



7/28/04

Errata: EP9315 - Silicon revision: E0

Reference EP9315 Data Sheet revision DS638PP2 dated July 2004

How to Determine the Silicon Revision of the Integrated Circuit

On the front of the integrated circuit, directly under the part number, is an alpha-numeric line. Characters 5 and 6 in this line represent the silicon revision of the chip. For example, this line indicates that the chip is a "E0" revision chip:

EFWAE0AM0340

This Errata is applicable only to the E0 revision of the chip.

Analog Touch Screen

Description 1

After power-on-reset, PENSTS in AR_SETUP2 register has the correct default value of "0". But after the first touch on the screen, PENSTS is stuck at "1" regardless if the screen is pressed or not.

Workaround

Configure the hardware so that as long as there is pressure on the touch surface, interrupts will occur periodically. This is done by setting the register ARXYMAXMIN so that the MIN values are 0x0 and the MAX values are 0xff. This causes the hardware to believe that while there is pressure on the surface, the pointing device is always moving. The frequency of interrupts is programmable in TSSETUP by adjusting the settling times and number of samples taken for each point. If a touch event takes longer than this time to occur, it is assumed that the touch surface has been released. For an example of this implementation, please see the source code provided with our Linux and WinCE Touch Screen drivers.

Ethernet

Description 1

The Ethernet controller does not correctly receive frames that have a size of 64 bytes.

Workaround

In order to receive frames of 64 bytes, enable the RCRCA bit in RxCTL. This will prevent the Ethernet controller from discarding the 64-byte-long frames.

Description 2

When there is inadequate AHB bus bandwidth for data to be transferred from the Ethernet controller FIFO to the receive descriptor, the Ethernet FIFO will overflow and cause the Ethernet controller to fail to receive any more packets.

This problem will also occur if the processor is too busy to service incoming packets in a timely manner. By the time that new receive descriptors are available, the data in the FIFO will contain frames that are corrupted.

It is the job of the system designer to ensure that there is adequate bandwidth for the applications being run.

Workaround

This is a rare occurrence, however at a system level it is important to reserve adequate bandwidth for the Ethernet controller. This can be accomplished by some of the following:

- Reducing the bandwidth use of other bus masters in the system.
- Lowering Ethernet rate to half duplex or 10Mbit if higher bandwidth is not required.
- Insuring that the Ethernet controller receive descriptor processing is given a high enough priority to ensure that the controller never runs out of receive descriptors.

HDLC

Description 1

When the final byte of a received packet is read into the DMA controller's buffer, the software will be notified by an HDLC RFC interrupt. However, the DMA controller may not have written the currently buffered part of the packet to memory, so that the last one to fifteen bytes of a packet may not be accessible.

Workaround

To insure that the DMA channel empties the buffer, do the following (in the HDLC interrupt handler, for example):

- 1) Note the values in the MAXCNTx and REMAIN registers for the DMA channel. The difference is the number of bytes read from the UART/HDLC, which is the size of the HDLC packet. Call this number N. Note that the BC field of the UART1HDLCRXInfoBuf register should also be N.
- 2) Temporarily disable the UART DMA RX interface by clearing the RXDMAE bit in the UART1DMACtrl register.
- 3) Wait until the difference between the CURRENTx and BASEx registers in the DMA channel is equal to N + 1.

At this point, the rest of the packet is guaranteed to have been written to memory. Using this method will cause an extra byte to be read from the UART by the DMA channel and also written to memory. This last byte should be ignored.

Raster

Description 1

If the raster engine is using single scan mode, two and two thirds per pixel mode (3 bits per pixel over an 8-bit bus) works correctly. If the raster engine is programmed to use two and two thirds pixels per clock shift mode with dual scan enabled, it will not generate valid timings for dual scan displays.

Workaround

There is no known workaround at this time.

Description 2

YCrCb formatted video will not produce the valid synchronization signals in 656 video mode.

Workaround

Design the system with an NTSC/PAL DAC that accepts RGB input signals.

SDRAM Controller

Description 1

Using the SDRAM controller in auto-precharge mode will produce system instability at external bus speeds greater than 50MHz.

