

Intel® Mobile Celeron™ Processor Specification Update

Release Date: 25 January 1999

Order Number: 244444-001

The Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

mbcl0199e

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

The Specification Update should be publicly available following the last shipment date for a period of time equal to the specific product's warranty period. Hardcopy Specification Updates will be available for one (1) year following End of Life (EOL). Web access will be available for three (3) years following EOL.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com/>

Copyright © Intel Corporation 1999.

* Third-party brands and names are the property of their respective owners.

CONTENTS

REVISION HISTORY	v
PREFACE.....	vi
Intel® Mobile Celeron™	
Processor Specification Update.....	1
GENERAL INFORMATION	3
ERRATA	11
DOCUMENTATION CHANGES	36
SPECIFICATION CLARIFICATIONS	48
SPECIFICATION CHANGES	70

INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

REVISION HISTORY

Date of Revision	Version	Description
January 1999	001	This document is the first Specification Update for the Intel® Mobile Celeron™ processor

INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

PREFACE

This document is an update to the specifications contained in the *Intel® Mobile Celeron™ Processor (BGA) Datasheet* (Order Number#245106-001), the *Intel® Mobile Celeron Processor in Mobile Modcule MMC-1* (Order Number #245101-001), the *Intel® Mobile Celeron Processor in Mobile Modcule MMC-2* (Order Number #245102-001) and the *Intel Architecture Software Developer's Manual, Volumes 1, 2 and 3* (Order Numbers 243190, 243191, and 243192, respectively). It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools. It contains Specification Changes, S-Specs, Errata, Specification Clarifications, and Documentation Changes.

Nomenclature

Specification Changes are modifications to the current published specifications for the Intel® Mobile Celeron™ processor or the Intel® Celeron™ Mobile Module. These changes will be incorporated in the next release of the specifications.

S-Specs are exceptions to the published specifications, and apply only to the units assembled under that s-spec.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

Errata are design defects or errors. Errata may cause the Intel® Mobile Celeron™ processor's or the Intel® Mobile Celeron Module's behavior to deviate from published specifications. Hardware and software designed to be used with any given processor must assume that all errata documented for that processor are present on all devices unless otherwise noted.

Identification Information

The Intel® Mobile Celeron™ processor or the Intel® Celeron Processor Mobile Module can be identified by the following values:

Family ¹	266- , 300- Model 5 ²
0110	0110

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
2. The Model corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.

The Intel® Mobile Celeron™ processor and the Intel Celeron Processor Mobile Module's second level (L2) cache size can be determined by the following register contents:

128-Kbyte Unified L2 Cache ¹	41h
---	-----

NOTE:

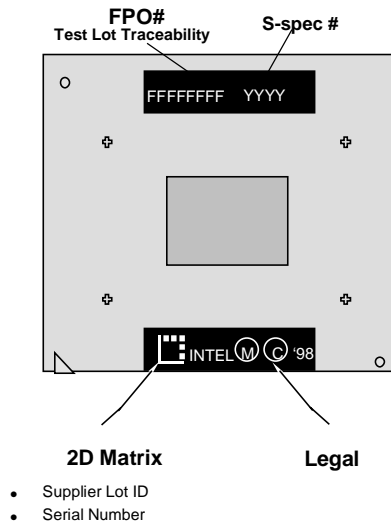
INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

1. For the Intel Mobile Celeron processor and the Intel® Celeron Processor Mobile Module, the unified L2 cache size will be returned as one of the cache/TLB descriptors when the CPUID instruction is executed with a 2 in the EAX register.

Intel® Mobile Celeron™
Processor Specification Update

GENERAL INFORMATION

Intel® Mobile Celeron™ Processor (BGA) Markings



INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Intel® Celeron™ Processor Mobile Module Markings

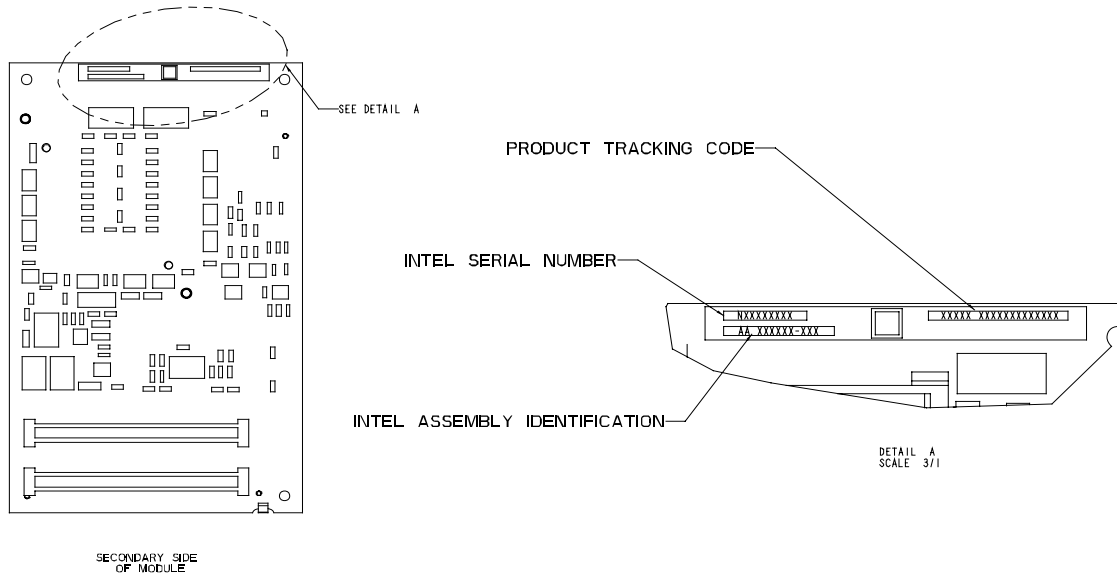
The Product Tracking Code (PTC) determines the Intel assembly level of the module. The PTC is on the secondary side of the module and provides the following information:

Example: **PMG33302001AA**

- The PTC will consist of 13 characters as identified in the above example and can be broken down as follows:

AABCCCDDEEEFF

- Definition:
 - AA - Processor Module = PM
 - B - Celeron Processor Mobile Module = H
 - CCC - Speed Identity = 300 and 266
 - DD - Cache Size = 01 (128KB)
 - EEE - Notifiable Design Revision (Start at 001)
 - FF - Notifiable Processor Revision (Start at AA)
- Note: For other Intel Mobile Modules, the second field (B) is defined as:
 - Pentium® II Processor Mobile Module (MMC-1) at 266/300 MHz = D
 - Pentium® II Processor Mobile Module (MMC-2) at 266/300 MHz = E
 - Pentium® II Processor Mobile Module (MMC-1) at 266PE/300PE/333/366 MHz = F
 - Pentium® II Processor Mobile Module (MMC-2) at 266PE/300PE/333/366 MHz = G
 - Intel® Celeron™ Processor Mobile Module (MMC-2) = I
 - Intel® Celeron™ Processor Mobile Module (MMC-1)**



INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

Intel® Mobile Celeron™ Processor Identification Information

S-Spec	Product Stepping	CPUID	Speed (MHz) Core/Bus	Integrated L2 Size (Kbytes)	Package	Notes
SL23Y	mcbA0	066Ah	266/66	128	BGA1	1
SL3AH	mcbA0	066Ah	300/66	128	BGA1	1

NOTES:

1. VCC_CORE is specified for 1.6V +/-135mV for these Intel® Mobile Celeron™ processors (BGA1 package).

Intel® Celeron™ Processor Mobile Module Identification Information

PTC	Product Stepping	CPUID	Speed (MHz) Core/Bus	Integrated L2 Size (Kbytes)	Package	Notes
PMH26601001AA	cmmA0	066Ah	266/66	128	MMC1	1,2
PMH30001001AA	cmmA0	066Ah	300/66	128	MMC1	1,2
PMI30001001AA	cmmA0	066Ah	300/66	128	MMC2	1,2

NOTES:

1. Vcore voltage is 1.6v
2. Voltage regulator comparator modification.

Summary Table of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the Intel® Mobile Celeron™ processors. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

CODES USED IN SUMMARY TABLE

X:	Specification Change, Erratum, Specification Clarification, or Documentation Change applies to the given processor stepping.
Doc:	Intel intends to update the appropriate documentation in a future revision.
Fix:	This erratum is intended to be fixed in a future stepping of the component.
Fixed:	This erratum has been previously fixed.
NoFix:	There are no plans to fix this erratum.
(No mark) or (blank box):	This item is fixed in or does not apply to the given stepping.
AP:	APIC related erratum.
SUB:	This column refers to errata on the Intel® Mobile Celeron™ processor or the Intel® Celeron™ Processor Mobile Module substrate.
Shaded:	This item is either new or modified from the previous version of the document.

Some of Intel's Specification Updates will be undergoing a numbering methodology change to reduce confusion when referring to errata which affect a specific product. Each Specification Update item will be prefixed with a capital letter to distinguish the product it refers to. The key below details the letters which will be used for the current Intel microprocessor Specification Updates:

A = Pentium® II processor

B = Mobile Pentium II processor

C = Intel® Celeron™ processor

D = Pentium II Xeon™ processor

H = Intel® Mobile Celeron™ processor

The Specification Updates for the Pentium processor, Pentium Pro processor, and other Intel products will not be implementing such a convention at this time.

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

NO.	mcbA0	cmmA0	SUB	Plans	ERRATA
H1	X	X		NoFix	FP Data Operand Pointer may be incorrectly calculated after FP access which wraps 64-Kbyte boundary in 16-bit code
H2	X	X		NoFix	Differences exist in debug exception reporting
H3	X	X		NoFix	Code fetch matching disabled debug register may cause debug exception
H4	X	X		NoFix	Double ECC error on read may result in BINIT#
H5	X	X		NoFix	FP inexact-result exception flag may not be set
H6	X	X		NoFix	BTM for SMI will contain incorrect FROM EIP
H7	X	X		NoFix	I/O restart in SMM may fail after simultaneous MCE
H8	X	X		NoFix	Branch traps do not function if BTMs are also enabled
H9	X	X		NoFix	Machine check exception handler may not always execute successfully
H10	X	X		NoFix	MCE due to L2 parity error gives L1 MCACOD.LL
H11	X	X		NoFix	LBERR may be corrupted after some events
H12	X	X		NoFix	BTMs may be corrupted during simultaneous L1 cache line replacement
H13	X	X		Fix	Potential early deassertion of LOCK# during split-lock cycles
H14	X	X		NoFix	A20M# may be inverted after returning from SMM and Reset
H15	X	X		Fix	Reporting of floating-point exception may be delayed
H16	X	X		NoFix	Near CALL to ESP creates unexpected EIP address
H17	X	X		Fix	Built-in self test always gives nonzero result
H18	X	X		Fix	Cache state corruption in the presence of page A/D-bit setting and snoop traffic
H19	X	X		Fix	Snoop cycle generates spurious machine check exception

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

NO.	mcbA0	cmmA0	SUB	Plans	ERRATA
H20	X	X		Fix	MOVD/MOVQ instruction writes to memory prematurely
F21	X	X		NoFix	Memory type undefined for non-memory operations
H22	X	X		NoFix	FP Data Operand Pointer may not be zero after power on or Reset
H23	X	X		NoFix	MOVD following zeroing instruction can cause incorrect result
H24	X	X		NoFix	Premature execution of a load operation prior to exception handler invocation
H25	X	X		NoFix	Read portion of RMW instruction may execute twice
H26	X	X		Fix	Intervening writeback may occur during locked transaction
H27	X	X		NoFix	MC2_STATUS MSR has model-specific error code and machine check architecture error code reversed
H28	X	X		NoFix	Mixed cacheability of lock variables is problematic in MP systems
H29		X		Fix	Thermal sensor may assert SMBALERT# incorrectly
H30	X	X		NoFix	MOV with debug register causes debug exception
H31	X	X		NoFix	Upper four PAT entries not usable with Mode B or Mode C paging
H32	X	X		Fix	Incorrect memory type may be used when MTRRs are disabled
H33	X	X		Fix	Misprediction in program flow may cause unexpected instruction execution
H34	X	X		NoFix	Data Breakpoint Exception in a Displacement Relative Near Call May Corrupt EIP
H35	X	X		NoFix	System bus ECC not functional with 2:1 ratio
H36	X	X		Fix	Fault on REP CMPS/SCAS operation may cause incorrect EIP
H37	X	X		NoFix	RDMSR and WRMSR to invalid MSR may not cause GP fault

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

NO.	mcbA0	cmmA0	SUB	Plans	ERRATA
H38	X	X		NoFix	SYSENTER/SYSEXIT instructions can implicitly load null segment selector to SS and CS registers
H39	X	X		NoFix	PRELOAD followed by EXTEST does not load boundary scan data
H40	X	X		NoFix	Far jump to new TSS with D-bit cleared may cause system hang
H41	X	X		NoFix	Incorrect chunk ordering may prevent execution of the machine check exception handler after BINIT#
H42	X	X		NoFix	Resume Flag May Not be Cleared After Debug Exception

NO.	mcbA0	cmmA0	SUB	Plans	DOCUMENTATION CHANGES
H1	X	X		Doc	Invalid arithmetic operations and masked responses to them relative to FIST/FISTP instruction
H2	X	X		Doc	FIDIV/FIDIVR m16int description
H3	X	X		Doc	PUSH does not pad with zeros
H4	X	X		Doc	DR7, bit 10 is reserved
H5	X	X		Doc	Additional states that are not automatically saved and restored
H6	X	X		Doc	Cache and TLB description correction
H7	X	X		Doc	SMRAM state save map contains documentation error
H8	X	X		Doc	OF and DF of the EFLAGS register are mislabeled as system flags
H9	X	X		Doc	CS:EIP pushed onto stack prior to code segment limit check
H10	X	X		Doc	Corrections to opcode maps
H11	X	X		Doc	Interrupt 13-general protection exception (#GP)
H12	X	X		Doc	Corrections to <i>Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference</i>
H13	X	X		Doc	MCI_ADDR MSR reference section correction

INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

NO.	mcbA0	cmmA0	SUB	Plans	DOCUMENTATION CHANGES
H14	X	X		Doc	FCOMI/FCOMIP/FUCOMI/FUCOMIP setting of flags relative to exceptions
H15	X	X		Doc	MemTypeGet() function example
H16	X	X		Doc	RSVD flag correction
H17	X	X		Doc	SMBus data setup time
H18	X	X		Doc	SMRAM State Save Map Documentation Correction

NO.	mcbA0	cmmA0	SUB	Plans	SPECIFICATION CLARIFICATIONS
H1	X	X		Doc	Writes to WC memory
H2	X	X		Doc	Multiple processors protocol and restrictions
H3	X	X		Doc	Critical sequence of events during a page fault exception
H4	X	X		Doc	Performance-monitoring counter issues
H5	X	X		Doc	POP[ESP] with 16-bit stack size
H6	X	X		Doc	Preventing caching
H7	X	X		Doc	Paging must be enabled before enabling the page global bit
H8	X	X		Doc	Interrupt recognition determines priority
H9	X	X		Doc	References to 2-Mbyte pages should include 4-Mbyte pages
H10	X	X		Doc	Modification of reserved areas in the SMRAM saved state map
H11	X	X		Doc	TLB flush necessary after PDPE change
H12	X	X		Doc	Exception handler wrong code bit clarification
H13	X	X		Doc	Switching to protected mode while in SMM

NO.	mcbA0	cmdA0	SUB	Plans	SPECIFICATION CHANGES
H1	X	X		Doc	WC buffer eviction data ordering

ERRATA***H1. FP Data Operand Pointer May Be Incorrectly Calculated After FP Access Which Wraps 64-Kbyte Boundary in 16-Bit Code***

PROBLEM: The FP Data Operand Pointer is the effective address of the operand associated with the last noncontrol floating-point instruction executed by the machine. If an 80-bit floating-point access (load or store) occurs in a 16-bit mode other than protected mode (in which case the access will produce a segment limit violation), the memory access wraps a 64-Kbyte boundary, and the floating-point environment is subsequently saved, the value contained in the FP Data Operand Pointer may be incorrect.

IMPLICATION: A 32-bit operating system running 16-bit floating-point code may encounter this erratum, under the following conditions:

- The operating system is using a segment greater than 64 Kbytes in size.
- An application is running in a 16-bit mode other than protected mode.
- An 80-bit floating-point load or store which wraps the 64-Kbyte boundary is executed.
- The operating system performs a floating-point environment store (FSAVE/FNSAVE/FSTENV/FNSTENV) after the above memory access.
- The operating system uses the value contained in the FP Data Operand Pointer.

Wrapping an 80-bit floating-point load around a segment boundary in this way is not a normal programming practice. Intel has not currently identified any software which exhibits this behavior.

WORKAROUND: If the FP Data Operand Pointer is used in an OS which may run 16-bit floating-point code, care must be taken to ensure that no 80-bit floating-point accesses are wrapped around a 64-Kbyte boundary.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H2. Differences Exist in Debug Exception Reporting

PROBLEM: There exist some differences in the reporting of code and data breakpoint matches between that specified by previous Intel processors' specifications and the behavior of the Intel® Mobile Celeron™ processor, as described below:

CASE 1:

The first case is for a breakpoint set on a MOVSS or POPSS instruction, when the instruction following it causes a debug register protection fault (DR7.gd is already set, enabling the fault). The processor reports delayed data breakpoint matches from the MOVSS or POPSS instructions by setting the matching DR6.bi bits, along with the debug register protection fault (DR6.bd). If additional breakpoint faults are matched during the call of the debug fault handler, the processor

sets the breakpoint match bits (DR6.bi) to reflect the breakpoints matched by both the MOVSS or POPSS breakpoint and the debug fault handler call. The Intel® Mobile Celeron™ processor only sets DR6.bd in either situation, and does not set any of the DR6.bi bits.

CASE 2:

In the second breakpoint reporting failure case, if a MOVSS or POPSS instruction with a data breakpoint is followed by a store to memory which crosses a 4- Kbyte page boundary, the breakpoint information for the MOVSS or POPSS will be lost. Previous processors retain this information across such a page split.

CASE 3:

If they occur after a MOVSS or POPSS instruction, the INT *n*, INTO, and INT3 instructions zero the DR6.bi bits (bits B0 through B3), clearing pending breakpoint information, unlike previous processors.

CASE 4:

If a data breakpoint and an SMI (System Management Interrupt) occur simultaneously, the SMI will be serviced via a call to the SMM handler, and the pending breakpoint will be lost.

CASE 5:

When an instruction which accesses a debug register is executed, and a breakpoint is encountered on the instruction, the breakpoint is reported twice.

IMPLICATION: When debugging or when developing debuggers for a Intel® Mobile Celeron™ processor-based system, this behavior should be noted. Normal usage of the MOVSS or POPSS instructions (e.g., following them with a MOV ESP) will not exhibit the behavior of cases 1-3. Debugging in conjunction with SMM will be limited by case 4.

WORKAROUND: Following MOVSS and POPSS instructions with a MOV ESP instruction when using breakpoints will avoid the first three cases of this erratum. No workaround has been identified for cases 4 or 5.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H3. Code Fetch Matching Disabled Debug Register May Cause Debug Exception

PROBLEM: The bits L0-3 and G0-3 enable breakpoints local to a task and global to all tasks, respectively. If one of these bits is set, a breakpoint is enabled, corresponding to the addresses in the debug registers DR0 - DR3. If at least one of these breakpoints is enabled, any of these registers are *disabled* (e.g., *L_n* and *G_n* are 0), and *RW_n* for the disabled register is 00 (indicating a breakpoint on instruction execution), normally an instruction fetch will not cause an instruction-breakpoint fault based on a match with the address in the disabled register(s). However, if the address in a disabled register matches the address of a code fetch which also results in a page fault, an instruction-breakpoint fault will occur.

IMPLICATION: While debugging software, extraneous instruction-breakpoint faults may be encountered if breakpoint registers are not cleared when they are disabled. Debug software which does not implement a code breakpoint handler will fail, if this occurs. If a handler is present, the fault will be serviced. Mixing data and code may exacerbate this problem by allowing disabled data breakpoint registers to break on an instruction fetch.

WORKAROUND: The debug handler should clear breakpoint registers before they become disabled.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H4. Double ECC Error on Read May Result in BINIT#

PROBLEM: For this erratum to occur, the following conditions must be met:

- Machine Check Exceptions (MCEs) must be enabled.
- A dataless transaction (such as a write invalidate) must be occurring simultaneously with a transaction which returns data (a normal read).
- The read data must contain a double-bit uncorrectable ECC error.

If these conditions are met, the Intel® Mobile Celeron™ processor will not be able to determine which transaction was erroneous, and instead of generating an MCE, it will generate a BINIT#.

IMPLICATION: The bus will be reinitialized in this case. However, since a double-bit uncorrectable ECC error occurred on the read, the MCE handler (which is normally reached on a double-bit uncorrectable ECC error for a read) would most likely cause the same BINIT# event.

WORKAROUND: Though the ability to drive BINIT# can be disabled in the Intel® Mobile Celeron™ processor, which would prevent the effects of this erratum, overall system behavior would not improve, since the error which would normally cause a BINIT# would instead cause the machine to shut down. No other workaround has been identified.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H5. FP Inexact-Result Exception Flag May Not Be Set

PROBLEM: When the result of a floating-point operation is not exactly representable in the destination format (1/3 in binary form, for example), an inexact-result (precision) exception occurs. When this occurs, the PE bit (bit 5 of the FPU status word) is normally set by the processor. Under certain rare conditions, this bit may not be set when this rounding occurs. However, other actions taken by the processor (invoking the software exception handler if the exception is unmasked) are not affected. This erratum can only occur if the floating-point operation which causes the precision exception is immediately followed by one of the following instructions:

- FST m32real
- FST m64real
- FSTP m32real
- FSTP m64real
- FSTP m80real
- FIST m16int
- FIST m32int
- FISTP m16int
- FISTP m32int
- FISTP m64int

Note that even if this combination of instructions is encountered, there is also a dependency on the internal pipelining and execution state of both instructions in the processor.

IMPLICATION: Inexact-result exceptions are commonly masked or ignored by applications, as it happens frequently, and produces a rounded result acceptable to most applications. The PE bit of the FPU status word may not always be set upon receiving an inexact-result exception. Thus, if these exceptions are unmasked, a floating-point error exception handler may not recognize that a precision exception occurred. Note that this is a “sticky” bit, e.g., once set by an inexact-result condition, it remains set until cleared by software.

WORKAROUND: This condition can be avoided by inserting two NOP instructions between the two floating-point instructions.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H6. BTM for SMI Will Contain Incorrect FROM EIP

PROBLEM: A system management interrupt (SMI) will produce a Branch Trace Message (BTM), if BTMs are enabled. However, the FROM EIP field of the BTM (used to determine the address of the instruction which was being executed when the SMI was serviced) will not have been updated for the SMI, so the field will report the same FROM EIP as the previous BTM.

IMPLICATION: A BTM which is issued for an SMI will not contain the correct FROM EIP, limiting the usefulness of BTMs for debugging software in conjunction with System Management Mode (SMM).

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H7. I/O Restart in SMM May Fail After Simultaneous MCE

PROBLEM: If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction becomes corrupted, the Intel® Mobile Celeron™ processor will signal a machine check exception (MCE). If the instruction is directed at a device which is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, upon attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is completed successfully, it will attempt to restart the I/O instruction, but will not have the correct machine state, due to the call to the MCE handler.

IMPLICATION: A simultaneous MCE and SMI# assertion may occur for one of the I/O instructions above. The SMM handler may attempt to restart such an I/O instruction, but will have corrupted state due to the MCE handler call, leading to failure of the restart and shutdown of the processor.

WORKAROUND: If a system implementation must support both SMM and MCEs, the first thing the SMM handler code (when an I/O restart is to be performed) should do is check for a pending MCE. If there is an MCE pending, the SMM handler should immediately exit via an RSM instruction and allow the machine check exception handler to execute. If there is not, the SMM handler may proceed with its normal operation.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H8. Branch Traps Do Not Function if BTMs Are Also Enabled

PROBLEM: If branch traps or branch trace messages (BTMs) are enabled alone, both function as expected. However, if both are enabled, only the BTMs will function, and the branch traps will be ignored.

IMPLICATION: The branch traps and branch trace message debugging features cannot be used together.

WORKAROUND: If branch trap functionality is desired, BTMs must be disabled.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H9. Machine Check Exception Handler May Not Always Execute Successfully

PROBLEM: An asynchronous machine check exception (MCE), such as a BINIT# event, which occurs during an access that splits a 4-Kbyte page boundary may leave some internal registers in an indeterminate state. Thus, MCE handler code may not always run successfully if an asynchronous MCE has occurred previously.

IMPLICATION: An MCE may not always result in the successful execution of the MCE handler. However, asynchronous MCEs usually occur upon detection of a catastrophic system condition that would also hang the processor. Leaving MCEs disabled will result in the condition which

caused the asynchronous MCE instead causing the processor to enter shutdown. Therefore, leaving MCEs disabled may not improve overall system behavior.

WORKAROUND: No workaround which would guarantee successful MCE handler execution under this condition has been identified.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H10. MCE Due to L2 Parity Error Gives L1 MCACOD.LL

PROBLEM: If a Cache Reply Parity (CRP) error, Cache Address Parity (CAP) error, or Cache Synchronous Error (CSER) occurs on an access to the Intel® Mobile Celeron™ processor's L2 cache, the resulting Machine Check Architectural Error Code (MCACOD) will be logged with '01' in the LL field. This value indicates an L1 cache error; the value should be '10', indicating an L2 cache error. Note that L2 ECC errors have the correct value of '10' logged.

IMPLICATION: An L2 cache access error, other than an ECC error, will be improperly logged as an L1 cache error in MCACOD.LL.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H11. LBER May Be Corrupted After Some Events

PROBLEM: The last branch record (LBR) and the last branch before exception record (LBER) can be used to determine the source and destination information for previous branches or exceptions. The LBR contains the source and destination addresses for the last branch or exception, and the LBER contains similar information for the last branch taken before the last exception. This information is typically used to determine the location of a branch which leads to execution of code which causes an exception. However, after a catastrophic bus condition which results in an assertion of BINIT# and the reinitialization of the buses, the value in the LBER may be corrupted. Also, after either a CALL which results in a fault or a software interrupt, the LBER and LBR will be updated to the same value, when the LBER should not have been updated.

IMPLICATION: The LBER and LBR registers are used only for debugging purposes. When this erratum occurs, the LBER will not contain reliable address information. The value of LBER should be used with caution when debugging branching code; if the values in the LBR and LBER are the same, then the LBER value is incorrect. Also, the value in the LBER should not be relied upon after a BINIT# event.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H12. BTMs May Be Corrupted During Simultaneous L1 Cache Line Replacement

PROBLEM: When Branch Trace Messages (BTMs) are enabled and such a message is generated, the BTM may be corrupted when issued to the bus by the L1 cache if a new line of data is brought into the L1 data cache simultaneously. Though the new line being stored in the L1 cache is stored correctly, and no corruption occurs in the data, the information in the BTM may be incorrect due to the internal collision of the data line and the BTM.

