

AMD Geode™ CS5535 Companion Device Silicon Revision A3 Specification Update



1.0 Scope

This document discusses known issues of silicon revision A3 of the AMD Geode™ CS5535 companion device.

Table 1-1 provides a summary of the issues. A detailed description of each issue, its impact, and a recommended resolution/fix follow.

The silicon revision number, 1.2, is printed on the chip (top-side of PBGA in the seventh line of the part marking).

Software Readable IDs for Rev A3

- CS5535 Silicon Revision:
 - MSR 51700017h[7:0] = 13h
 - PCI Index 08h[15:0] = 0013h
 - TAP Controller Instruction 8FFFFAh = 13h

- Device (Module) Revisions:
 - GLIU: MSR 51010000h[7:0] = 01h
 - GLPCI_SB: MSR 51700000h[7:0] = 00h
 - ACC: MSR 51500000h[7:0] = 00h
 - ATAC: MSR 51300000h[7:0] = 00h
 - USBC1: MSR 51600000h[7:0] = 00h
 - USBC2: MSR 51200000h[7:0] = 00h
 - DD: MSR 51400000h[7:0] = 00h
 - GLCP: MSR 51700000h[7:0] = 00h

All references to the AMD Geode GX processor indicate the GX 533@1.1W processor*, GX 500@1.0W processor*, and/or GX 466@0.9W processor*.

Note: This is revision D of this document. The change from revision C (dated February 2006) is the document was made non confidential.

Table 1-1. Issue Summary

Issue# ¹	Description
1	PIT deviation from industry standard operation
2	Debug port hangs system if used
3	Incorrect operation of PIT count mode 3 value load
4	Restrictions on AEOI bit use in the PIC
5	DMA rotating mode does not rotate correctly
6	PMC thermal alarm recovery behavior
8	PMC Sleep_Clock control disable side effect
15	GLIU Data Compare cannot generate a hit
37	GLIU Error MSR unexpected address flag failure
40	UART registers can not be accessed during DMA cycles
44	8237 software DMA memory write from LPC fails
54	I/O cycles with address 10000h or higher hangs the system
57	No over-voltage protection on USB
58	LBARs allow specification of an I/O base at 10000h or higher
70	PIC fails if processed interrupt still high when clock goes off
75	ATAC PRD may prefetch incorrectly
76	ACC PRD may prefetch incorrectly
82	UDMA write CRC corruption after device termination

*The AMD Geode GX 533@1.1W processor operates at 400 MHz, the AMD Geode GX 500@1.0W processor operates at 366 MHz, and the AMD Geode GX 466@0.9W processor operates at 333 MHz. Model numbers reflect performance as described here: <http://www.amd.com/connectivitysolutions/geodegxbenchmark>.

Table 1-1. Issue Summary (Continued)

Issue# ¹	Description
84	ATAC SSMI during DMA does not work correctly
85	Main GLIU clock domain does not awaken when statistics counter generates an ASMI
89	Byte accesses to ACPI PM1_STS and PM1_EN (ACPI I/O Offset 00h and 02h) and hang
90	WORD and BYTE writes to PMC registers may corrupt other PMC register writes
92	KEL HceStatus[3] does not report writes to I/O Port 064h
93	KEL Snoop corrupts HCE state
94	Possible PCI memory read failure on back-to-back fetches
96	PIC does not support Polling mode
97	Automatic clock gating for UART clock domain does not work in CEIR mode
98	IRDY# and DEVSEL# do not meet t _{VAL} minimum timing
99	SSMI may not be synchronous if IRQ occurs at the same time
100	PIO access to IDE drive with UDMA bus master active fails
101	Most wake event inputs do not work from the ACPI S3 state
102	RTC CMOS RAM accesses incorrectly
103	I/Os to LPC devices do not work during block and demand mode legacy DMA
104	Clock gating in the USB Controllers does not work
105	ATAPI packet command after bus master enable to CDROM drive fails
106	AuxOutputfull bit in KEL incorrectly affected by read from Port 60h
107	tACK does not affect UDMA timing as expected
108	ATAC can hang if PRD address is not DWORD aligned
109	LPC Aborts on Long_Sync or Ready_Sync for I/O writes and reads
111	NEGACK status bit behaves incorrectly in the SMB controller
112	Retry counter does not distinguish between accesses in GLPCI
113	Write to GPIO High Bank Feature Bit Registers failed after suspend
114	SMBus Hold Time too short
115	Missing LOWBAT_FLAG after abnormal standby
116	Sleep_X is not deasserted if delay is larger than PM_SED delay

1. Issue numbers may not be sequential since issues are omitted when resolved.

2.0 Issues

1. PIT deviation from industry standard operation

Description: *The Indispensable PC Hardware Book* says: "If you output a counter latch command and later a read-back command to determine the counting mode without having read the counter value before, then the PIT first supplies the status byte with the first IN instructions and only afterwards the byte(s) that indicate the counter value". Effectively, a read-back command should cancel an uncompleted counter latch command.

For the PIT in the CS5535, this 'cancellation' does not happen. The PIT first supplies the counter value.

Implications: The issue only occurs with a 'questionable' programming sequence under which a new command is issued before completing the current command. Therefore, the negative effects of this deviation are negligible.

Resolution: Do not issue a read-back command until the previous latch command has completed.