Workaround

Do not turn on the auto-precharge feature of the SDRAM controller if the external bus speed will be greater than 50 MHz.

MaverickCrunch™

Description 1

Under certain circumstances, data in coprocessor registers or in memory may be corrupted. The following sequence of instructions will cause the corruption:

- 1) Let the first instruction be any coprocessor instruction that is not executed, for any of the following reasons:
 - It fails its condition code check.
 - It appears in the processor pipeline but is not executed due to a taken branch, an exception or an interrupt.
- 2) Assume that this first instruction is stalled by the coprocessor due to an internal dependency.
- 3) Let the second instruction be any coprocessor load or store 64/double: `cfldr64`, `cfldrd`, `cfstr64`, `cfstrd`.

If the second instruction is a load, the upper word in the target register will generally get an incorrect value. If the second instruction is a store, the word immediately following the second target memory location will be written; that is, instead of just writing two consecutive 32-bit words (a 64-bit value or a double value) to memory, a third 32-bit word immediately following this will be written, leading to memory corruption.

Consider a simple example with a store instruction:

```
cfaddne    c0, c1, c2          ; assume this does not execute
cfstr64    c3, [r2, #0x0]
```

Three words will be written to memory. The correct values will appear at the memory location pointed to by r2, and r2 + 0x4. Another value will be written at r2 + 0x8.

Consider now an example with a load instruction:

```
cfaddne    c0, c1, c2          ; assume this does not execute
cflldr     c3, [r2, #0x0]
```

The final value in c3 will be incorrect. The lower 32 bits will be correct, while the upper 32 bits will be incorrect.

Finally, consider a case where a branch occurs:

```
target
cflldr     c3, [r2, #0x0]
b         target
nop
cfadd     c0, c1, c2          ; though in pipeline, this does not execute
```

Note that the above examples assume that the cfaddne or cfadd would busy-wait (for whatever reason) if actually executed. If not, the execution of the following instruction would be correct.

Workaround

The simplest workaround is to insure that no two such instructions ever appear in the instruction stream consecutively. Specifically, a conditional coprocessor instruction should not precede a load/store 64/double. Simply inserting another ARM or coprocessor instruction accomplishes this:

```
cfaddne    c0, c1, c2          ; assume this does not execute
nop        ; inserted extra instruction here
cflldr     c3, [r2, #0x0]
```

Cases where branches may be taken also needs to be handled; in this particular case, the first instruction is moved earlier in the instruction stream by exchanging it with the previous one:

```
target
cflldr     c3, [r2, #0x0]
b         target
cfadd     c0, c1, c2          ; though in pipeline, this does not execute
nop
```

To avoid this error in exception and interrupt handlers, the first instruction in an interrupt or exception handler should not be a coprocessor instruction. Since the first instruction is a branch, this error will not appear.

Description 2

Under certain circumstances, incorrect values may be used for arithmetic calculations or stored in memory. The error appears as follows.

- 1) Execute a coprocessor instruction whose target is one of the coprocessor general purpose register c0 through c15.
- 2) Let the second instruction be an instruction with the same target, but not be executed for one of the following reasons:
 - It fails its condition code check.
 - It appears in the processor pipeline but is not executed due to a taken branch, an exception or an interrupt.
- 3) Execute a third instruction at least one of whose operands is the target of the previous two instructions.

For example, assume no pipeline interlocks other than the dependencies involving register c0 in the following instruction sequence:

```
cfadd32      c0, c1, c2
cfsub32ne   c0, c3, c4      ; assume this does not execute
cfstr32     c0, [r2, #0x0]
```

In this particular case, the incorrect value stored at the address in r2 is the previous value in c0, not the expected one resulting from the cfadd32.

Workaround

Insure that this sequence of instructions does not occur.