IMPLICATION: Although BTMs may not be entirely reliable due to this erratum, the conditions necessary for this boundary condition to occur have only been exhibited during focused simulation testing. Intel has currently not observed this erratum in a system level validation environment.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H13. Potential Early Deassertion of LOCK# During Split-Lock Cycles

PROBLEM: During a split-lock cycle there are four bus transactions: 1st ADS# (a partial read), 2nd ADS# (a partial read), 3rd ADS# (a partial write), and the 4th ADS# (a partial write). Due to this erratum, LOCK# may deassert one clock after the 4th ADS# of the split-lock cycle instead of after the 4th RS# assertion corresponding to the 4th ADS# has been sampled. The following sequence of events are required for this erratum to occur:

1. A lock cycle occurs (split or nonsplit).
2. Five more bus transactions (assertion of ADS#) occur.
3. A split-lock cycle occurs and BNR# toggles after the 3rd ADS# (partial write) of the split-lock cycle. This in turn delays the assertion of the 4th ADS# of the split-lock cycle. BNR# toggling at this time could most likely happen when the bus is set for an IOQ depth of 2.

When all of these events occur, LOCK# will be deasserted in the next clock after the 4th ADS# of the split-lock cycle.

IMPLICATION: This may affect chipset logic which monitors the behavior of LOCK# deassertion.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H14. A20M# May Be Inverted After Returning From SMM and Reset

PROBLEM: This erratum is seen when software causes the following events to occur:

1. The assertion of A20M# in real address mode.
2. After entering the 1-Mbyte address wrap-around mode caused by the assertion of A20M#, there is an assertion of SMI# intended to cause a Reset or remove power to the processor. Once in the SMM handler, software saves the SMM state save map to an area of nonvolatile

memory from which it can be restored at some point in the future. Then software asserts RESET# or removes power to the processor.

3. After exiting Reset or completion of power-on, software asserts SMI# again. Once in the SMM handler, it then retrieves the old SMM state save map which was saved in event 2 above and copies it into the current SMM state save map. Software then asserts A20M# and executes the RSM instruction. After exiting the SMM handler, the polarity of A20M# is inverted.

IMPLICATION: If this erratum occurs, A20M# will behave with a polarity opposite from what is expected (e.g., the 1-Mbyte address wrap-around mode is enabled when A20M# is deasserted, and does not occur when A20M# is asserted).

WORKAROUND: Software should save the A20M# signal state in nonvolatile memory before an assertion of RESET# or a power down condition. After coming out of Reset or at power on, SMI# should be asserted again. During the restoration of the old SMM state save map described in event 3 above, the entire map should be restored, except for bit 5 of the byte at offset 7F18h. This bit should retain the value assigned to it when the SMM state save map was created in event 3. The SMM handler should then restore the original value of the A20M# signal.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H15. Reporting of Floating-Point Exception May Be Delayed

PROBLEM: The Intel® Mobile Celeron™ processor normally reports a floating-point exception for an instruction when the next floating-point or MMX™ technology instruction is executed. The assertion of FERR# and/or the INT 16 interrupt corresponding to the exception may be delayed until the floating-point or MMX technology instruction *after* the one which is expected to trigger the exception, if the following conditions are met:

1. A floating-point instruction causes an exception.
2. Before another floating-point or MMX™ technology instruction, any one of the following occurs:
 - a. A subsequent data access occurs to a page which has not been marked as accessed, or
 - b. Data is referenced which crosses a page boundary, or
 - c. A possible page-fault condition is detected which, when resolved, completes without faulting.
3. The instruction causing event 2 above is followed by a MOVQ or MOVD store instruction.

IMPLICATION: This erratum only affects software which operates with floating-point exceptions unmasked. Software which requires floating-point exceptions to be visible on the next floating-point or MMX technology instruction, and which uses floating-point calculations on data which is then used for MMX technology instructions, may see a delay in the reporting of a floating-point

instruction exception in some cases. Note that mixing floating-point and MMX technology instructions in this way is not recommended.

WORKAROUND: Inserting a WAIT or FWAIT instruction (or reading the floating-point status register) between the floating-point instruction and the MOVQ or MOVD instruction will give the expected results. This is already the recommended practice for software.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H16. Near CALL to ESP Creates Unexpected EIP Address

PROBLEM: As documented, the CALL instruction saves procedure linking information in the procedure stack and jumps to the called procedure specified with the destination (target) operand. The target operand specifies the address of the first instruction in the called procedure. This operand can be an immediate value, a general purpose register, or a memory location. When accessing an absolute address indirectly using the stack pointer (ESP) as a base register, the base value used is the value in the ESP register before the instruction executes. However, when accessing an absolute address directly using ESP as the base register, the base value used is the value of ESP *after* the return value is pushed on the stack, not the value in the ESP register *before* the instruction executed.

IMPLICATION: Due to this erratum, the processor may transfer control to an unintended address. Results are unpredictable, depending on the particular application, and can range from no effect to the unexpected termination of the application due to an exception. Intel has observed this erratum only in a focused testing environment. Intel has not observed any commercially available operating system, application, or compiler that makes use of or generates this instruction.

WORKAROUND: If the other seven general purpose registers are unavailable for use, and it is necessary to do a CALL via the ESP register, first push ESP onto the stack, then perform an *indirect* call using ESP (e.g., CALL [ESP]). The saved version of ESP should be popped off the stack after the call returns.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H17. Built-in Self Test Always Gives Nonzero Result

PROBLEM: The Built-in Self Test (BIST) of the Intel® Mobile Celeron™ processor does not give a zero result to indicate a passing test. Regardless of pass or fail status, bit 6 of the BIST result in the EAX register after running BIST is set.

IMPLICATION: Software which relies on a zero result to indicate a passing BIST will indicate BIST failure.

WORKAROUND: Mask bit 6 of the BIST result register when analyzing BIST results.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H18. Cache State Corruption in the Presence of Page A/D-bit Setting and Snoop Traffic

PROBLEM: If an operating system uses the Page Access and/or Dirty bit feature implemented in the Intel architecture and there is a significant amount of snoop traffic on the bus, while the processor is setting the Access and/or Dirty bit the processor may inappropriately change a single L1 cache line to the modified state.

IMPLICATION: The occurrence of this erratum may result in cache incoherency, which may cause parity errors, data corruption (with no parity error), unexpected application or operating system termination, or system hangs.

WORKAROUND: It is possible for BIOS code to contain a workaround for this erratum.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H19. Snoop Cycle Generates Spurious Machine Check Exception

PROBLEM: The processor may incorrectly generate a Machine Check Exception (MCE) when it processes a snoop access that does not hit the L1 data cache. Due to an internal logic error, this type of snoop cycle may still check data parity on undriven data lines. The processor generates a spurious machine check exception as a result of this unnecessary parity check.

IMPLICATION: A spurious machine check exception may result in an unexpected system halt if Machine Check Exception reporting is enabled in the operating system.

WORKAROUND: It is possible for BIOS code to contain a workaround for this erratum. This workaround would fix the erratum; however, the data parity error will still be reported.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H20. MOVD/MOVQ Instruction Writes to Memory Prematurely

PROBLEM: When an instruction encounters a fault, the faulting instruction should not modify any CPU or system state. However, when the MMX technology store instructions MOVD and MOVQ encounter any of the following events, it is possible for the store to be committed to memory even though it should be canceled:

1. If CR0.EM = 1 (Emulation bit), then the store could happen prior to the triggered invalid opcode exception.
2. If the floating-point Top-of-Stack (FP TOS) is not zero, then the store could happen prior to executing the processor assist routine that sets the FP TOS to zero.
3. If there is an unmasked floating-point exception pending, then the store could happen prior to the triggered unmasked floating-point exception.
4. If CR0.TS = 1 (Task Switched bit), then the store could happen prior to the triggered Device Not Available (DNA) exception.

- If the MOVD/MOVQ instruction is restarted after handling any of the above events, then the store will be performed again, overwriting with the expected data. The instruction will not be restarted after event 1. The instruction will definitely be restarted after events 2 and 4. The instruction may or may not be restarted after event 3, depending on the specific exception handler.

IMPLICATION: This erratum causes unpredictable behavior in an application if MOVD/MOVQ instructions are used to manipulate semaphores for multiprocessor synchronization, or if these MMX instructions are used to write to uncacheable memory or memory mapped I/O that has side effects, e.g., graphics devices. This erratum is completely transparent to all applications that do not have these characteristics. When each of the above conditions are analyzed:

1. Setting the CR0.EM bit forces all floating-point/MMX™ instructions to be handled by software emulation. The MOVD/MOVQ instruction, which is an MMX instruction, would be considered an invalid instruction. Operating systems typically terminates the application after getting the expected invalid opcode fault.
2. The FP TOS not equal to 0 case only occurs when the MOVD/MOVQ store is the first MMX instruction in an MMX technology routine and the previous floating-point routine did not clean up the floating-point states properly when it exited. Floating-point routines commonly leave TOS to 0 prior to exiting. For a store to be executed as the first MMX instruction in an MMX technology routine following a floating-point routine, the software would be implementing instruction level intermixing of floating-point and MMX instructions. Intel does not recommend this practice.
3. The unmasked floating-point exception case only occurs if the store is the first MMX technology instruction in an MMX technology routine and the previous floating-point routine exited with an unmasked floating-point exception pending. Again, for a store to be executed as the first MMX instruction in an MMX technology routine following a floating-point routine, the software would be implementing instruction level intermixing of floating-point and MMX instructions. Intel does not recommend this practice.
4. Device Not Available (DNA) exceptions occur naturally when a task switch is made between two tasks that use either floating-point instructions and/or MMX instructions. For this erratum, in the event of the DNA exception, data from the prior task may be temporarily stored to the present task's program state.

WORKAROUND: Do not use MMX instructions to manipulate semaphores for multiprocessor synchronization. Do not use MOVD/MOVQ instructions to write directly to I/O devices if doing so triggers user visible side effects. An OS can prevent old data from being stored to a new task's program state by cleansing the FPU explicitly after every task switch. Follow Intel's recommended programming paradigms in the Intel Architecture Optimization Manual for writing MMX technology programs. Specifically, do not mix floating-point and MMX instructions. When transitioning to new a MMX technology routine, begin with an instruction that does not depend on the prior state of either the MMX technology registers or the floating-point registers, such as a load or PXOR mm0, mm0. Be sure that the FP TOS is clear before using MMX instructions.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H21. Memory Type Undefined for Nonmemory Operations

PROBLEM: The Memory Type field for nonmemory transactions such as I/O and Special Cycles are undefined. Although the Memory Type attribute for nonmemory operations logically should (and usually does) manifest itself as UC, this feature is not designed into the implementation and is therefore inconsistent.

IMPLICATION: Bus agents may decode a non-UC memory type for nonmemory bus transactions.

WORKAROUND: Bus agents must consider transaction type to determine the validity of the Memory Type field for a transaction.

STATUS: For the steppings affected, see the Summary Table of Changes at the beginning of this section.

H22. FP Data Operand Pointer May Not Be Zero After Power On or Reset

PROBLEM: The FP Data Operand Pointer, as specified, should be reset to zero upon power on or Reset by the processor. Due to this erratum, the FP Data Operand Pointer may be nonzero after power on or Reset.

IMPLICATION: Software which uses the FP Data Operand Pointer and count on its value being zero after power on or Reset without first executing a FINIT/FNINIT instruction will use an incorrect value, resulting on incorrect behavior of the software.

WORKAROUND: Software should follow the recommendation in Section 8.2 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide* (Order Number 243192). This recommendation states that if the FPU will be used, software-initialization code should execute an FINIT/FNINIT instruction following a hardware reset. This will correctly clear the FP Data Operand Pointer to zero.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H23. MOVD Following Zeroing Instruction Can Cause Incorrect Result

PROBLEM: An incorrect result may be calculated after the following circumstances occur:

1. A register has been zeroed with either a SUB reg, reg instruction or an XOR reg, reg instruction,
2. A value is moved with sign extension into the same register's lower 16 bits; or a signed integer multiply is performed to the same register's lower 16 bits,
3. This register is then copied to an MMX™ technology register using the MOVD instruction prior to any other operations on the sign-extended value.

Specifically, the sign may be incorrectly extended into bits 16-31 of the MMX technology register. Only the MMX technology register is affected by this erratum. The erratum only occurs when the 3 following steps occur in the order shown. The erratum may occur with up to 40 intervening instructions that do not modify the sign-extended value between steps 2 and 3.

1. XOR EAX, EAX
or SUB EAX, EAX
2. MOVSBX AX, BL
or MOVSBX AX, byte ptr <memory address> or MOVSBX AX, BX
or MOVSBX AX, word ptr <memory address> or IMUL BL (AX implicit, opcode F6 /5)
or IMUL byte ptr <memory address> (AX implicit, opcode F6 /5) or IMUL AX, BX (opcode 0F AF /r)
or IMUL AX, word ptr <memory address> (opcode 0F AF /r) or IMUL AX, BX, 16 (opcode 6B /r ib)
or IMUL AX, word ptr <memory address>, 16 (opcode 6B /r ib) or IMUL AX, 8 (opcode 6B /r ib)
or IMUL AX, BX, 1024 (opcode 69 /r iw)
or IMUL AX, word ptr <memory address>, 1024 (opcode 69 /r iw) or IMUL AX, 1024 (opcode 69 /r iw)
or CBW
3. MOVD MM0, EAX

Note that the values for immediate byte/words are merely representative (i.e., 8, 16, 1024) and that any value in the range for the size may be affected. Also, note that this erratum may occur with “EAX” replaced with any 32-bit general purpose register, and “AX” with the corresponding 16-bit version of that replacement. “BL” or “BX” can be replaced with any 8-bit or 16-bit general purpose register. The CBW and IMUL (opcode F6 /5) instructions are specific to the EAX register only.

