings* – Updated/Changed the "Once the onset of PLL drift occurs, ..." paragraph • Advisory 2.1.13, *Boot: ECC Data Error in Spare Area Causes NAND Boot Failure* – Added *new*
Changes that were made from SPRZ314F to SPRZ314G Added new section: Section 2, Silicon Revision 2.3 Usage Notes and Known Design Exceptions to Functional Specifications: Section 2.1, Usage Notes for Silicon Revision 2.3: • Moved all Silicon Revision 2.1 Usage Notes to Silicon Revision 2.3 • Added new Usage Note: Section 2.1.7, McASP: Inactive Slot Usage Note Section 2.2, Silicon Revision 2.3 Known Design Exceptions to Functional Specifications: • Moved all Silicon Revision 2.1 Advisories to Silicon Revision 2.3 • Added new Advisory: Advisory 2.3.25, USB0: CPU gets Stale Receive Data from the Data Buffer located in External Memory • Added new Advisory: Advisory 2.3.26, USB0: Early DMA Completion in DMA Receive Mode and More Than One Endpoint is Transferring Data • Added new Advisory: Advisory 2.3.27, USB0: DMA Hung up in Frequent Teardowns
Changes that were made from SPRZ314G to SPRZ314H Updated/Changed Section 2.2, Silicon Revision 2.3 Usage Notes and Known Design Exceptions to Functional Specifications: • Corrected Advisory numbers of the form "2.1.x" to "2.3.x" Section 3.2, Silicon Revision 2.3 Usage Notes and Known Design Exceptions to Functional Specifications: • Moved Advisory 2.1.10, USB0 PLL Mean Frequency Can Drift Across Large Temperature Swings from Advisory 2.1.10, Silicon Revision 2.3 Usage Notes and Known Design Exceptions to Functional Specifications

Copyright © 2010–2014, Texas Instruments Incorporated

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46, latest issue, and to discontinue any product or service per JESD48, latest issue. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products (also referred to herein as “components”) are sold subject to TI’s terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its components to the specifications applicable at the time of sale, in accordance with the warranty in TI’s terms and conditions of sale of semiconductor products. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by applicable law, testing of all parameters of each component is not necessarily performed.

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

TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI components or services are used. Information published by TI regarding third-party products or services does not constitute a license 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 significant portions 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. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

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

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards which anticipate dangerous consequences of failures, monitor failures and their consequences, lessen the likelihood of failures that might cause harm and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI’s goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed a special agreement specifically governing such use.

Only those TI components which TI has specifically designated as military grade or “enhanced plastic” are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components which have ***not*** been so designated is solely at the Buyer's risk, and that Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components as meeting ISO/TS16949 requirements, mainly for automotive use. In any case of use of non-designated products, TI will not be responsible for any failure to meet ISO/TS16949.

Products		**Applications**	
Audio	www.ti.com/audio	Automotive and Transportation	www.ti.com/automotive
Amplifiers	amplifier.ti.com	Communications and Telecom	www.ti.com/communications
Data Converters	dataconverter.ti.com	Computers and Peripherals	www.ti.com/computers
DLP® Products	www.dlp.com	Consumer Electronics	www.ti.com/consumer-apps
DSP	dsp.ti.com	Energy and Lighting	www.ti.com/energy
Clocks and Timers	www.ti.com/clocks	Industrial	www.ti.com/industrial
Interface	interface.ti.com	Medical	www.ti.com/medical
Logic	logic.ti.com	Security	www.ti.com/security
Power Mgmt	power.ti.com	Space, Avionics and Defense	www.ti.com/space-avionics-defense
Microcontrollers	microcontroller.ti.com	Video and Imaging	www.ti.com/video
RFID	www.ti-rfid.com		
OMAP Applications Processors	www.ti.com/omap	**TI E2E Community**	e2e.ti.com
Wireless Connectivity	www.ti.com/wirelessconnectivity		

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2014, Texas Instruments Incorporated