Another solution is to insure that the first and third instructions are sufficiently far apart in the instruction stream by placing five other instructions between them:

```
cfadd32      c0, c1, c2
nop          ; inserted extra instruction here
nop          ; inserted extra instruction here
cfsub32ne   c0, c3, c4      ; assume this does not execute
nop          ; inserted extra instruction here
nop          ; inserted extra instruction here
nop          ; inserted extra instruction here
cfstr32     c0, [r2, #0x0]
```

The five intervening instructions need not be nops and may appear before or after the second instruction.

Note that it is the instruction stream as executed by the processor, not the instructions as they appear in the source code, which is relevant. Hence, cases where the program flow changes between the first and third instruction must be considered.

The first five instructions of an exception or interrupt handler should not be coprocessor instructions.

Description 3

Under certain circumstances, data in coprocessor general purpose registers or in memory may be corrupted. The error appears as follows.

- 1) Let the first instruction be a serialized instruction that does not execute for one of the following reasons:
 - It fails its condition code check.
 - It appears in the processor pipeline but is not executed due to a taken branch, an exception or an interrupt.

An instruction is serialized if either exceptions are enabled, or the instruction is a DSPSC move (cfmv32sc or cfmvsc32).

2) The immediately following instruction is a two-word coprocessor load or store (cflldr64, cflldr, cfstr64, or cfstrd).

In the case of a load, only the lower 32 bits (the first word) will be loaded into the target register. For example:

```
cfadd32ne    c0, c1, c2    ; assume this does not execute
cflldr64     c3, [r2, #0x0]
```

The lower 32 bits of c3 will correctly become what is at the memory address in r2, but the upper 32 bits of c3 will not become what is at address r2 + 0x4.

In the case of a store, only the lower 32 bits (the first word) will be stored into memory. For example:

```
cfadd32ne    c4, c5, c6    ; assume this does not execute
cfstr64      c3, [r2, #0x0]
```

The lower 32 bits of c3 will be correctly written to the memory address in r2, but the upper 32 bits of c3 will not be written.

Workaround

Separating the first and second instruction by one instruction will avoid this error whether or not the coprocessor is operating in serialized or unserialized mode. For example:

```
                                ; load sequence
cfadd32ne    c0, c1, c2    ; assume this does not execute
nop          ; inserted extra instruction here
cflldr64     c3, [r2, #0x0] ; store sequence
cfadd32ne    c4, c5, c6    ; assume this does not execute
nop          ; inserted extra instruction here
cfstr64      c3, [r2, #0x0]
```

Note that the effect of branches should also be accounted for, as it is the instruction stream as seen by the coprocessor that matters, not the order of instructions in the source code. The two instructions following a taken branch may be seen by the coprocessor and then not executed, and would be treated exactly as the first instruction above.

The asynchronous invocation of interrupt/exception handlers will not expose this error, as their first instruction is always a branch.

Description 4

When no exceptions are enabled (coprocessor is operating unserialized) and forwarding is enabled, memory can be corrupted when two types of instructions appear in the instruction stream with a particular relative timing.

1) Execute an instruction that is a data operation (not a move between ARM and coprocessor registers) whose destination is a general purpose register c0 through c15.

2) Execute an instruction that is a two-word coprocessor store, either cfstr64 or cfstrd, where the destination register of the first instruction is the source of the store instruction (that is, the second instruction stores the result of the first one to memory).

3) Finally, the first and second instruction must appear to the coprocessor with the correct relative timing;

this timing is not simply proportional to the number of intervening instructions and is difficult to predict in general.

The result is that the lower 32 bits of the result stored to memory will be correct, but the upper the 32 bits will be wrong. The value appearing in the target register will still be correct.

Workaround

One workaround is to operate the coprocessor without forwarding enabled, with a possible decrease in performance.

Another is to operate in serialized mode by enabling at least one exception, with significantly reduced performance.

Another workaround is to insure that at least seven instructions appear between the first and second instructions that cause the error.

Another possible workaround is to insure that the second instruction appears earlier in the instruction stream or early enough to avoid the error. In general, this is complex to determine, though if no instructions separate the first and second instruction, the error will never manifest itself.