In the example, EAX is forced to contain 0 by the XOR or SUB instructions. Since the MOVSBX, IMUL and CBW instructions listed should modify only bits 15:8 of EAX by sign extension, bits 31:16 of EAX should always contain 0. This implies that when MOVD copies EAX to MM0, bits 31:16 of MM0 should also be 0. Under certain scenarios, bits 31:16 of MM0 are not 0, but are replicas of bit 15 (the 16th bit) of AX. This is noticeable when the value in AX after the MOVSBX, IMUL or CBW instruction is negative, i.e., bit 15 of AX is a 1.

When AX is positive (bit 15 of AX is a 0), MOVD will always produce the correct answer. If AX is negative (bit 15 of AX is a 1), MOVD may produce the right answer or the wrong answer depending on the point in time when the MOVD instruction is executed in relation to the MOVSBX, IMUL or CBW instruction.

IMPLICATION: The effect of incorrect execution will vary from unnoticeable, due to the code sequence discarding the incorrect bits, to an application failure. If the MMX technology-enabled application in which MOVD is used to manipulate pixels, it is possible for one or more pixels to

exhibit the wrong color or position momentarily. It is also possible for a computational application that uses the MOVD instruction in the manner described above to produce incorrect data. Note that this data may cause an unexpected page fault or general protection fault.

WORKAROUND: There are two possible workarounds for this erratum:

1. Rather than using the MOVXSX-MOVD or CBW-MOVD pairing to handle one variable at a time, use the sign extension capabilities (PSRAW, etc.) within MMX™ technology for operating on multiple variables. This would result in higher performance as well.
2. Insert another operation that modifies or copies the sign-extended value between the MOVXSX/IMUL/CBW instruction and the MOVD instruction as in the example below:
XOR EAX, EAX (or SUB EAX, EAX)
MOVXSX AX, BL (or other MOVXSX, other IMUL or CBW instruction)
*MOV EAX, EAX
MOVD MM0, EAX

*Note: MOV EAX, EAX is used here as it is fairly generic. Again, EAX can be any 32-bit register.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H24. Premature Execution of a Load Operation Prior to Exception Handler Invocation

PROBLEM: This erratum can occur with any of the following situations:

1. If an instruction that performs a memory load causes a code segment limit violation
2. If a waiting floating-point instruction or MMX™ instruction that performs a memory load has a floating-point exception pending, or
3. If an MMX instruction that performs a memory load and has either CR0.EM =1 (Emulation bit set), or a floating-point Top-of-Stack (FP TOS) not equal to 0, or a DNA exception pending.

If any of the above circumstances occur it is possible that the load portion of the instruction will have executed before the exception handler is entered.

IMPLICATION: In normal code execution where the target of the load operation is to write back memory there is no impact from the load being prematurely executed, nor from the restart and subsequent re-execution of that instruction by the exception handler. If the target of the load is to uncached memory that has a system side-effect, restarting the instruction may cause unexpected system behavior due to the repetition of the side-effect.

WORKAROUND: Code which performs loads from memory that has side-effects can effectively workaround this behavior by using simple integer-based load instructions when accessing side-effect memory and by ensuring that all code is written such that a code segment limit violation cannot occur as a part of reading from side-effect memory.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H25. Read Portion of RMW Instruction May Execute Twice

PROBLEM: When the Intel® Mobile Celeron™ processor executes a read-modify-write (RMW) arithmetic instruction, with memory as the destination, it is possible for a page fault to occur during the execution of the store on the memory operand after the read operation has completed but before the write operation completes.

If the memory targeted for the instruction is UC (uncached), memory will observe the occurrence of the initial load before the page fault handler and again if the instruction is restarted.

IMPLICATION: This erratum has no effect if the memory targeted for the RMW instruction has no side-effects. If, however, the load targets a memory region that has side-effects, multiple occurrences of the initial load may lead to unpredictable system behavior.

WORKAROUND: Hardware and software developers who write device drivers for custom hardware that may have a side-effect style of design should use simple loads and simple stores to transfer data to and from the device. Then, the memory location will simply be read twice with no additional implications.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H26. Intervening Writeback May Occur During Locked Transaction

PROBLEM: During a transaction which has the LOCK# signal asserted (e.g., a locked transaction), there is a potential for an explicit writeback caused by a previous transaction to complete while the bus is locked. The explicit writeback will only be issued by the processor which has locked the bus, and the lock signal will not be deasserted until the locked transaction completes, but the atomicity of a lock may be compromised by this erratum. Note that the explicit writeback is an expected cycle, and no memory ordering violations will occur. This erratum is, however, a violation of the bus lock protocol.

IMPLICATION: A chipset or third-party agent (TPA) which tracks bus transactions in such a way that locked transactions may only consist of a read-write or read-read-write-write locked sequence, with no transactions intervening, may lose synchronization of state due to the intervening explicit writeback. Systems using chipsets or TPAs which can accept the intervening transaction will not be affected.

WORKAROUND: The bus tracking logic of all devices on the system bus should allow for the occurrence of an intervening transaction during a locked transaction.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H27. MC2_STATUS MSR Has Model-Specific Error Code and Machine Check Architecture Error Code Reversed

PROBLEM: The *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, documents that for the MCi_STATUS MSR, bits 15:0 contain the MCA (machine-check architecture) error code field and bits 31:16 contain the model-specific error code field. However, for the MC2_STATUS MSR, these bits have been reversed. For the MC2_STATUS MSR, bits 15:0 contain the model-specific error code field and bits 31:16 contain the MCA error code field.

IMPLICATION: A machine check error may be decoded incorrectly if this erratum on the MC2_STATUS MSR is not taken into account.

WORKAROUND: When decoding the MC2_STATUS MSR, reverse the two error fields.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H28. Mixed Cacheability of Lock Variables Is Problematic in MP Systems

PROBLEM: This errata only affects multiprocessor systems where a lock variable address is marked cacheable in one processor and uncacheable in any others. The processors which have it marked uncacheable may stall indefinitely when accessing the lock variable. The stall is only encountered if:

- One processor has the lock variable cached, and is attempting to execute a cache lock.
- If the processor which has that address cached has it cached in its L2 only.
- Other processors, meanwhile, issue back to back accesses to that same address on the bus.

IMPLICATION: MP systems where all processors either use cache locks or consistent locks to uncacheable space will not encounter this problem. If, however, a lock variable's cacheability varies in different processors, and several processors are all attempting to perform the lock simultaneously, an indefinite stall may be experienced by the processors which have it marked uncacheable in locking the variable (if the conditions above are satisfied). Intel has only encountered this problem in focus testing with artificially generated external events. Intel has not currently identified any commercial software which exhibits this problem.

WORKAROUND: Follow a homogenous model for the memory type range registers (MTRRs), ensuring that all processors have the same cacheability attributes for each region of memory; do not use locks whose memory type is cacheable on one processor, and uncacheable on others. Avoid page table aliasing, which may produce a nonhomogenous memory model.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H29. Thermal Sensor May Assert SMBALERT# Incorrectly

PROBLEM: The Intel® Mobile Celeron™ Processor Mobile Module has a thermal sensor that monitors the processor core's temperature. Please note that desktop systems could have a similar

thermal device. The thermal sensor asserts SMBALERT# if the processor temperature exceeds the temperature limits set in the Alarm Threshold Registers (T_{HIGH} , T_{LOW}). It also sets the corresponding Status Register bits to identify the cause of the interrupt. Figure 1 gives one example of the how the SMBALERT# signal could be used in a system.

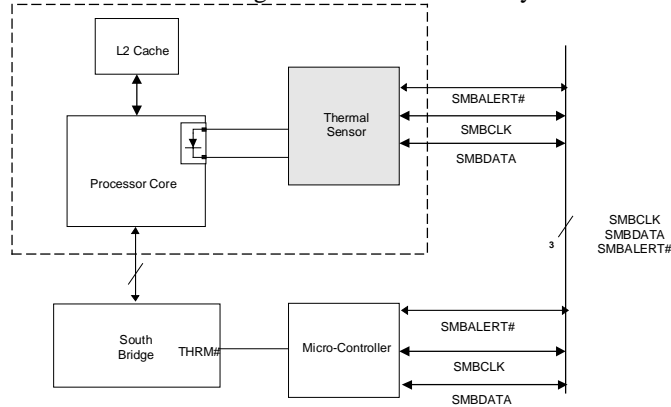


Figure 1. An Example of Microcontroller Driven Thermal Management

Under the conditions described below, the thermal sensor incorrectly generates the SMBALERT# interrupt. *All* of the following conditions must be met to trigger a false interrupt:

1. The thermal sensor must be in auto-convert mode.
2. The absolute value of the difference between the current temperature reading and the T_{HIGH} or T_{LOW} limit value must be less than or equal to 8 °C.
3. The current temperature reading must be different from the previous reading.

With a false assertion of SMBALERT#, the corresponding bit in the Status Register (L_{HIGH} , L_{LOW} , R_{HIGH} , and R_{LOW}) also will be incorrect.

IMPLICATION: There is no system impact from this erratum if temperature polling is used for processor thermal management. If the SMBALERT# interrupt is employed to manage processor thermal sensing, then servicing the false interrupt may result in premature system action depending on the software and hardware implementations used. The rate of the false interrupts is less than the auto-convert rate of the thermal sensor.

WORKAROUND: Three different (mutually exclusive) workarounds are possible:

1. Before servicing an interrupt from the thermal sensor, read and compare the processor thermal reading with the threshold limits (T_{HIGH} or T_{LOW}). Figures 2 and 3 provide basic flowcharts for the implementation of this workaround in an interrupt driven system.
2. If the firmware implemented polls the Status Register only, then before taking any action, re-read the temperature register and do a comparison with the alarm threshold limits (T_{HIGH} or T_{LOW}) to determine if the value is actually still within the temperature window.

- Use a temperature polling scheme to monitor the processor temperature.

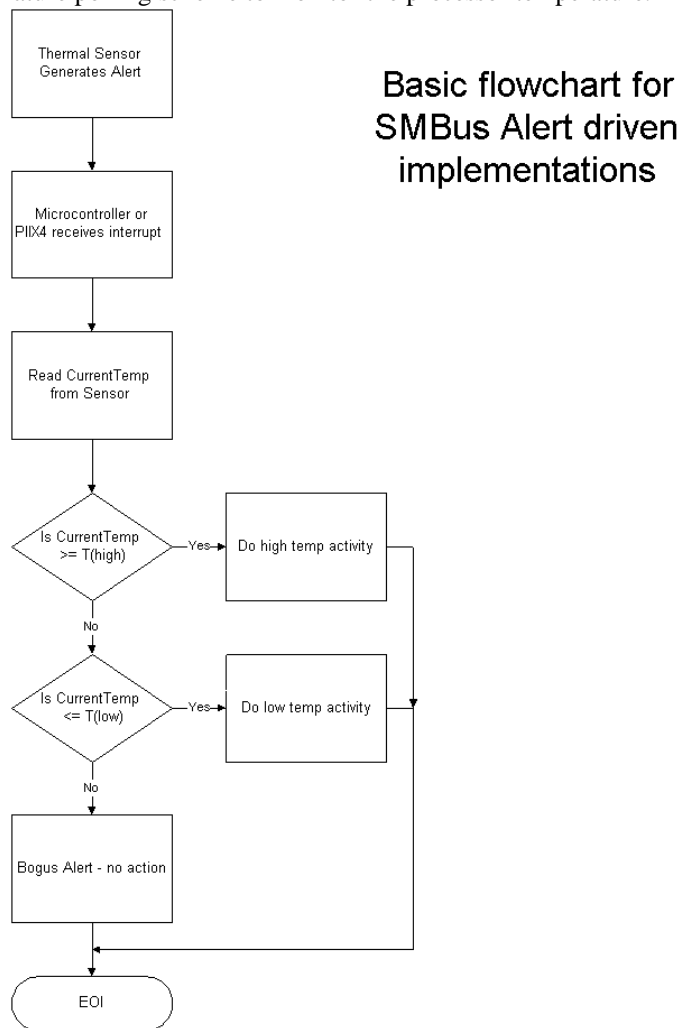


Figure 2. Workaround Flowchart: SMBALERT#-Driven System

Basic flowchart for system interrupt driven implementations

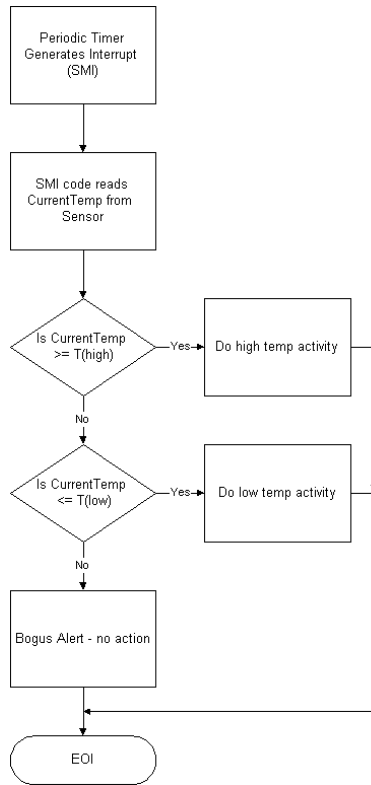


Figure 3. Workaround Flowchart: SMI#-Driven System

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H30. MOV With Debug Register Causes Debug Exception

PROBLEM: When in V86 mode, if a MOV instruction is executed on debug registers, a general-protection exception (#GP) should be generated, as documented in the *Intel Architecture Software Developer’s Manual, Volume 3: System Programming Guide*, Section 14.2. However, in the case when the general detect enable flag (GD) bit is set, the observed behavior is that a debug exception (#DB) is generated instead.