2. Debug port hangs system if used

Description: The GeodeLink™ Interface Unit (GLIU) incorporates a debug port that allows internal GeodeLink traffic to be copied to an external port so it can be viewed. The GLIU uses a special packet type to indicate debug. The debug port should detect this packet type and always forward it to the LPC bus. The debug port detects the packet type but hangs.

Implications: Since there is no 'internal only' GeodeLink traffic in the CS5535, the loss of this feature is minimal. All CS5535 internal traffic exits or enters the CS5535 via the PCI bus. Hence, all internal traffic is also external.

Resolution: None. Do not use the debug port.

3. Incorrect operation of PIT count mode 3 value load

Description: A new count value should not affect the current counting. However, the new count value is immediately loaded without finishing the counting process of the original count value.

Implications: Software that relies on this operation will not operate properly.

Resolution: None. Avoid reliance on this property of the PIT.

4. Restrictions on AEOI bit use in the PIC

Description: Upon receiving the first interrupt acknowledge pulse, the priority encoder should transfer the highest priority bit from the IRR to the ISR and clear the corresponding bit in the IRR. However, the XPIC (Extended PIC) does not clear the IRR bit until one clock after the second interrupt acknowledge pulse. When the PIC performs an auto-end-of-interrupt (AEOI), this causes two interrupts to be sent to the CPU instead of one.

Implications: The AEOI bit is bit 1 of Initialization Control Word Four (ICW4). When AEOI is enabled in the master PIC, the master PIC automatically issues a non-specific EOI command on the rising edge of the last interrupt acknowledge pulse. The AEOI bit can only be enabled on the master PIC. AEOI should only be enabled when a nested multilevel interrupt structure is not needed.

PCs need a nested multilevel interrupt structure. DOS uses non-specific EOIs. Microsoft® Windows®, Linux, BeOS, and FreeBSD all use specific EOIs.

This feature is not used on PCs.

Resolution: None.

5. DMA rotating mode does not rotate correctly

Description: There are four transfer channels/requests. Given a case where each request arrives one clock cycle later with respect to the previous request, then the transfers should be performed in the same order as the incoming requests. With rotating priority, the most recent channel to be serviced becomes the lowest priority, with the priorities of the other channels rotated accordingly. This design is supposed to prevent any one channel from monopolizing the system. However, the priorities do not always rotate as expected.

Implications: The transfers on multiple channels do not complete successfully. However, the sequence of the transfers do not rotate correctly. Since legacy devices using such techniques are so slow compared to existing system speeds, and since there are so few multiple legacy DMA operations, there is no significant impact of incorrect rotation.

Resolution: None.

6. PMC thermal alarm recovery behavior

Description: In the event of a thermal alarm, the system goes into Faulted Standby state as expected. However, there is no recovery mechanism available for the PMC to recover on its own because the thermal alarm resides in the Work domain instead of the Standby domain.

During a Faulted Standby state, the thermal alarm timer is blocked by the Work/Standby interface. Specifically, the thermal alarm ball is in the Working power domain. Hence the alarm state outside the component can not be detected while in Standby state.

Once out of Standby, thermal alarm (THRM_ALARM#) again comes into play. If it is still asserted, its timer will start again.

Implications: Thermal alarm causes the power to be removed from the system. This should cool the system and cause the alarm state to de-assert. After a delay, the system will power up again. Generally, the alarm has de-asserted and the system returns to normal. If the alarm has not de-asserted within a specified delay, then the Standby state is forced again. This process repeats until the thermal alarm condition is cleared (e.g., the temperature goes down).

Resolution: None.

8. PMC Sleep_Clock control disable side effect

Description: The use of SLEEP_X, SLEEP_Y, and SLP_CLK_EN# is optional and subject to optional delays. All of these options are specified as independent. However, if any option is desired, the Sleep Clock option must be enabled.

Implications: Operation may not be as expected if the Sleep Clock option is not properly enabled.

Resolution: If the Sleep Clock option is not desired and if one or more of the other options above are desired, then disable the Sleep Clock at the GPIO but enable it in the appropriate PMC register. A disable at the GPIO has the same system effect as a disable in the PMC. Thus always enable Sleep Clock at the appropriate PMC register and then enable/disable actual use via the GPIO controls.

15. GLIU Data Compare cannot generate a hit

Description: The GLIU contains a debug feature that generates a signal if a specified data pattern is detected in a data packet passing through it. This signal is passed to the GLCP as an "error" for debug. The detection operation does not work.

Implications: The detection operation is designed to support debug efforts only and hence this failure has no operational impact. The operation is controlled by the following registers:

DA_COMPARE_VAL_LO,
DA_COMPARE_VAL_HI,
DA_COMPARE_MASK_LO,
and DA_COMPARE_MASK_HI (MSR 510100D0h-510100D3h). Due to this issue, these registers should be left in their default state and not written.

Resolution: None.

37. GLIU Error MSR unexpected address flag failure

Description: When the GLIU receives an MSR packet request with an address that does not match an implemented register, this is considered an "Unexpected Address Error". If the "Unexpected Address" mask bit in the GLIU's GLD_MSR_ERROR is in the enabled state (MSR 51010003h[1] = 0) and such an error occurs, then two actions should happen: 1) the "Unexpected Address" flag (MSR 51010003h[33]) should become active; 2) the "Exception" flag in the response packet should be set. However, neither of these actions occur.