Note that branches and interrupts/exceptions must be considered, because it is the instruction stream seen by the processor and coprocessor that can expose this error, including the effects of branches, asynchronous interrupts and exceptions. To avoid this error due to interrupts and exceptions, simply do not allow the first seven instructions in an exception or interrupt handler to be coprocessor instructions.

Description 5

When operating in serialized mode, `cfrshl32` and `cfrshl64` do not work properly. The instructions shift by an unpredictable amount, but cause no other side effects.

The coprocessor is in serialized mode when:

- At least one exception is enabled by setting one of the following bits in the DSPSC: IXE, UFE, OFE, or IOE.
- Serialization is not specifically disabled by setting bit AEXC in the DSPSC.

Workaround

One workaround is to avoid these instructions. With this approach, an alternative instruction sequence may accomplish the shift with the following steps:

- Move the data to be shifted to ARM register(s)
- Shift the data using non-coprocessor instructions
- Move the shifted data back to the coprocessor.

Another workaround is to never operate in serialized mode. With this approach, synchronous exceptions are not possible.

Description 6

If an interrupt occurs during the execution of `cfldr32` or `cfmv64lr`, the instruction may not sign extend the result correctly.

Either instruction places a 32 bit value into the lower half of a MaverickCrunch general purpose register and sign extends the high (32nd) bit through the upper half of the register. An IRQ or FIQ may cause the ARM processor to interrupt the execution of any instruction on any cycle. If this happens to either of these instructions at the right time, it will properly load the low 32 bits of the target register, but instead of sign extending it will replicate the low 32 bit into the upper 32 bits. Code that depends on sign extension will fail to operate correctly.

Workaround

Possible workarounds include:

- Disable interrupts when executing `cfldr32` or `cfmv64lr` instructions.
- Avoid executing these two instructions.
- Do not depend on the sign extension to occur; that is, ignore the upper word in any calculations involving data loaded using these instructions.
- Add extra code to sign extend the lower word after it is loaded by explicitly forcing the upper word to be all zeroes or all ones, as appropriate. It is possible to do this selectively in the exception/interrupt handler code. If the instruction preceding the interrupted instruction can be determined, and it is a `cfldr32` or `cfmv64lr`, the instruction may be re-executed or explicitly sign extended before returning from interrupt/exception.

Description 7

The coprocessor can incorrectly update one of its destination accumulators even if the coprocessor instruction should not have been executed or is canceled by the ARM processor. This error can occur if the following is true:

- 1) The first instruction must be a coprocessor compare instruction, `cfcmp32`, `cfcmp64`, `cfcmps`, and `cfcmpd`.
- 2) The second instruction must have an accumulator as a destination:
 - Moves to accumulators: `cfmva32`, `cfmva64`, `cfmval32`, `cfmvam32`, `cfmvah32`.
 - Arithmetic into accumulators: `cfmadd32`, `cfmadda32`, `cfmsub32`, `cfmsuba32`.
- 3) The second instruction is not executed for one of the following reasons:
 - It fails its condition code check.
 - It appears in the processor pipeline but is not executed due to a taken branch, an exception or an interrupt.

Example 1: In this case the second instruction may modify `a2` even if the condition is not matched.

```
cfcmp32    r15, c0, c5
cfmva64ne  a2, c8
```

Example 2: In this case the second instruction may modify `a2` even if an interrupt or exception causes it to be canceled and re-executed after the interrupt/exception handler returns.

```
cfcmp32    r15, c0, c5
cfmadda    a2, a2, c0, c1
```


Workaround

The workaround for this issue is to insure that at least one other instruction appears between these instructions. For example, possible fixes for the instructions sequences above are:

```
cfcmp32      r15, c0, c5
nop
cfmva64ne    a2, c8
                                     and
cfcmp32      r15, c0, c5
nop
cfmadda      a2, a2, c0, c1
```

Description 8

If a data abort occurs on an instruction preceding a coprocessor data path instruction that writes to one of the accumulators a0 - a3, the accumulator may be updated even though the instruction was canceled. Instructions that write the accumulators are: cfmva32, cfmva64, cfmval32, cfmvam32, cfmvah32, cfmadd32, cfmsub32, cfmadda32 and cfmsuba32.