IMPLICATION: With debug-register protection enabled (e.g., the GD bit set), when attempting to execute a MOV on debug registers in V86 mode, a debug exception will be generated instead of the expected general-protection fault.

WORKAROUND: In general, operating systems do not set the GD bit when they are in V86 mode. The GD bit is generally set and used by debuggers. The debug exception handler should check that the exception did not occur in V86 mode before continuing. If the exception did occur in V86 mode, the exception may be directed to the general-protection exception handler.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H31. Upper Four PAT Entries Not Usable With Mode B or Mode C Paging

PROBLEM: The Page Attribute Table (PAT) contains eight entries, which must all be initialized and considered when setting up memory types for the Intel® Mobile Celeron™ processor. However, in Mode B or Mode C paging, the upper four entries do not function correctly for 4-Kbyte pages. Specifically, bit seven of page table entries that translate addresses to 4-Kbyte pages should be used as the upper bit of a three-bit index to determine the PAT entry that specifies the memory type for the page. When Mode B (CR4.PSE = 1) and/or Mode C (CR4.PAE) are enabled, the processor forces this bit to zero when determining the memory type regardless of the value in the page table entry. The upper four entries of the PAT function correctly for 2-Mbyte and 4-Mbyte large pages (specified by bit 12 of the page directory entry for those translations).

IMPLICATION: Only the lower four PAT entries are useful for 4KB translations when Mode B or C paging is used. In Mode A paging (4-Kbyte pages only), all eight entries may be used. All eight entries may be used for large pages in Mode B or C paging.

WORKAROUND: None identified.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H32. Incorrect Memory Type May be Used When MTRRs Are Disabled

PROBLEM: If the Memory Type Range Registers (MTRRs) are disabled without setting the CR0.CD bit to disable caching, and the Page Attribute Table (PAT) entries are left in their default setting, which includes UC- memory type (PCD = 1, PWT = 0; see the *Addendum—Intel Architecture Software Developer’s Manual, Volume 3: System Programming Guide*, for details), data for entries set to UC- will be cached as if the memory type were writeback (WB). Also, if the page tables are set to a memory type other than UC-, then the effective memory type used will be that specified by the page tables and PAT. Any regions of memory normally forced to UC by the MTRRs (such as the VGA video region) may now be incorrectly cached and speculatively accessed.

Even if the CR0.CD bit is correctly set when the MTRRs are disabled and the PAT is left in its default state, then retries and out of order retirement of UC accesses may occur, contrary to the strong ordering expected for these transactions.

IMPLICATION: The occurrence of this erratum may result in the use of incorrect data and unpredictable processor behavior when running with the MTRRs disabled. Interaction between the mouse, cursor, and VGA video display leading to video corruption may occur as a symptom of this erratum as well.

WORKAROUND: Ensure that when the MTRRs are disabled, the CR0.CD bit is set to disable caching. This recommendation is described in *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*. If it is necessary to disable the MTRRs, first clear the PAT register before setting the CR0.CD bit, flushing the caches, and disabling the MTRRs to ensure that UC memory type is always returned and strong ordering is maintained.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H33. Misprediction in Program Flow May Cause Unexpected Instruction Execution

PROBLEM: To optimize performance through dynamic execution technology, the P6 architecture has the ability to predict program flow. In the event of a misprediction, the processor will normally clear the incorrect prediction, adjust the EIP to the correct location, and flush out any instructions it may have fetched from the misprediction. In circumstances where a branch misprediction occurs, the correct target of the branch has already been opportunistically fetched into the streaming buffers, and the L2 cycle caused by the evicted cache line is retried by the L2 cache, the processor may fail to flush out the retirement unit before the speculative program flow is committed to a permanent state.

IMPLICATION: The results of this erratum may range from no effect to unpredictable application or OS failure. Manifestations of this failure may result in:

- Unexpected values in EIP,
- Faults or traps (e.g., page faults) on instructions that do not normally cause faults,
- Faults in the middle of instructions, or
- Unexplained values in registers/memory at the correct EIP.

WORKAROUND: It is possible for BIOS code to contain a workaround for this erratum.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H34. Data Breakpoint Exception in a Displacement Relative Near Call May Corrupt EIP

PROBLEM: If a data breakpoint is programmed at the memory location where the stack push of a near call is performed, the processor will update the stack and ESP appropriately, but may skip the code at the destination of the call. Hence, program execution will continue with the next instruction immediately following the call, instead of the target of the call.

IMPLICATION: The failure mechanism for this erratum is that the call would not be taken; therefore, instructions in the called subroutine would not be executed. As a result, any code relying on the execution of the subroutine will behave unpredictably.

WORKAROUND: Do not program a data breakpoint exception on the stack where the push for the near call is performed.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H35. System Bus ECC Not Functional With 2:1 Ratio

PROBLEM: If a processor is underclocked at a core frequency to system bus frequency ratio of 2:1 and system bus ECC is enabled, the system bus ECC detection and correction will negatively affect internal timing dependencies.

IMPLICATION: If system bus ECC is enabled, and the processor is underclocked at a 2:1 ratio, the system may behave unpredictably due to these timing dependencies.

WORKAROUND: All bus agents that support system bus ECC must disable it when a 2:1 ratio is used.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H36. Fault on REP CMPS/SCAS Operation May Cause Incorrect EIP

PROBLEM: If either a General Protection Fault, Alignment Check Fault or Machine Check Exception occur during the first iteration of a REP CMPS or a REP SCAS instruction, an incorrect EIP may be pushed onto the stack of the event handler if all the following conditions are true:

- The event occurs on the initial load performed by the instruction(s),
- The condition of the zero flag before the repeat instruction happens to be opposite of the repeat condition (e.g., REP/REPE/REPZ CMPS/SCAS with ZF = 0 or RENE/REPNZ CMPS/SCAS with ZF = 1), and
- The faulting micro-op and a particular micro-op of the REP instruction are retired in the retirement unit in a specific sequence.

The EIP will point to the instruction following the REP CMPS/SCAS instead of pointing to the faulting instruction.

IMPLICATION: The result of the incorrect EIP may range from no effect to unexpected application/OS behavior.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H37. RDMSR or WRMSR To Invalid MSR Address May Not Cause GP Fault

PROBLEM: The RDMSR and WRMSR instructions allow reading or writing of MSRs (Model Specific Registers) based on the index number placed in ECX. The processor should reject access to any reserved or unimplemented MSRs by generating #GP(0). However, there are some invalid MSR addresses for which the processor will not generate #GP(0).

IMPLICATION: For RDMSR, undefined values will be read into EDX:EAX. For WRMSR, undefined processor behavior may result.

WORKAROUND: Do not use invalid MSR addresses with RDMSR or WRMSR.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H38. SYSENTER/SYSEXIT Instructions Can Implicitly Load “Null Segment Selector” to SS and CS Registers

PROBLEM: According to the processor specification, attempting to load a null segment selector into the CS and SS segment registers should generate a General Protection Fault (#GP). Although loading a null segment selector to the other segment registers is allowed, the processor will generate an exception when the segment register holding a null selector is used to access memory. However, the SYSENTER instruction can implicitly load a null value to the SS segment selector. This can occur if the value in SYSENTER_CS_MSR is between FFF8h and FFFBh when the SYSENTER instruction is executed. This behavior is part of the SYSENTER/SYSEXIT instruction definition; the content of the SYSTEM_CS_MSR is always incremented by 8 before it is loaded into the SS. This operation will set the null bit in the segment selector if a null result is generated, but it does not generate a #GP on the SYSENTER instruction itself. An exception will be generated as expected when the SS register is used to access memory, however.

The SYSEXIT instruction will also exhibit this behavior for both CS and SS when executed with the value in SYSENTER_CS_MSR between FFF0h and FFF3h, or between FFE8h and FFEbh, inclusive.

IMPLICATION: These instructions are intended for operating system use. If this erratum occurs (and the OS does not ensure that the processor never has a null segment selector in the SS or CS segment registers), the processor’s behavior may become unpredictable, possibly resulting in system failure.

WORKAROUND: Do not initialize the SYSTEM_CS_MSR with the values between FFF8h and FFFBh, FFF0h and FFF3h, or FFE8h and FFEbh before executing SYSENTER or SYSEXIT.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H39. PRELOAD Followed by EXTEST Does Not Load Boundary Scan Data

PROBLEM: According to the IEEE 1149.1 Standard, the EXTEST instruction would use data “typically loaded onto the latched parallel outputs of boundary-scan shift-register stages using the

SAMPLE/PRELOAD instruction prior to the selection of the EXTEST instruction.” As a result of this erratum, this method cannot be used to load the data onto the outputs.

IMPLICATION: Using the PRELOAD instruction prior to the EXTEST instruction will not produce expected data after the completion of EXTEST.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H40. Far Jump to New TSS With D-bit Cleared May Cause System Hang

PROBLEM: A task switch may be performed by executing a far jump through a task gate or to a new Task State Segment (TSS) directly. Normally, when such a jump to a new TSS occurs, the D-bit (which indicates that the page referenced by a Page Table Entry (PTE) has been modified) for the PTE which maps the location of the previous TSS will already be set and the processor will operate as expected. However, if the D-bit is clear at the time of the jump to the new TSS, the processor will hang.

IMPLICATION: If an OS is used which can clear the D-bit for system pages, and which jumps to a new TSS on a task switch, then a condition may occur which results in a system hang. Intel has not identified any commercial software which may encounter this condition; this erratum was discovered in a focused testing environment.

WORKAROUND: Ensure that OS code does not clear the D-bit for system pages (including any pages that contain a task gate or TSS). Use task gates rather than jumping to a new TSS when performing a task switch.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H41. Incorrect Chunk Ordering May Prevent Execution of the Machine Check Exception Handler After BINIT#

PROBLEM: If a catastrophic bus error is detected which results in a BINIT# assertion, and the BINIT# assertion is propagated to the processor’s L2 cache at the same time that data is being sent to the processor, then the data may become corrupted in the processor’s L1 cache.

IMPLICATION: Since BINIT# assertion is due to a catastrophic event on the bus, the corrupted data will not be used. However, it may prevent the processor from executing the Machine Check Exception (MCE) handler, causing the system to hang.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

H42. Resume Flag May Not Be Cleared After Debug Exception

PROBLEM: The Resume Flag (RF) is normally cleared by the processor after executing an instruction which causes a debug exception (#DB). In the process of determining whether the RF needs to be cleared after executing the instruction, the processor uses an internal register containing stale data. The stale data may unpredictably prevent the processor from clearing the RF.

IMPLICATION: If this erratum occurs, further debug exceptions will be disabled.

WORKAROUND: None identified at this time.

STATUS: For the steppings affected see the Summary Table of Changes at the beginning of this section.

DOCUMENTATION CHANGES

The Documentation Changes listed in this section apply to the *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3*. All Documentation Changes will be incorporated into a future version of the appropriate Intel® Mobile Celeron™ processor documentation.

Note: The *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3* applies to all P6 family processors, and therefore some of the Documentation Changes in this section may not pertain to the Intel Mobile Celeron processor specifically.

H1. Invalid Arithmetic Operations and Masked Responses to Them Relative to FIST/FISTP Instruction

The *Intel Architecture Software Developer's Manual, Volume 1: Basic Architecture*, Table 7-20 show “Invalid Arithmetic Operations and the Masked Responses to Them.” The table entry corresponding to the FIST/FISTP condition is missing, and is shown below:

Condition	Masked Response
FIST/FISTP instruction when input operand $\langle \rangle$ MAXINT for destination operand size.	Return MAXNEG to destination operand.

When FIST/FISTP instruction is executed with input operand $\langle \rangle$ and the destination operand size is MAXINT, the floating-point zero-divide exception will return MAXNEG to the destination operand as its masked response.

H2. FIDIV/FIDIVR m16int Description

The *Intel Architecture Software Developer's Manual, Volume 1: Basic Architecture*, pages 3-118 and 3-122, show in the Description column for the FIDIV *m16int* instruction as “Divide ST(0) by *m64int* by ST(0) and store the result in ST(0)” and FIDIVR *m16int* instruction as “Divide *m64int* by ST(0) and store the result in ST(0).” In both of these cases, *m64int* should be replaced with *m16int*.

H3. PUSH Does Not Pad With Zeros

The *Intel Architecture Software Developer's Manual, Volume 1: Basic Architecture*, page 4-3, contain a section regarding stack alignment. The last sentence in the first paragraph of this section, reads “If a 16-bit value is pushed onto a 32-bit wide stack, the value is automatically padded with zeros out to 32-bits.” This sentence should be removed. The PUSH instruction does not pad with zeros.

H4. DR7, Bit 10 is Reserved

The *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, shows Figure 14-1, "Debug Registers." Bit 10 of DR7 should be "Reserved" instead of "1".

H5. Additional States That Are Not Automatically Saved and Restored

In Section 11.4.1 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, the end of section lists the registers that are not automatically saved and restored following an SMI and the RSM instruction, respectively. The last two paragraphs should be as follows:

The following state is not automatically saved and restored following an SMI and the RSM instruction, respectively:

- Debug registers DR0 through DR3.
- The FPU registers.
- The MTRRs.
- Control register CR2.
- The model-specific registers (for the P6 family and Pentium® processors), or test registers TR3 through TR7 (for the Pentium and Intel486™ processors).
- The state of the trap controller.
- The Machine-Check architecture registers.
- The APIC internal interrupt state (ISR, IRR, etc.).
- The Microcode Update state.