Implications: No functional impact. This is a rare failure that only occurs in debug (i.e., does not occur in an operational system).

Resolution: None.

40. UART registers can not be accessed during DMA cycles

Description: The first x86 implementations (8086, 80286, 80386, 80486, etc.) had HOLD/H LDA signals that directly or indirectly connected to the 8237 DMA components. During block or demand DMA modes, the 8237 would hold the system bus until DMA transfers were complete. Thus, any software attempts to access the UARTs during DMA were naturally blocked.

Current architectures such as the AMD Geode GX processor and the CS5535 companion device do not hold the main CPU during legacy DMA transfers since it would have a substantial system performance impact. Thus, it is possible for software to attempt to access the UARTs during legacy DMA. Under the CS5535 companion device implementation, such write accesses are discarded and read accesses return all ones.

Implications: Software that accesses the UARTs during 8237 block or demand mode DMA may experience unexpected results.

Resolution: Software should not access the UARTs during the modes detailed above. Wait until DMA completion as signaled by DMA terminal count IRQ. Alternatively, do not use block or demand modes. single byte transfer mode yields the same performance and does hold the Local bus. Thus, the UARTs can be accessed during DMA.

44. 8237 software DMA memory write from LPC fails

Description: The 8237 DMA Controller contains a legacy feature called "Software DMA". This is a diagnostic feature designed to generate memory access cycles. The actual data moved with the DMA is a "don't care", specifically, software DMA has no operational use. Conceptually, the combinations supported are "don't care" data from UART or LPC bus written to memory and "don't care" data read from memory to UART or LPC bus. However, the LPC bus "don't care" data written to memory combination does not work. Under the failing combination, the LPC Controller issues a DMA acknowledge to the LPC bus. However, since no device actually asked for a DMA, no device responds. The LPC Controller then aborts the DMA request and does not write to memory. The overall software DMA operation will complete but no "don't care" data is written to memory.

Implications: No operational impact.

Resolution: Do not use the failing combination in diagnostic tests.

54. I/O cycles with address 10000h or higher hangs the system

Description: Any I/O address 10000h or higher on the PCI bus should not be claimed by the CS5535. Assuming no other PCI device made a claim, such an address would master abort. Such an address can be generated by an x86 processor by performing a 16-bit access to FFFFh which will increment to 10000h.

However, the CS5535 claims addresses at or above 10000h. These default to the Diverse Device which hangs for such addresses.

Implications: Operationally, there is no impact because no "normal" software would make such an access. This issue is likely only to occur under new software code that is not fully debugged.

Resolution: Do not generate PCI I/O cycles with addresses at or above 10000h.

57. No over-voltage protection on USB

Description: The following paragraph is from *USB Specification 1.1, Section 7.1.1 USB Driver Characteristics*:

"USB devices must be capable of withstanding continuous exposure to the waveforms shown in Figure 7-1 while in any drive state. These waveforms are applied directly into each USB data pin from a voltage source with an output impedance of 39. The open-circuit voltage of the source shown in Figure 7-1 is based on the expected worst-case overshoot and undershoot."

The above paragraph and referenced figure constitute an over-voltage protection specification. The CS5535 component does not meet this specification.

Implications: Use of external over-voltage protection circuits is required.

Resolution: Use an appropriate over-voltage protection circuit. Such a circuit will be detailed in a forthcoming update of this specification update document.

58. LBARs allow an I/O base at 10000h or higher

Description: The ATA-5 Controller (ATAC) and Diverse Device (DD) allow I/O Local Base Address Register (LBAR) values at 10000h or higher. Such addresses are illegal under an x86 system. Such settings should not be used.

Implications: Operationally, LBAR values of 10000h or higher are not used by the BIOS or OS. The fact that an LBAR can be set to an illegal value has no operational impact.

Resolution: Do not use LBAR values of 10000h or higher.

70. PIC fails if processed interrupt still high when clock goes off

Description: This issue applies when a PIC interrupt input is configured for "edge" operation and the interrupt source remains high long after the interrupt has been serviced. For example, PIT Timer 0 is connected to PIC Interrupt 0. This timer is used as a continuously running interval or periodic timer. The timer is normally programmed to produce a square wave signal with at least a 50% high duty cycle and frequency less than 250 Hz. The low-to-high edge of this signal generates the periodic interrupt. This means the interrupt source remains high hundreds of ms after the interrupting edge. This is more than enough time for the interrupt to be generated and serviced. By itself, this is not a problem.

The Diverse Device (DD) features automatic hardware clock gating on the Local bus clock. If there is no other activity in the DD, the Local bus clock automatically turns off after a given interrupt has been generated and serviced. If the interrupt source of an edge-configured PIC input remains high after the clock goes off, the PIC can not return to an Idle state. The Idle state is associated with the interrupt input being low. From Idle, any low-to-high edge produces a new interrupt. Because the PIC can not return to Idle, any new low-to-high edges are not recognized as interrupts. Since the interrupt is used to turn the clock back on, the clock never comes back on. Thus, the system potentially hangs due to the lack of an interrupt.