For example:

```
str          r7, [r0, #0x1d] ; assume this causes a data abort
cfmadda32    a0, a2, c0, c1
```

The second instruction will update a0 even though it should be canceled due to the data abort on the previous instruction.

Workaround

A complete software workaround requires ensuring that data aborts do not occur due to any instruction immediately preceding the coprocessor instructions described in this errata. The only way to ensure this is to not allow memory operations immediately preceding these types of instructions. For example, the fixes for the instructions above are:

```
str          r7, [r0, #0x1d] ; assume this causes a data abort
nop
cfmadda32    a0, a2, c0, c1
```

Description 9

The coprocessor will erroneously update an accumulator if the coprocessor instruction that updates an accumulator is canceled and is followed by a coprocessor instruction that is *not* a data path instruction. Coprocessor data path instructions include any instruction that does not move data to or from memory or to or from the ARM registers. This error will occur under the following conditions:

- 1) The first instruction must update a coprocessor accumulator. These include:
 - Moves to accumulators: cfmva32, cfmva64, cfmval32, cfmvam32, cfmvah32.
 - Arithmetic into accumulators: cfmadd32, cfmadda32, cfmsub32, cfmsuba32.
- 2) The first instruction is not executed for one of the following reasons:
 - It fails its condition code check.
 - It appears in the processor pipeline but is not executed due to a taken branch, an exception or an interrupt.

3) The second instruction is not a coprocessor datapath instruction.

For example:

```
cfmva64ne    a2, c3
cfmvr64l     r4, c15
```

If the first instruction should not execute or is interrupted, it may incorrectly update a2.

Workaround

Because any instruction may be canceled due to an asynchronous interrupt, the most general software workaround is to insure that no instruction that updates an accumulator is followed immediately by a non-datapath coprocessor instruction. For example, the fix for the instruction sequence above is:

```
cfmva64ne    a2, c3
nop
cfmvr64l     r4, c15
```

Description 10

An instruction that writes a result to an accumulator may cause corruption of any of the four accumulators when the coprocessor is operating in serialized mode. Instructions that write an accumulator are: cfmva32, cfmva64, cfmval32, cfmvam32, cfmvah32, cfmadd32, cfmsub32, cfmadda32 and cfmsuba32.

The coprocessor is operating in serialized mode when at least one exception is enabled by setting one of the following bits in the DSPSC: IXE, UFE, OFE, IOE and when serialization is not specifically disabled by setting bit AEXC in the DSPSC.

For example, the following sequence of instructions may corrupt a2 if the second instruction is not executed.

```
cfmadda32    a0, a2, c0, c1
cfmadda32ne  a2, c3, c0, c1
```

Workaround

The only workaround for this issue is to operate the coprocessor in unserialized mode.

Description 11

An erroneous memory transfer to or from MaverickCrunch registers can occur given the following conditions are satisfied:

- 1) Two consecutive MaverickCrunch load/store instructions appear in the instruction stream.
- 2) The first instruction is a two-word load or store, one of cfldr64, cfstr64, cfldrd, and cfstrd, that fails its condition code check and does not busy-wait.
- 3) The second instruction is MaverickCrunch load or store, which is executed and does not busy-wait.

When the error occurs, the result is either MaverickCrunch register or memory corruption. Here are several examples:

```
cfstr64ne    c0, [r0, #0x0]    ; assume does not execute
cfldrs       c2, [r2, #0x8]    ; could corrupt c2!
```

```
cfldr64ne    c0, [r0, #0x0]    ; assume does not execute
cfldr64ne    c2, [r2, #0x8]    ; could corrupt memory!

cfldr64ne    c0, [r0, #0x0]    ; assume does not execute
cfldr64ne    c2, [r2, #0x8]    ; could corrupt c2!
```

Workaround

The software workaround involves avoiding a pair of consecutive instructions with these properties. For example, if a conditional MaverickCrunch two-word load/store appears, insure that the following instruction is not a coprocessor load/store:

```
cfstr64ne    c0, [r0, #0x0]    ; assume does not execute
nop          ; separate two instructions
cfldr64ne    c2, [r2, #0x8]    ; c2 will be ok
```

Another workaround is to insure that the first instruction is not conditional:

```
cfstr64      c0, [r0, #0x0]    ; executes
cfldr64ne    c2, [r2, #0x8]    ; c2 will be ok
```

Note: If both instructions depend on the same condition code, the error should not occur, as either both or neither will execute.