If an SMI is used to power down the processor, a power-on reset will be required before returning to SMM, which will reset much of this state back to its default values. So an SMI handler that is going to trigger power down should first read these registers listed above directly, and save them (along with the rest of RAM) to nonvolatile storage. After the power-on reset, the continuation of the SMI handler should restore these values, along with the rest of the system's state. Anytime the SMI handler changes these registers in the processor it must also save and restore them.

NOTE

A small subset of the MSRs (such as the time-stamp counter and performance-monitoring counter) are not arbitrarily writeable and therefore cannot be saved and restored. SMM-based power-down and restoration should only be performed with operating systems that do not use or rely on the values of these registers. Operating system developers should be aware of this fact and ensure that their operating-system assisted power-down and restoration software is immune to unexpected changes in these register values.

H6. Cache and TLB Description Correction

In the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference*, Table 3-7, the correct description for descriptor value 02H should be as follows:

Descriptor Value	Cache or TLB Description
02H	Instruction TLB: 4-Mbyte Pages, fully associative, 2 entries

Also, the third bullet after the table should be as follows:

- Bytes 1, 2 and 3 of register EAX indicate that the processor contains the following:
 - 01H–A 32-entry instruction TLB (4-way set associative) for mapping 4-Kbytes pages.
 - 02H–A **2**-entry instruction TLB (**fully** associative) for mapping 4-Mbyte pages.
 - 03H–A 64-entry data TLB (4-way set associative) for mapping 4-Kbyte pages.

For the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, Table 9-1, the following corrections should be made:

Cache or Buffer	Characteristics
Instruction TLB (Large Pages)	<ul style="list-style-type: none"> - P6 family processors: 2 entries, fully associative. - Pentium® processor: Uses same TLB as used for 4-Kbyte pages. - Intel486™ processor: None (large pages not supported).

H7. SMRAM State Save Map Contains Documentation Errors

In the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, Chapter 11, "System Management Mode," Table 11-1 incorrectly documents the SMBASE+Offset for IDT Base and GDT Base for Intel® Mobile Celeron™ processors.

The storage locations for these parameters are model specific (i.e., they may differ between the Pentium processor, the Pentium Pro processor, Pentium II processor, Intel Celeron processor, Intel Mobile Celeron processor, and other P6 family proliferations). These entries in the tables above will be changed to Reserved. Hardware and software may not rely on the contents of these Reserved regions.

H8. OF and DF of the EFLAGS Register are Mislabeled as System Flags

In Table 3-7 of the *Intel Architecture Software Developer's Manual, Volume 1: Basic Architecture*, the Overflow Flag (OF) and Direction Flag (DF) are both incorrectly labeled as System Flags. The Overflow Flag should be labeled as a Status Flag and the Direction Flag should be labeled as a Control Flag.

H9. CS:EIP Pushed Onto Stack Prior to Code Segment Limit Check

The *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference*, Section 3.4, contain a detailed definition of the CALL instruction. In this definition, all instances where the instruction pointer is checked to ensure it is within the acceptable code segment limit followed by the CS:EIP register being pushed on the stack are in error. CS:EIP is pushed on the stack prior to the check of the instruction pointer. This means that in the case of a GP#(0) being generated due to an out-of-range instruction pointer, these values will be present on the stack.

H10. Corrections to Opcode Maps

In Appendix A, "Opcode Map," in the *Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference*, one and two byte opcode maps are documented. The following tables are intended to replace those tables in their entirety.

Table A-1. One-Byte Opcode Map¹

	0	1	2	3	4	5	6	7
0	ADD						PUSH	POP
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv	ES	ES

Table A-1. One-Byte Opcode Map (Continued)

	0	1	2	3	4	5	6	7
1	ADC						PUSH	POP
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv	SS	SS
2	AND							DAA
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv	=ES	
3	XOR							AAA
	Eb,Gb	Ev,Gv	Gb,Eb	Gb,Ev	AL,Ib	eAX,Iv	=SS	
4	INC general register							
	eAX	eCX	EDX	eBX	eSP	eBP	eSI	eDI
5	PUSH general register							
	eAX	eCX	EDX	eBX	eSP	eBP	eSI	eDI
6	PUSHA/ PUSHA D	POPA/ POPAD	BOUND	ARPL			Operand	Address
			Gv,Ma	Ew,Gw	=FS	=GS	Size	Size
7	Short-displacement jump on condition (Jb)							
	JO	JNO	JB/JNAE /JC	JNB/ JAE/JNC	JZ/JE	JNZ/ JNE	JBE/ JNA	JNBE/ JA
8	Imm Group 1 ²			Imm Group 1 ²	TEST		XCHG	

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

	Eb,Ib	Ev,Iv	Ev,Ib	Ev,Ib	Eb,Gb	Ev,Gv	Eb,Gb	Ev,Gv
9	NOP	XCHG word or double-word register with eAX						
		eCX	EDX	eBX	eSP	eBP	eSI	eDI
A	MOV				MOVSB	MOVSW	CMPSB	CMPSW
	AL,Ob	eAX,Ov	Ob,AL	Ov,eAX	Xb,Yb	Xv,Yv	Xb,Yb	Xv,Yv
B	MOV immediate byte into byte register							
	AL	CL	DL	BL	AH	CH	DH	BH
C	Shift Group 2 ²		RET near		LES	LDS	MOV	
	Eb,Ib	Ev,Ib	Iw		Gv,Mp	Gv,Mp	Eb,Ib	Ev,Iv
D	Shift Group 2 ²				AAM	AAD		XLAT/ XLATB
	Eb,1	Ev,1	Eb,CL	Ev,CL				

Table A-1. One-Byte Opcode Map (Continued)

	0	1	2	3	4	5	6	7
E	LOOPN E/ LOOPN Z	LOOPE/ LOOPZ	LOOP	JCXZ/ JECXZ	IN		OUT	
	Jb	Jb	Jb	Jb	AL,Ib	eAX,Ib	Ib,AL	Ib,eAX
F	LOCK		REPNE	REP/ REPE	HLT	CMC	Unary Group 32	
							Eb	Ev
	8	9	A	B	C	D	E	F
0	OR						PUSH	2-byte
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv	CS	Escape
1	SBB						PUSH	POP
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv	DS	DS
2	SUB							DAS
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv		
3	CMP							AAS
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	AL,Ib	eAX,Iv	=DS	
4	DEC General-Purpose Register							
	eAX	eCX	EDX	eBX	eSP	eBP	eSI	eDI
5	POP Into General-Purpose Register							
	eAX	eCX	EDX	eBX	eSP	eBP	eSI	eDI
6	PUSH	IMUL	PUSH	IMUL	INSB	INSW/D	OUTSB	OUTSW/ D
	Iv	Gv,Ev,Iv	Lb	Gv,Ev,Ib	Yb,DX	Yv,DX	Dx,Xb	DX,Xv

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

7	Short-Displacement Jump on Condition (Jb)							
	JS	JNS	JP/JPE	JNP/JPO	JL/JNGE	JNL/JGE	JLE/JNG	JNLE/JG
8	MOV					LEA	MOV	POP
	Eb,Gb	Ev,Gv	Gb,Eb	Gv,Ev	Ew,Sw	Gv,M	Sw,Ew	Ev
9	CBW/ CWDE	CWD/ CDQ	CALL	FWAIT	PUSHF/ PUSHFD	POPF/ POPFd	SAHF	LAHF
			Ap		Fv	Fv		

Table A-1. One-Byte Opcode Map (Continued)

	8	9	A	B	C	D	E	F
A	TEST		STOS/ STOSB	STOS/S TOSW/S TOTSD	LODSB	LODSW/ LODSD	SCAS/ SCACSB	SCASW/ SCASD/ SCAS
	AL,Ib	eAX,Iv	Yb,AL	Yv,eAX	AL,Xb	eAX,Xv	AL,Yb	eAX,Yv
B	MOV Immediate Word or Double Into Word or Double Register							
	eAX	eCX	eDX	eBX	eSP	eBP	eSI	eDI
C	ENTER	LEAVE	RET far	RET far	INT 3	INT	INTO	IRET
	Iw, Ib		Iw			lb		
D	ESC (Escape to Coprocessor Instruction Set)							
E	CALL	JMP			IN		OUT	
	Jv	Jv	Ap	Jb	AL,DX	eAX,DX	DX,AL	DX,eAX
F	CLC	STC	CLI	STI	CLD	STD	Group 4 ²	Group 5 ²

NOTES:

1. All blanks in the opcode map are reserved and should not be used. Do not depend on the operation of these undefined opcodes.
2. Bits 5, 4, and 3 of ModR/M byte used as an opcode extension (see Section A.4.).

Table A-2. Two Byte Opcode Map (First byte is 0FH)¹

	0	1	2	3	4	5	6	7
0	Group 6 ²	Group 7 ²	LAR	LSL			CLTS	
			Gv,Ew	Gv,Ew				
1								
2	MOV							
	Rd,Cd	Rd,Cd	Cd,Rd	Dd,Rd				
3	WRMSR	RDTSC	RDMSR	RDPMC				

Table A-2. Two Byte Opcode Map (Continued)

4	CMOVO	CMOVN O	CMOVB/ CMOVC/ CMOVNAE	CMOVAE/ CMOVNB/ CMOVNC	CMOVE / CMOVZ	CMOVN E/CMO VNZ	CMOVB E/CMO VNA	CMOVA/ CMOVNBE
	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev
5								
6	PUNPCKLBW	PUNPCKLWD	PUNPCKLDQ	PACKSSDW	PCMPGTB	PCMPGTW	PCMPGTD	PACKUSWB
	Pq, Qd	Pq, Qd	Pq, Qd	Pq, Qd	Pq, Qd	Pq, Qd	Pq, Qd	Pq, Qd
7	Group A ²				PCMPEQB	PCMPEQW	PCMPEQD	EMMS
		PSHIM W ³	PSHIMD 3	PSHIMQ 3	Pq, Qd	Pq, Qd	Pq, Qd	
8	Long-Displacement Jump on Condition (Jv)							
	JO	JNO	JB/JNAE / JC	JAE/JNB /JNC	JE/JZ	JNE/JNZ	JBE/JNA	JA/JNBE
9	Byte Set on condition (Eb)							
	SETO	SETNO	SETB/ SETC/ SETNA	SETAE/ SETNB/ SETNC	SETE/ SETG/ SETZ	SETNE/ SETNZ	SETBE/ SETNA	SETA/ SETNBE
A	PUSH	POP	CPUID	BT	SHLD	SHLD		
	FS	FS		Ev,Gv	Ev,Gv,I b	Ev,Gv,C L		
B	CMPXCHG	CMPXCHG	LSS	BTR	LFS	LGS	MOVZX	
	Eb,Gb	Ev,Gv	Mp	Ev,Gv	Mp	Mp	Gv,Eb	Gv,Ew
C	XADD	XADD						Group 9 ²
	Eb,Gb	Ev,Gv						
D		PSRLW	PSRLD	PSRLQ		PMULL W		
		Pq, Qd	Pq, Qd	Pq, Qd		Pq, Qd		
E		PSRAW	PSRAD			PMULH W		
		Pq, Qd	Pq, Qd			Pq, Qd		
F		PSLLW	PSLLD	PSLLQ		PMADDWD		
		Pq, Qd	Pq, Qd	Pq, Qd		Pq, Qd		

Table A-2. Two Byte Opcode Map (Continued)

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

	8	9	A	B	C	D	E	F
0	INVD	WBINVD		UD24				
1								
2								
3								
4	CMOVS	CMOVNS	CMOVP/ CMOVPE	CMOVN P/CMOV PO	CMOVL/ CMOVNGE	CMOVG E/CMO VNL	CMOVL E/CMO VNG	CMOVG/ CMOVNLE
	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev	Gv, Ev
5								
6	PUNPC KHBW	PUNPC KHWD	PUNPC KHDQ	PACKSS DW			MOVD	MOVQ
	Pq,Qd	Pq,Qd	Pq,Qd	Pq,Qd			Pd,Ed	Pq,Qq
7							MOVD	MOVQ
							Ed,Pd	Qq,Pq
8	Long-Displacement Jump on Condition (Jv)							
	JS	JNS	JP/JPE	JNP/JPO	JL/JNGE	JNL/JGE	JLE/JNG	JNLE/JG
	Byte set on condition (Eb)							
9	SETS	SETNS	SETP/ SETPE	SETNP/ SETPO	SETL/ SETNG E	SETNL/ SETGE	SETLE/ SETNG	SETNLE
	Eb	Eb	Eb	Eb	Eb	Eb	Eb	Eb
A	PUSH	POP	RSM	BTS	SHRD	SHRD		IMUL
	GS	GS		Ev,Gv	Ev,Gv,Ib	Ev,Gv,C L		Gv,Ev
B		Invalid Opcode ⁴	Group 8 ²	BTC	BSF	BSR	MOVSX	
			Ev,lb	Ev,Gv	Gv,Ev	Gv,Ev	Gv,Eb	Gv,Ew

Table A-2. Two Byte Opcode Map (Continued)

	8	9	A	B	C	D	E	F
--	----------	----------	----------	----------	----------	----------	----------	----------

C	BSWAP							
	EAX	ECX	EDX	EBX	ESP	EBP	ESI	EDI
D	PSUBUSB	PSUBUSW		PAND	PADDUSB	PADDUSW		PANDN
	Pq,Qq	Pq,Qq		Pq,Qq	Pq,Qq	Pq,Qq		Pq,Qq
E	PSUBSB	PSUBS W		POR	PADDS B	PADDS W		PXOR
	Pq,Qq	Pq,Qq		Pq,Qq	Pq,Qq	Pq,Qq		Pq,Qq
F	PSUBB	PSUBW	PSUBD		PADDB	PADDW	PADDD	
	Pq,Qq	Pq,Qq	Pq,Qq		Pq,Qq	Pq,Qq	Pq,Qq	

NOTES:

1. All blanks in the opcode map are reserved and should not be used. Do not depend on the operation of these undefined opcodes.
2. Bits 5, 4, and 3 of ModR/M byte used as an opcode extension (see Section A.4.).
3. These abbreviations are not actual mnemonics. When shifting by immediate shift counts, the PSHMD mnemonic represents the PSLLD, PSRAD, and PSRLD instructions, PSHMW represents the PSLW, PSRAW, and PSRLW instructions, and PSHMQ represents the PSLQ and PSRLQ instructions. The instructions that shift by immediate counts are differentiated by the ModR/M bytes (see Section A.4.).
4. Use the 0F0B opcode (UD2 instruction) or the 0FB9H opcode when deliberately trying to generate an invalid opcode exception (#UD).