Unrelated activity elsewhere in the DD could turn on the Local bus clock while the interrupt source is low and allow the PIC to return to Idle. However, if a source, such as in the PIT for example, goes low and then back high while the clock is off, the PIC will not recover even if the Local bus clock comes back on due to an unrelated activity. A PIC edge-configured input can only return to Idle if the clock is currently on or goes on while the source input is low.

Implications: Interrupts may be missed. If an interrupt is missed, then the system may hang or exhibit other unexpected behavior.

Resolution: Do not turn on automatic clock gating for the DD Local bus.

75. ATAC PRD may prefetch incorrectly

Description: The recommended setting for the ATAC GLD_MSR_CONFIG (MSR 51300001h) is 00000000_0048F000h. This setting allows the ATAC master to determine the GeodeLink Adaptor (GA) prefetch setting when it does a Physical Region Descriptor (PRD) access. The master requests a prefetch value of two (010 binary) for PRD requests. This is a prefetch of eight bytes. If the PRD table is 8-byte aligned in main memory, this is OK. If not aligned, this setting will result in the first four bytes of the next PRD to be prefetched. This could introduce an error.

Implications: The error “exposure” of this issue is as follows:

1) UDMA write (memory read): None. The data fetches cause the prefetch to be discarded.

2) UDMA read (memory write):

a) 8-byte aligned PRD tables. None. GA prefetches just the right amount of data.

b) 4-byte aligned PRD tables. GA prefetches first four bytes of next PRD. This is a hazard only if the software is changing PRD entries associated with UDMA currently running. With well-behaved software, this does not happen because all PRD entries should be defined before the UDMA operation is started.

Note that any new UDMA operation involves a write to the ATAC registers. Any I/O write discards all prefetched data unconditionally.

Resolution: Insure PRD tables are defined before starting a UDMA operation.

76. ACC PRD may prefetch incorrectly

Description: The recommended setting for the ACC GLD_MSR_CONFIG (MSR 51500001h) is 00000000_0008F000h. This setting allows the ACC master to determine the GeodeLink Adaptor (GA) prefetch setting when it does a Physical Region Descriptor (PRD) access. The master requests a prefetch value of two (010 binary) for PRD requests. This is a prefetch of eight bytes. If the PRD table is 8-byte aligned in main memory, this is OK. If not aligned, this setting results in the first four bytes of the next PRD to be prefetched. This could introduce an error.

Implications: The error “exposure” of this issue is as follows:

1) Audio playback (memory read): None. The data fetches cause the prefetch to be discarded.

2) Audio record (memory write):

a) 8-byte aligned PRD tables. None. GA prefetches just the right amount of data.

b) 4-byte aligned PRD tables. GA will prefetch first four bytes of next PRD. This is a hazard only if the software is changing the next PRD entry associated with the current PRD entry.

Note that any new audio operation involves a write to the native registers. Any I/O write discards all prefetched data unconditionally.

Resolution: Insure PRD table changes being made on the fly are at least two tables ahead of the current PRD.

82. UDMA write CRC corruption after device termination

Background: This issue manifests itself within the context of a UDMA write operation but is actually due to a corrupted state residue (host_pio_term set) from a previous UDMA read operation. Within the context of the UDMA read operation that introduces this corrupted state, there is no error. Additionally, the UDMA read does not introduce the corrupted state unless there has been a host termination near the end of the operation. In summary, the issue detailed does not occur unless there has been a preceding UDMA read operation with a host termination.

Description: During a UDMA write to an IDE device, it is possible for the IDE device to terminate the transaction. When this happens, the ATAC passes a CRC to the device for the portion of the overall UDMA completed. A new CRC is then generated for the next portion of the UDMA write.

For the first device termination of the over all UDMA operation, the ATAC passes the correct CRC. However, the ATAC does not, in general, start a new CRC for the next portion, but continues from the last. Thus, the CRC for the next portion will be in error.

For the record, if there happens to be a host initiated IDE device I/O before the UDMA resumption, then the CRC is restarted correctly. Specifically, a host termination after a device termination will cause a new CRC to be started. While this cures the problem, this rare set of events can in no way be depended upon to cure this problem.

Implications: In general UDMA writes can not be used due to the CRC errors introduced.

Resolution: The root cause of this issue is the corruption that occurs due to a host termination in a proceeding UDMA read. Use the UDMA I/O blocking feature associated with bit 8 of the ATAC Error MSR (MSR 51300003h) to prevent host terminations.

84. ATAC SSMI during DMA does not work correctly

Description: This issue concerns the ATAC GLD_MSR_SMI (MSR 51300002h). If bit 0 (IDE_PIO_SSMI_EN) is set in this register and an IDE I/O operation is attempted during an IDE master DMA operation, then the I/O should be blocked and an SSMI generated. An SSMI is generated by setting an SSMI bit in the response packet. Furthermore, bit 32 (IDE_PIO_SSMI_FLAG) is set in this same MSR to indicate the SSMI occurred. This applies regardless of read or write operation.

The operation works as detailed above, except the I/O to the IDE is not blocked.

Implications: This feature is not required operationally. It is only useful for debug and workarounds. There is no operational impact of this issue.

Resolution: Do not use this feature.