Contacting Cirrus Logic Support

For all product questions and inquiries contact a Cirrus Logic Sales Representative.
To find one nearest you go to www.cirrus.com

IMPORTANT NOTICE

"Preliminary" product information describes products that are in production, but for which full characterization data is not yet available. Cirrus Logic, Inc. and its subsidiaries ("Cirrus") believe that the information contained in this document is accurate and reliable. However, the information is subject to change without notice and is provided "AS IS" without warranty of any kind (express or implied). Customers are advised to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgment, including those pertaining to warranty, patent infringement, and limitation of liability. No responsibility is assumed by Cirrus for the use of this information, including use of this information as the basis for manufacture or sale of any items, or for infringement of patents or other rights of third parties. This document is the property of Cirrus and by furnishing this information, Cirrus grants no license, express or implied under any patents, mask work rights, copyrights, trademarks, trade secrets or other intellectual property rights. Cirrus owns the copyrights associated with the information contained herein and gives consent for copies to be made of the information only for use within your organization with respect to Cirrus integrated circuits or other products of Cirrus. This consent does not extend to other copying such as copying for general distribution, advertising or promotional purposes, or for creating any work for resale.

CERTAIN APPLICATIONS USING SEMICONDUCTOR PRODUCTS MAY INVOLVE POTENTIAL RISKS OF DEATH, PERSONAL INJURY, OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE ("CRITICAL APPLICATIONS"). CIRRUS PRODUCTS ARE NOT DESIGNED, AUTHORIZED OR WARRANTED FOR USE IN AIRCRAFT SYSTEMS, MILITARY APPLICATIONS, PRODUCTS SURGICALLY IMPLANTED INTO THE BODY, LIFE SUPPORT PRODUCTS OR OTHER CRITICAL APPLICATIONS (INCLUDING MEDICAL DEVICES, AIRCRAFT SYSTEMS OR COMPONENTS AND PERSONAL OR AUTOMOTIVE SAFETY OR SECURITY DEVICES). INCLUSION OF CIRRUS PRODUCTS IN SUCH APPLICATIONS IS UNDERSTOOD TO BE FULLY AT THE CUSTOMER'S RISK AND CIRRUS DISCLAIMS AND MAKES NO WARRANTY, EXPRESS, STATUTORY OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR PARTICULAR PURPOSE, WITH REGARD TO ANY CIRRUS PRODUCT THAT IS USED IN SUCH A MANNER. IF THE CUSTOMER OR CUSTOMER'S CUSTOMER USES OR PERMITS THE USE OF CIRRUS PRODUCTS IN CRITICAL APPLICATIONS, CUSTOMER AGREES, BY SUCH USE, TO FULLY INDEMNIFY CIRRUS, ITS OFFICERS, DIRECTORS, EMPLOYEES, DISTRIBUTORS AND OTHER AGENTS FROM ANY AND ALL LIABILITY, INCLUDING ATTORNEYS' FEES AND COSTS, THAT MAY RESULT FROM OR ARISE IN CONNECTION WITH THESE USES.

Cirrus Logic, Cirrus, MaverickCrunch, MaverickKey, and the Cirrus Logic logo designs are trademarks of Cirrus Logic, Inc. All other brand and product names in this document may be trademarks or service marks of their respective owners.

Microsoft and Windows are registered trademarks of Microsoft Corporation.

Microwire is a trademark of National Semiconductor Corp. National Semiconductor is a registered trademark of National Semiconductor Corp.

Texas Instruments is a registered trademark of Texas Instruments, Inc.

Motorola is a registered trademark of Motorola, Inc.

LINUX is a registered trademark of Linus Torvalds.