H11. Interrupt 13-General Protection Exception (#GP)

In Section 5.12 of the *Intel Architecture Software Developer’s Manual, Volume 3: System Programming Guide*, a description of the exception interrupts is provided. In the description section of Interrupt 13-General Protection Exception (#GP), the last bullet applies if the PAE and/or PSE flags are set, rather than just the PAE flag as reported in the documentation.

H12. Corrections to Intel Architecture Software Developer’s Manual, Volume 2: Instruction Set Reference

- The following typographical errors and other documentation errors will be corrected in the next revision of the *Intel Architecture Software Developer’s Manual, Volume 2: Instruction Set Reference*. A list of significant changes is given below. Note that other changes may be made, and not all significant changes may be listed here.
- Page 3-79: The example for the DAA instruction is incorrect, and should read:


```
ADD AL, BL Before: AL=79H BL=35H EFLAGS(0SZAPC)=XXXXXX
           After: AL=AEH BL=35H EFLAGS(0SZAPC)=110000
DAA       Before: AL=2EH BL=35H EFLAGS(0SZAPC)=110000
           After: AL=04H BL=35H EFLAGS(0SZAPC)=x00101
```
- Page 3-236: The TASK-RETURN parameters are (* PE=1, VM=0, NT=1 *).
- Page 3-350: The second paragraph of the description should begin “The current **operand-size** attribute...”

H13. MCI_ADDR MSR Reference Section Correction

The first sentence of Section 12.3.2.3 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, contain a reference to a previous section, but incorrectly identify the referenced section number. The first sentence should read: "The MCI_ADDR MSR contains the address of the code or data memory location that produced the machine-check error if the ADDR_V flag in the MCI_STATUS register is set (see Section 12.3.2.2, "MCI_STATUS MSR")."

H14. FCOMI/FCOMIP/FUCOMI/FUCOMIP Setting of Flags Relative to Exceptions

Page 3-112 of the Intel Architecture Software Developer's Manual, Volume 2: Instruction Set Reference, shows a table for FCOMI/FCOMIP/FUCOMI/FUCOMIP comparison results, where the last entry in the table "Unordered" has an asterisk (*) beside it referencing a table note that reads: "Note: * Flags not set if unmasked invalid-arithmetic operand (#IA) exception is generated." However, this note should read: "Note: * Flags are set regardless, whether there is an unmasked invalid operand (#IA) exception generated or not."

H15. MemTypeGet() Function Example

Example 9-2 of Section 9.11.7.1 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, contains pseudocode that uses the MemTypeGet() function. The line that reads: "IF (BASE + SIZE) wrap 4-Gbyte address space THEN return INVALID" is incorrect. This line should read: "IF (BASE + SIZE) wrap 64-Gbyte address space THEN return INVALID."

H16. RSVD Flag Correction

Figure 5-7 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, contains a definition of the RSVD flag. The definition is reversed and should read as follows:

RSVD	1	The page fault occurred because a 1 was detected in one of the reserved bit positions of	a page table entry or directory entry that was marked present.
	0	The fault was not caused by a reserved bit violation.	

H17. SMBus Data Setup Time

The SMBus data signal has a min setup time (min Tsu: DAT) of 250ns, as published in the System Management Bus Specification rev 1.0. This parameter is currently specified to 800 ns for the Mobile Pentium II processor (mini-cartridge) and Pentium II Processor mobile module products (see table below).

SMBus Data Min Setup

Parameter	Description	Spec	Change to
min Tsu: DAT	min SMBus Data Setup Time	250 ns	800 ns

H18. SMRAM State Save Map Documentation Correction

In the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, chapter 11, "System Management Mode," section 11.4.1, a description of the register status is provided. It is stated,

"The following registers are saved (but not readable) and restored upon exiting SMM:

- Control register CR4."

It states:

"The following registers are saved (but not readable) and restored upon exiting SMM: Control register CR4 (CR4 is set to "0" while in the SMM handler)."

SPECIFICATION CLARIFICATIONS

The Specification Clarifications listed in this section apply to the Intel® Mobile Celeron™ Processor (BGA) Datasheet (Order Number#245106-001), or the Intel® Mobile Celeron Processor in Mobile Module MMC-1 (Order Number #245101-001), or Intel® Mobile Celeron Processor in Mobile Module MMC-2 (Order Number #245102-001). All Specification Changes will be incorporated into a future version on the appropriate Intel Celeron processor documentation.

NOTE

The *Intel Architecture Software Developer's Manual, Volumes 1, 2, and 3*, applies to all P6 family processors, and therefore some of the specification clarifications in this section may not pertain to the Intel® Mobile Celeron™ processor or the Intel Celeron™ Processor Mobile Module specifically.

H1. Writes to WC Memory

Section 9.3 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, identifies that “Writes” to a region of WC memory “may be delayed and combined in the write buffer to reduce memory accesses.” This sentence should state that “Writes” to a region of WC memory “may be delayed and combined in the write buffer to reduce memory accesses. The writes may be delayed until the next occurrence of a buffer or processor serialization event, e.g., CPUID execution, a read or write to uncached memory, interrupt occurrence, LOCKed instruction execution, etc., if the WC buffer is partially filled.”

H2. Multiple Processor Protocol and Restrictions

Section 7.6.1 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, contain inconsistencies which will be clarified as follows:

7.6.1 Protocol Requirements and Restrictions

The MP protocol imposes the following requirements and restrictions on the system:

- An APIC clock (APICLK) must be provided on all systems based on the P6 family processors (excluding mobile processors and modules).
- All interrupt mechanisms must be disabled for the duration of the MP protocol algorithm, including the window of time between the assertion of INIT# or receipt of an INIT IPI by the application processors and the receipt of a STARTUP IPI by the application processors. That is, requests generated by interrupting devices must not be seen by the local APIC unit (on board the processor) until the completion of the algorithm. Failure to disable the interrupt mechanisms may result in processor shutdown.

- The MP protocol should be initiated only after a hardware reset. After completion of the protocol algorithm, a flag is set in the APIC base MSR of the BSP (APIC_BASE.BSP) to indicate that it is the BSP. This flag is cleared for all other processors. If a processor or the system is subject to an INIT sequence (either through the INIT# pin or an INIT IPI), then the MP protocol is not re-executed. Instead, each processor examines its BSP flag to determine whether the processor should boot or wait for a STARTUP IPI.

H3. Critical Sequence of Events During a Page Fault Exception

Section 3.6.4 of the *Intel Architecture Software Developer’s Manual, Volume 3: System Programming Guide*, will be clarified as follows:

If the processor generates a page-fault exception, the operating system must carry out the following operations in this order:

1. Copy the page from disk storage into physical memory if needed.
2. Load the page address into the page-table or page-directory entry and set its present flag. Other bits, such as the dirty and accessed bits, may also be set at this time.
3. Invalidate the current page table entry in the TLB (see Section 3.7, “Translation Lookaside Buffers (TLBs)” for a discussion of TLBs and how to invalidate them).
4. Return from the page fault handler to restart the interrupted program or task.

H4. Performance-Monitoring Counter Issues

The following table replaces Table A-1 of the *Intel Architecture Software Developer’s Manual, Volume 3: System Programming Guide*. The only changes to this new table are enhanced descriptions of the events counted.

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
Data Cache Unit (DCU)	43H	DATA_ MEM_ REFS	00H	All loads from any memory type. All stores to any memory type. Each part of a split is counted separately. The internal logic counts not only external memory loads and stores, but also internal retries.	

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				<p>Note: 80-bit floating-point accesses are double counted, since they are decomposed into a 16-bit exponent load and a 64-bit mantissa load. Memory accesses are only counted when they are actually performed. E.g., a load that gets squashed because a previous cache miss is outstanding to the same address, and which finally gets performed, is only counted once. Does not include I/O accesses, or other nonmemory accesses.</p>	
	45H	DCU_LINES_IN	00H	Total lines allocated in the DCU.	
	46H	DCU_M_LINES_IN	00H	Number of M state lines allocated in the DCU.	
	47H	DCU_M_LINES_OUT	00H	Number of M state lines evicted from the DCU. This includes	

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				evictions via snoop HITM, intervention or replacement.	
	48H	DCU_MISS_OUTSTANDING	00H	Weighted number of cycles while a DCU miss is outstanding, incremented by the number of outstanding cache misses at any particular time. Cacheable read requests only are considered. Uncacheable requests are excluded. Read-for-ownerships are counted as well as line fills, invalidates, and stores.	An access that also misses the L2 is short-changed by 2 cycles. (e.g., if count is N cycles, should be N+2 cycles.) Subsequent loads to the same cache line will not result in any additional counts. Count value not precise, but still useful.
Instruction Fetch Unit (IFU)	80H	IFU_IFETCH	00H	Number of instruction fetches, both cacheable and noncacheable. Including UC fetches.	
	81H	IFU_IFETCH_MISSES	00H	Number of instruction fetch misses. All instruction fetches that do not hit the IFU, e.g., that produce memory	

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				requests. Includes UC accesses.	
	85H	ITLB_ MISS	00H	Number of ITLB misses.	
	86H	IFU_ MEM_ STALL	00H	Number of cycles instruction fetch is stalled, for any reason. Includes IFU cache misses, ITLB misses, ITLB faults and other minor stalls.	
	87H	ILD_ STALL	00H	Number of cycles that the instruction length decoder is stalled.	
L2 Cache ¹	28H	L2_ IFETCH	MESI 0FH	Number of L2 instruction fetches. This event indicates that a normal instruction fetch was received by the L2. The count includes only L2 cacheable instruction fetches; it does not include UC instruction fetches. It does not include ITLB miss accesses.	
	29H	L2_ LD	MESI 0FH	Number of L2 data loads. This event indicates that a normal, unlocked, load	

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				memory access was received by the L2. It includes only L2 cacheable memory accesses; it does not include I/O accesses, other nonmemory accesses, or memory accesses such as UC/WT memory accesses. It does include L2 cacheable TLB miss memory accesses.	
	2AH	L2_ST	MESI 0FH	Number of L2 data stores. This event indicates that a normal, unlocked, store memory access was received by the L2. Specifically, it indicates that the DCU sent a read-for-ownership request to the L2. It also includes Invalid to Modified requests sent by the DCU to the L2. It includes only L2 cacheable store memory accesses; it does not include	

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				I/O accesses, other nonmemory accesses, or memory accesses like UC/WT stores. It includes TLB miss memory accesses.	
	24H	L2_LINES_IN	00H	Number of lines allocated in the L2.	
	26H	L2_LINES_OUT	00H	Number of lines removed from the L2 for any reason.	
	25H	L2_M_LINES_INM	00H	Number of modified lines allocated in the L2.	
	27H	L2_M_LINES_OUTM	00H	Number of modified lines removed from the L2 for any reason.	
	2EH	L2_RQSTS	MESI 0FH	Total number of L2 requests.	
	21H	L2_ADS	00H	Number of L2 address strobes.	
	22H	L2_DBUS_B USY	00H	Number of cycles during which the L2 cache data bus was busy.	
	23H	L2_DBUS_B USY_RD	00H	Number of cycles during which the data bus was busy transferring read data from L2 to the processor.	