85. Main GLIU clock domain does not awaken when statistics counter generates an ASMI

Description: When the GLIU statistics counters are in use, they generate ASMI's under certain conditions. These counter functions work correctly. The primary GLIU logic and these counters operate on different clock domains and hence have separate controls in the GLIU GLD_MSR_PM (MSR 51010004h). They are domains zero and one respectively. If hardware clock gating is enabled for domain zero and a statistics counter generates an ASMI, the domain zero clock should wakeup to pass on the ASMI, but it does not.

Implications: If the statistics counters are used in a mode that generates an ASMI, then any given ASMI could be lost if hardware clock gating is enabled for domain zero.

Resolution: If the statistics counters are used in a mode that generates an ASMI, then leave hardware clock gating off for domain zero.

89. Byte accesses to ACPI PM1_STS and PM1_EN (ACPI I/O Offset 00h and 02h) and hang

Description: Byte reads and writes to ACPI registers PM1_EN and PM1_STS do not complete. The current implementation (rev A3) supports only WORD and DWORD reads/writes to PM1_EN and PM1_STS.

Nature of the failure: PM1_STS is located at ACPI I/O Offset 00h and PM1_EN is located at ACPI I/O Offset 02h from the DIVIL MSR_LBAR ACPI (MSR 5140000Eh). They are located on a 32-bit bus so PM1_STS is associated with lanes 1 and 0 while PM1_EN is associated with lanes 3 and 2.

The following access rules apply:

1) PM1_STS access requires byte lanes 1 and 0 to be active. If only one lane is active, the access will hang.

2) PM1_EN access requires byte lanes 3 and 2 to be active. If only one lane is active, the access will hang.

3) Alternatively, if any 3 or 4 lanes are active, PM1_STS and PM1_EN will be accessed simultaneously as 32-bit operation. If read, both registers are read. If write, only the register(s) associated with (1&0) and/or (3&2) will be written. Regardless, the access will not hang.

4) Address bits[1:0] are “don't care” with regards to the PMC.

5) Any BYTE, WORD, or DWORD access to ACPI I/O Offset 04h hits reserved space and will not hang.

Any access disregarding the rules above will hang.

Implications: The PM1_STS and PM1_EN registers are 16 bits, so there is no “natural” need for byte access. However, the ACPI specification allows it. Therefore, byte access is possible but has not been observed in application software.

Resolution: Do not use byte access to these registers. Alternatively, use AMD BIOS that traps bytes and performs the intended function via system management emulation software.

90. WORD and BYTE writes to PMC registers may corrupt other PMC register writes

Description: The PMC operational clock is different and slower than the Diverse Device (DD) Local bus clock. A shadow register technique is used to hold write data from DD I/O space temporarily before it is moved to the operational space. The shadow register is updated based on the byte enables associated with the 32-bit write. However, when the operational register is updated with the shadow register contents, the byte lane information is not used. As a result, the register is updated at 32-bit granularity irrespective of the write size.

Additionally, the shadow register for the inactive byte lane retains the previous write data. Thus, a corruption to a future write is possible. There are two categories of corruption: local and cross. Local corruption can only occur to the same register on a future write. Cross corruption can occur to an unrelated register. All the PMC registers are listed below by category.

Local	Address
PM1_STS	ACPI I/O Offset 00h
PM1_EN	ACPI I/O Offset 02h
GPE0_STS	ACPI I/O Offset 18h
GPE0_EN	ACPI I/O Offset 1Ch
PM_WKD	PMS I/O Offset 30h
PM_WKXD	PMS I/O Offset 34h
PM_RD	PMS I/O Offset 38h
PM_WKXA	PMS I/O Offset 3Ch
PM_FSD	PMS I/O Offset 40h
PM_TSD	PMS I/O Offset 44h
PM_PSD	PMS I/O Offset 48h
PM_NWKD	PMS I/O Offset 4Ch
PM_AWKD	PMS I/O Offset 50h
PM_SSC	PMS I/O Offset 54h

Cross	Address
PM1_CNT	ACPI I/O Offset 08h
PM2_CNT	ACPI I/O Offset 0Ch
PM_TMR	ACPI I/O Offset 10h
PM_SSD	PMS I/O Offset 00h
PM_SCXA	PMS I/O Offset 04h
PM_SCYA	PMS I/O Offset 08h
PM_SODA	PMS I/O Offset 0Ch
PM_SCLK	PMS I/O Offset 10h
PM_SED	PMS I/O Offset 14h
PM_SCXD	PMS I/O Offset 18h
PM_SCYD	PMS I/O Offset 1Ch
PM_SIDD	PMS I/O Offset 20h

Implications: The ACPI specification allows WORD and BYTE accesses to the ACPI defined registers PM1_STS, PM1_EN, PM1_CNT, and PM2_CNT. However, byte accesses result in the register corruption described above. The ACPI specification allows DWORD, WORD, and BYTE accesses to ACPI defined registers GPE0_STS and GPE0_EN. However, WORD and BYTE accesses result in the register corruption described above. For the proprietary non-ACPI registers, the CS5535 specification restricts accesses to DWORD only.

Resolution: An SMI I/O trap can be setup for the ACPI defined registers. This trap can control accesses to the specified registers and prevent an incorrectly sized access.

92. KEL HceStatus[3] does not report writes to I/O Port 064h

Description: When KEL emulation logic is enabled (HCE_Control EmulationEnable bit is set) (KEL Memory Offset 100h[0] = 1), writes to I/O Ports 060h and 064h should set InputFull in HCE_Status (KEL Memory Offset 10Ch[1] = 1). This works correctly. Additionally, CmdData (KEL Memory Offset 10Ch[3]) should be 0 on writes to I/O Port 060h and 1 on writes to I/O Port 064h. However, CmdData is always 0.

Implications: Emulation software can not determine the difference between writes to I/O Ports 060h and 064h using the KEL.

Resolution: Emulation Events (EEs) are processed with ASMLs. At the processor level, an ASMI traps the current instruction. The determination between I/O Ports 060h and 064h is made using information provided by the trap.

93. KEL Snoop corrupts HCE state

Description: Under Snoop mode (see KEL MSR 5140001Fh[0]), writes to I/O Port 060h and 064h should never change the state of any KEL HCE register. However, writes have the same effect on register state as in Emulation mode (KEL Memory Offset 100h[0] = 1).

Implications: Snoop mode is not usable without system software work-arounds.

Resolution: System software must track the use of Snoop mode. On entry, system software must store the HCE state. On exit from Snoop mode, system software must restore the HCE state.

94. Possible PCI memory read failure on back-to-back fetches

Description: This issue applies to memory reads made from the PCI bus to the CS5535 as a slave. If single 32-bit back-to-back reads to consecutive addresses in the same cache line are made one clock cycle apart, then the second read does not address the correct location.

This issue was discovered in simulation space using artificial address generators. Back-to-back memory reads can not be issued to the CS5535 from the Geode GX processor under real world environments.

Implications: No impact on real world systems.

Resolution: None.

96. PIC does not support Polling mode

Description: The standard 8259 PIC (Programmable Interrupt Controller) has a Polling mode that is not supported by the CS5535. Specifically, the OCW3 bit 2 (I/O port 20h/A0h [2]) is marked reserved, indicating that the Polling mode is not supported.

Implications: This mode is not normally used in x86 systems.

Resolution: None.

97. Automatic clock gating for UART clock domain does not work in CEIR mode

Description: In CEIR mode with clock gating on (i.e., MSR 51400004h[15:14] and [13:12] = 01), the UART clock becomes inactive when the RX_FIFO is empty and the incoming data is a long stream of 1s even though the RXACT bit (UART I/O Offset 07h[5]) is set. This affects both UART1 and UART2.

Implications: CEIR mode will not work correctly with clock gating enabled.

Resolution: Do not turn on automatic clock gating for these domains. Device power will therefore be slightly higher than nominal. Note that wakeup from this device is still possible even with clock gating disabled.

98. IRDY# and DEVSEL# do not meet t_{VAL} minimum timing

Description: IRDY# and DEVSEL# do not meet the t_{VAL} minimum value of 2 ns with V_{CORE} = 1.575V at 0°C as described in Table 6-8. "PCI, SUSP#, SUSPA#, and RESET_OUT# Timing Parameters" of the AMD Geode™ CS5535 Companion Device Data Book. The guaranteed t_{VAL} for IRDY# and DEVSEL# is 1.5 ns.

Implications: IRDY# and DEVSEL# violate the PCI timing specification. No system failures are expected since there is a sufficient Hold time margin of 0 ns for the devices receiving the signals as stated in the PCI specification.

Resolution: If a problem does arise, place small capacitors (10 to 50 pF) on the failing signal to slow down the edge rate by 500 ps to bring the signal back into spec. There is sufficient maximum t_{VAL} margin to support the slower edge rate.

99. SSMI may not be synchronous if IRQ occurs at the same time

Description: An SSMI event can be enabled on specific I/O addresses. The purpose of the SSMI is to trap the I/O and let a VSM (Virtual System Module) alter the contents if the I/O is a read, or change or block the data if the I/O is a write. The code immediately after the I/O does not execute until the VSM has executed.

If the CS5535 receives an interrupt request event and SSMI event at the same time, it may not signal the SMI event to the CPU Core through the serial packet mechanism before the completion of the I/O that is signaling the SSMI event. This means that the I/O will be allowed to complete without causing the SSMI to be taken. The SSMI will eventually be signaled, but the I/O that caused it has long since been executed.

Implications: The CS5535 can not generate SSMIs reliably.

Resolution: The Geode GX processor has the same capabilities. Use the processor to create all desired SSMIs.

100. PIO access to CD-ROM drive with UDMA bus master active fails

Description: Normal UDMA procedure is to send the command packet to the IDE device in PIO mode, followed by enabling the IDE bus master in the CS5535 to complete the desired transfer. While it is not normal, it is acceptable to enable the IDE bus master and then send the command packet to the IDE device in PIO mode. In the CS5535, if the bus master is enabled first, the UDMA transaction will fail. However, it will not result in a hang. The bus master can be reset. This reverse order programming was observed in a CD-ROM driver in an older OS (Microsoft Windows 98). In the Windows 98 failure, the driver, upon detecting the transfer failure, reverts to PIO mode for all future transactions. This reverse operation has not been observed in any other drivers (CD-ROM or HD) in any other OS.

Implications: UDMA operation, where the CS5535 bus master is enabled followed by a PIO command packet, will fail. Normal expected operation, PIO command packet followed by bus master enable, is the required procedure for successful UDMA operation.

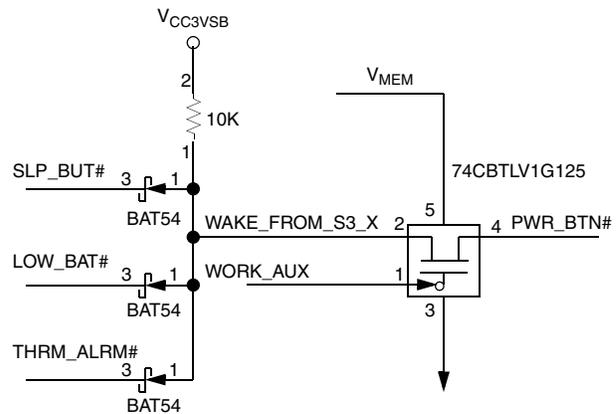
Resolution: Use the normal UDMA procedure, as defined above, in all IDE drivers that do UDMA.

101. Most wake event inputs do not work from the ACPI S3 state

Description: When in the ACPI S3 state (Save-to-RAM), transitions on GPIO[27:24] (SLP_BUT#, THRM_ALRM#, and LOW_BAT#) are not recognized. When enabled, transitions on the above inputs are supposed to wake the system from the S3 state.

Implications: GPIO[27:24] (SLP_BUT#, THRM_ALRM#, and LOW_BAT#) inputs can not be used as wakeup events when in ACPI S3 state.

Resolution: Use the PWRBTN# input. If it is necessary to distinguish multiple wake events, external circuitry can be used to OR them together. The external circuitry must be powered by VIO_STB and V_{MEM} as shown in the figure below. The individual inputs should still be connected to the appropriate CS5535 inputs for software to determine the source.



Wake from S3 Workaround

102. RTC CMOS RAM accesses incorrectly

Description: There are two RAM locations: legacy RAM accessed through I/O 70h/71h (address/data), and extended accessed through I/O 72h/73h (address/data). If data is written into the extended data port (I/O 73h) followed immediately by a read of the legacy data port (I/O 71h), the data read from the legacy data port will be the data written into the extended data port. If a write to an address port is done in between, correct data is then read.

Implications: RTC CMOS RAM cannot be reliably read if writes and reads to the RAM are performed in the above described manner. Risk of failed access is low because the access methodology described above is not normally done.

Resolution: None.

103. I/Os to LPC devices do not work during block and demand mode legacy DMA

Description: Block and demand mode legacy DMA to LPC can block I/O cycles to LPC for extended periods of time. When the I/Os are attempted, they may inadvertently complete without actually accessing the LPC device and return invalid data.

Implications: Block and demand mode legacy DMA can not be reliably used to devices on LPC.

Resolution: Use single transfer mode legacy DMA instead. SMM mode can be used to trap block and demand configurations and convert them to single transfer mode if necessary.

104. Clock gating in the USB Controllers does not work

Description: If clock gating is enabled (MSR 51600004h[5,4] = 01 and MSR 51200004h[5,4] = 01) in the USB Controllers, reads to some registers will fail when the controllers are idle.

Implications: The USB Controllers can not be used reliably if clock gating is enabled.

Resolution: Disable clock gating to the USB Controllers (MSR 51600004h[5,4] = 00 and MSR 51200004h[5,4] = 00).

105. ATAPI packet command after bus master enable to CD-ROM drive fails

Description: If IDE bus mastering has been started and an interrupt is signaled on the IDE interface from an IDE device, the IDE bus master engine in the CS5535 will reset. This will cause the data transfer to stop. Generally the operating system will time out and conclude that bus mastering is not working and operate the IDE interface in PIO mode.

Implications: This soft failure is a result of two unexpected actions. It is not expected behavior for an IDE device to signal an interrupt before the completion of a bus master transfer and it is not expected behavior that signaling an interrupt will reset the bus master engine. The failure was observed in only one IDE device: a CD-ROM drive with a specific firmware revision in the drive. The manufacturer has revised the firmware and with this new revision, the interrupt is no longer signaled unexpectedly. This is a minor issue that is unlikely to affect systems.

Resolution: None.

106. AuxOutputfull bit in KEL incorrectly affected by read from Port 60h

Description: The AuxOutputfull bit in the KEL_HCE_STS register (KEL Memory Offset 10Ch[5]) is incorrectly cleared to 0 on a read from Port 60h. The specification indicates that only a direct write to this bit can change the state of this bit.

Implications: If AuxOutputFull is used to distinguish between mouse and keyboard traffic in a non-interrupt environment, and if reads from Port 60h occur before checking the AuxOutputFull bit, incorrect PS/2 keyboard/mouse emulation on USB can occur. Some test software may utilize this method. However, this is not the usual situation. Generally, systems manage mouse and keyboard with separate interrupts.

Resolution: Use interrupts to manage the USB keyboard and mouse when used in PS/2 emulation.

107. tACK does not affect UDMA timing as expected

Description: The tACK bits in the Channel 0 Drive x DMA registers (MSRs 51300021h [3:0] and 51300023h [3:0]) have no affect on the setup to IDE_DACK timing. It is fixed at two 66 MHz clocks (30 ns).

Implications: No timing violation occurs due to this issue. 30 ns satisfies the setup to IDE_DACK timing requirement in all UDMA modes.

Resolution: None required.

108. ATAC can hang if PRD address is not DWORD aligned

Description: The CS5535 allows for WORD aligned PRD sizes and start addresses. The ATAC can hang if WORD alignment is used.

Implications: The CS5535 allows standard drivers to use word alignment, which could cause ATAC hangs. Extensive testing with the Windows XP standard driver has not produced any failures, however, this is not a guarantee. For the other supported OSs, patches could be required to make sure that WORD aligned PRDs are avoided.

Resolution: Make all ATAC drivers use DWORD alignment.

109. LPC Aborts on Long_Sync or Ready_Sync for I/O writes and reads

Description: Under random conditions the CS5535 companion device's LPC controller may abort a transaction when it observes only two consecutive clocks without a defined (non-reserved) SYNC pattern after the turn-around phase.

Implications: This behavior has only been observed with LPC-to-ISA bridges, which may drive defined (non-reserved) SYNC patterns not earlier than the third clock after the turn-around phase. In that case, write transactions are discarded and read transactions return a value of all ones.

Resolution: A Programmable Logic Device (PLD) can be programmed and connected to the LPC bus in a way that works around the above mentioned issue. See the *AMD Geode™ CS5535/CS5536 Companion Devices LPC-to-ISA Bridge* (publication #33329) application note.

111. NEGACK status bit behaves incorrectly in the SMB controller

Description: When the SMB controller tries to access a non-existing device, it sets the NEGACK bit, SMB I/O Offset 01h[4], to 1 after it detects no acknowledge at the ninth clock. The specification says there are only two ways to clear this status bit: write 1 to this bit or disable the SMB controller. Sometimes this bit is cleared by a read of this register, but not always.

Implications: When the SMB controller tries to access a non-existing device, sometimes the NEGACK bit is not set immediately after the SMB controller detects no acknowledge at the ninth clock. It takes a long time to set this bit. The controller continues to send the non-existing device address before it sets this bit.

Resolution: Write twice to register SMB_STS (SMB I/O Offset 01h) when clearing the NEGACK bit. The first write must clear the NEGACK bit by writing a 1 to SMB_STS[4]. For the second, write a 0x00 to SMB_STS to "mask" the issue in the NEGACK state machine. Since all writable bits in SMB_STS are W1C bits (write-one-to-clear), the second write does not affect any other logic.

112. Retry counter does not distinguish between accesses in GLPCI

Description: The Retry Transaction Limit (MSR 51000010h[51:49]) determines the number of retries the master should attempt for a given transaction. The implementation counts not only its own retries, but also any other retry on the PCI bus.

Implications: The master may stop retrying earlier than expected.

Resolution: Program a higher value for 51000010h[51:49] or disable this functionality by setting to 0.

113. Write to GPIO High Bank Feature Bit Registers failed after suspend

Description: Atomic write transactions to the atomic GPIO High Bank Feature Bit Registers should only affect the bits selected by the atomic write transaction. However, after a suspend, an atomic write transaction to any of the atomic GPIO High Bank Feature Bit Registers affects the selected bit correctly but clears all non-selected bits of the accessed register.

Implications: Unexpected behavior results after the suspend if these registers are written.

Resolution: After a suspend, an atomic write transaction to any of the atomic GPIO High Bank Feature

Bit Registers must be executed as a Read-Modify-Write transaction to all bits of the register.

114. SMBus Hold Time to short

Description: SMB Specification Version 2.0 requires a minimal data hold time T_{HD:DAT} of 300 nanoseconds. For the CS5535, the minimum data hold time T_{HD:DAT} is 280 nanoseconds.

Implications: None expected.

Resolution: None.

115. Missing LOWBAT_FLAG after abnormal standby

Description: Asserting LOW_BAT# in Normal or Restart Standby state causes the system to enter a Faulted Standby state, but the status bit (LOWBAT_FLAG) in the PM_SSC register (PMS I/O Offset 54h[5]) fails to be set.

Implications: After exit from Standby, the previous occurrence of a low-battery condition cannot be detected via the LOWBAT_FLAG bit.

Resolution: None.

116. Sleep_X is not deasserted if delay is larger than PM_SED delay

Description: When the system is transitioning out of the Sleep state, the PM_SCXD, PM_SCYD, and PM_IN_SLPCTL registers' (PMS I/O Offset 18h, 1Ch, and 20h[29:0]) delay counters are reset upon the de-assertion of SUSPA#. In consequence of the premature counter reset, the respective control signals, SLEEP_X (ball C2) and SLEEP_Y (ball J3), are not de-asserted and the PCI/IDE input signals are not re-enabled.

Implications: When the system is transitioning out of the Sleep state, the external circuitry controlled by SLEEP_X and SLEEP_Y remains in the Sleep state and the PCI and IDE input signaling remains non-functional.

Resolution: Program the PM_SCXD, PM_SCYD, and PM_IN_SLPCTL delay counters (PMS I/O Offset 18h, 1Ch, and 20h[29:0]) to values less than the value of the PM_SED delay counter (PMS I/O Offset 14h[29:0]).

© 2006 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.



www.amd.com

TRADEMARKS

AMD, the AMD Arrow logo, and combinations thereof, and Geode and GeodeLink are trademarks of Advanced Micro Devices, Inc.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States and/or other jurisdictions.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.