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
External Bus Logic (EBL) ²	62H	BUS_DRDY_CLOCKS	00H (Self) 20H (Any)	Number of clocks during which DRDY# is asserted. Essentially, utilization of the external system data bus.	Unit Mask = 00H counts bus clocks when the processor is driving DRDY#. Unit Mask = 20H counts in processor clocks when any agent is driving DRDY#.
	63H	BUS_LOCK_CLOCKS	00H (Self) 20H (Any)	Number of clocks during which LOCK# is asserted on the external system bus.	Always counts in processor clocks.
	60H	BUS_REQ_OUTSTANDING	00H (Self)	Number of bus requests outstanding. This counter is incremented by the number of cacheable read bus requests outstanding in any given cycle.	Counts only DCU full-line cacheable reads, not RFOs, writes, instruction fetches, or anything else. Counts “waiting for bus to complete” (last data chunk received).
	65H	BUS_TRAN_BRD	00H (Self) 20H (Any)	Number of burst read transactions.	
	66H	BUS_TRAN_RFO	00H (Self) 20H (Any)	Number of completed read for ownership transactions.	
	67H	BUS_TRANS_WB	00H (Self)	Number of completed write	

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
			20H (Any)	back transactions.	
	68H	BUS_ TRAN_ IFETCH	00H (Self) 20H (Any)	Number of completed instruction fetch transactions.	
	69H	BUS_ TRAN_ INVAL	00H (Self) 20H (Any)	Number of completed invalidate transactions.	
	6AH	BUS_ TRAN_ PWR	00H (Self) 20H (Any)	Number of completed partial write transactions.	
	6BH	BUS_ TRANS_P	00H (Self) 20H (Any)	Number of completed partial transactions.	
	6CH	BUS_ TRANS_IO	00H (Self) 20H (Any)	Number of completed I/O transactions.	
	6DH	BUS_ TRAN_DEF	00H (Self) 20H (Any)	Number of completed deferred transactions.	
	6EH	BUS_ TRAN_ BURST	00H (Self) 20H (Any)	Number of completed burst transactions.	
	70H	BUS_ TRAN_ ANY	00H (Self) 20H (Any)	Number of all completed bus transactions. Address bus utilization can be calculated	

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				knowing the minimum address bus occupancy. Includes special cycles, etc.	
	6FH	BUS_TRAN_MEM	00H (Self) 20H (Any)	Number of completed memory transactions.	
	64H	BUS_DATA_RCV	00H (Self)	Number of bus clock cycles during which this processor is receiving data.	
	61H	BUS_BNR_DRV	00H (Self)	Number of bus clock cycles during which this processor is driving the BNR# pin.	
	7AH	BUS_HIT_DRV	00H (Self)	Number of bus clock cycles during which this processor is driving the HIT# pin.	Includes cycles due to snoop stalls. The event counts correctly, but the BPMi pins function as follows based on the setting of the PC bits (bit 19 in the PerfEvtSel0 and PerfEvtSel1 registers). If the core clock to bus clock ratio is 2:1 or 3:1, and a PC bit is set, the

INTEL®MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
					<p>BPM_i pins will be asserted for a single clock when the counters overflow. If the PC bit is clear, the processor toggles the BPM_i pins when the counter overflows. If the clock ratio is not 2:1 or 3:1, the BPM_i pins will not function for these performance-monitoring counter events.</p>
	7BH	BUS_HITM_DRV	00H (Self)	Number of bus clock cycles during which this processor is driving the HITM# pin.	<p>Includes cycles due to snoop stalls.</p> <p>The event counts correctly, but the BPM_i pins function as follows based on the setting of the PC bits (bit 19 in the PerfEvtSel0 and PerfEvtSel1 registers). If the core clock to bus clock ratio is 2:1 or 3:1, and a PC bit is set, the BPM_i pins will be asserted for a single clock when</p>

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
					the counters overflow. If the PC bit is clear, the processor toggles the BPM _i pins when the counter overflows. If the clock ratio is not 2:1 or 3:1, the BPM _i pins will not function for these performance-monitoring counter events.
	7EH	BUS_SNOOP_STALL	00H (Self)	Number of clock cycles during which the bus is snoop stalled.	
Floating Point Unit	C1H	FLOPS	00H	Number of computational floating-point operations retired. Excludes floating-point computational operations that cause traps or assists. Includes floating-point computational operations executed by the assist handler. Includes internal sub-operations of complex floating-point instructions	Counter 0 only

INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				like transcendentals. Excludes floating-point loads and stores.	
	10H	FP_COMP_OPS_EXE	00H	Number of computational floating-point operations executed. The number of FADD, FSUB, FCOM, FMULs, integer MULs and IMULs, FDIVs, FPREMs, FSQRTS, integer DIVs and IDIVs. Note not the number of cycles but, the number of operations. This event does not distinguish an FADD used in the middle of a transcendental flow from a separate FADD instruction.	Counter 0 only
	11H	FP_ASSIST	00H	Number of floating-point exception cases handled by microcode.	Counter 1 only. This event includes counts due to speculative execution.
	12H	MUL	00H	Number of multiplies. Note: includes integer	Counter 1 only

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				and well FP multiplies and is speculative.	
	13H	DIV	00H	Number of divides. Note: includes integer and FP multiplies and is speculative.	Counter 1 only
	14H	CYCLES_DIV_BUSY	00H	Number of cycles that the divider is busy, and cannot accept new divides. Note: includes integer and FP divides, FPREM, FPSQRT, etc., and is speculative.	Counter 0 only
Memory Ordering	03H	LD_BLOCKS	00H	Number of store buffer blocks. Includes counts caused by preceding stores whose addresses are unknown, preceding stores whose addresses are known to conflict, but whose data is unknown and preceding stores that conflicts with the load, but which incompletely overlap the load.	

INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
	04H	SB_ DRAINS	00H	Number of store buffer drain cycles. Incremented during every cycle the store buffer is draining. Draining is caused by serializing operations like CPUID, synchronizing operations like XCHG, Interrupt acknowledgment as well as other conditions such as cache flushing.	
	05H	MIS- ALIGN_ MEM_REF	00H	Number of misaligned data memory references. Incremented by 1 every cycle during which either the proc load or store pipeline dispatches a misaligned uop. Counting is performed if it's the first half or second half, or if it is blocked, squashed or misses. Note in this	It should be noted that MISALIGN_MEM_REF is only an approximation, to the true number of misaligned memory references. The value returned is roughly proportional to the number of misaligned memory accesses, e.g., the size of the problem.

INTEL® MOBILE Celeron™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				context misaligned means crossing a 64 bit boundary.	
In-struction De-coding and Retirement	C0H	INST_RETIRED	00H	Number of instructions retired.	A hardware interrupt received during/after the last iteration of the REP STOS flow causes the counter to undercount by 1 instruction.
	C2H	UOPS_RETIRED	00H	Number of UOPs retired.	
	D0H	INST_DECOD-ER	00H	Number of instructions decoded.	
Inter-rupts	C8H	HW_INT_RX	00H	Number of hardware interrupts received.	
	C6H	CYCLES_INT_MASKED	00H	Number of processor cycles for which interrupts are disabled.	
	C7H	CYCLES_INT_PENDING_AND_MASKED	00H	Number of processor cycles for which interrupts are disabled and interrupts are pending.	
Bran-ches	C4H	BR_INST_RETIRED	00H	Number of branch instructions	

INTEL® MOBILE CELERON™ PROCESSOR SPECIFICATION UPDATE

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				retired.	
	C5H	BR_MISS_P RED_ RETIRED	00H	Number of mispredicted branches retired.	
	C9H	BR_ TAKEN_ RETIRED	00H	Number of taken branches retired.	
	CAH	BR_MISS_P RED_ TAKEN_ RET	00H	Number of taken mispredictions branches retired.	
	E0H	BR_INST_ DECOD-ED	00H	Number of branch instructions decoded.	
	E2H	BTB_ MISSES	00H	Number of branches that for which the BTB did not produce a prediction	
	E4H	BR_BOGUS	00H	Number of bogus branches.	
	E6H	BA-CLEAR	00H	Number of time BACLEAR is asserted. This is the number of times that a static branch prediction was made, where the branch decoder decided to make a branch prediction because the BTB did not.	
Stalls	A2H	RE- SOURCE_S TALLS	00H	Incremented by one during every cycle that there is	

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
				a resource related stall. Includes register renaming buffer entries, memory buffer entries. Does not include stalls due to bus queue full, too many cache misses, etc. In addition to resource related stalls, this event counts some other events. Includes stalls arising during branch misprediction recovery, e.g., if retirement of the mispredicted branch is delayed and stalls arising while store buffer is draining from synchronizing operations.	
	D2H	PARTIAL_RAT_STALLS	00H	Number of cycles or events for partial stalls. Note Includes flag partial stalls.	
Segment Register Loads	06H	SEGMENT_REG_LOADS	00H	Number of segment register loads	

Unit	Event Number	Mnemonic Event Name	Unit Mask	Description	Comments
Clocks	79H	CPU_CLK_UN-HALTED	00H	Number of cycles during which the processor is not halted.	
MMX™ Unit	B0H	MMX_INST R_EXEC	00H	Number of MMX Instructions Executed	Available in Intel® Celeron™, Pentium® II and Pentium II Xeon™ processors only. Does not account for MOVQ and MOVD stores from register to memory.

NOTES:

- Several L2 cache events, where noted, can be further qualified using the Unit Mask (UMSK) field in the PerfEvtSel0 and PerfEvtSel1 registers. The lower 4 bits of the Unit Mask field are used in conjunction with L2 events to indicate the cache state or cache states involved. The Pentium® II processor identifies cache states using the “MESI” protocol and consequently each bit in the Unit Mask field represents one of the four states: UMSK[3] = M (8H) state, UMSK[2] = E (4H) state, UMSK[1] = S (2H) state, and UMSK[0] = I (1H) state. UMSK[3:0] = MESI (FH) should be used to collect data for all states; UMSK = 0H, for the applicable events, will result in nothing being counted.
- All of the external bus logic (EBL) events, except where noted, can be further qualified using the Unit Mask (UMSK) field in the PerfEvtSel0 and PerfEvtSel1 registers. Bit 5 of the UMSK field is used in conjunction with the EBL events to indicate whether the processor should count transactions that are self generated (UMSK[5] = 0) or transactions that result from any processor on the bus (UMSK[5] = 1).

H5. POP[ESP] with 16-bit Stack Size

In the *Intel Architecture Software Developer’s Manual, Volume 2: Instruction Set Reference*, the section regarding “POP–Pop a Value from the Stack,” the following note:

“If the ESP register is used as a base register for addressing a destination operand in memory, the POP instruction computes the effective address of the operand after it increments the ESP register.”

is incomplete, and should read as follows:

“If the ESP register is used as a base register for addressing a destination operand in memory, the POP instruction computes the effective address of the operand after it increments the ESP register. For the case of a 16-bit stack where ESP wraps to 0h as a result of the POP instruction, the resulting location of the memory write is processor family specific.”

In Section 17.23.1 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, a new section will be added:

A POP-to-memory instruction, Which Uses the Stack Pointer (ESP) as a Base Register.

For a POP-to-memory instruction that meets the following conditions:

1. The stack segment size is 16-bits,
2. Any 32-bit addressing form with the SIB byte specifying ESP as the base register, and
3. The initial stack pointer is FFFCh(32-bit operand) or FFFEh (16-bit operand) and will wrap around to 0h as a result of the POP operation.

The result of the memory write is processor family specific. For example, in Pentium II and Pentium Pro processors the result of the memory write is to SS:0h plus any scaled index and displacement. In Pentium and i486™ processors, the result of the memory write may be either a stack fault (real mode or protected mode with stack segment size of 64-Kbytes), or write to SS:10000h plus any scaled index and displacement (protected mode and stack segment size exceeds 64-Kbytes).

H6. Preventing Caching

Section 9.5.2 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, document the procedure to prevent the L1 and L2 caches from performing all caching operations. However, this procedure differs from that given in Section 9.11.8, "Multiple-Processor Considerations." The correct procedure that should be used is as follows:

1. Enter the no-fill cache mode. (Set the CD flag in control register CR0 to 1 and the NW flag to 0.)
2. Flush all caches using the WBINVD instruction.
3. Disable the MTRRs and set the default memory type to uncached, or set all MTRRs for the uncached memory type (see the discussion of the TYPE field and the E flag in Section 11.11.2.1, "MTRRdefType Register").

The caches must be flushed when the CD flag is cleared to insure system memory coherency. If the caches are not flushed in step 2, cache hits on reads will still occur and data will be read from valid cache lines.

H7. Paging Must Be Enabled Before Enabling the Page Global Bit

In Section 2.5 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, the following line should be added to the text describing the Page Global Enable bit (PGE).

"In addition, the bit must not be enabled before paging is enabled via CR0.PG. Program correctness may be affected by reversing this sequence and processor performance will be impacted."

H8. *Interrupt Recognition Determines Priority*

The interrupt priority documented in Table 5-2 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, reflects the order in which interrupts will be serviced upon simultaneous recognition by the processor (for example, when multiple interrupts are pending at an instruction boundary). These tables do not necessarily reflect the order in which interrupts will be recognized by the processor if received simultaneously at the processor pins.

H9. *References to 2-Mbyte Pages Should Include 4-Mbyte Pages*

Generically, “large pages” refers to either 2-Mbyte or 4-Mbyte pages. In Section 3.8 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, 2-Mbyte pages are often referenced alone, when the behavior of 4-Mbyte pages is identical; these references should include all large pages.

H10. *Modification of Reserved Areas in the SMRAM Saved State Map*

If data is incorrectly written to reserved areas of the saved state map, the processor will enter the shutdown state. This can also occur if invalid state information is saved in the SMRAM (such as if illegal combinations of bits are written to CR0 or CR4 before an SMI is serviced). CR4 is not distinctly part of the saved state map, as implied in Section 11.3.1.1 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*.

H11. *TLB Flush Necessary After PDPE Change*

As described in Section 3.7 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, the operating system is required to invalidate the corresponding entry in the TLB after any change to a page-directory or page-table entry. However, if the physical address extension (PAE) feature is enabled to use 36-bit addressing, a new table is added to the paging hierarchy, called the page directory pointer table (as per Section 3.8, “Physical Address Extension”). If an entry is changed in this table (to point to another page directory), the TLBs must then be flushed by writing to CR3.

H12. *Exception Handler Error Code Bit Clarification*

Section 5.11 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, describe the bit definitions for the error code pushed onto the stack of the exception handler. The explanation of the EXT bit 0 will be changed to read as follows: External event (bit 0). When set, indicates that an event external to the program caused the exception, such as a hardware interrupt.

H13. Switching to Protected Mode While in SMM

Should the System Management Mode (SMM) code developer require a transition to protected mode while in SMM, a change is required to the sequence of events used to switch to protected mode as documented in Section 8.8.1 of the *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*.

Items 3 and 4 of this section state:

3. Execute a MOV CR0 instruction that sets the PE flag (and optionally the PG flag) in control register CR0.
4. Immediately following the MOV CR0 instruction, execute a far JMP or far CALL instruction. (This operation is typically a far jump or call to the next instruction in the instruction stream.)

Random failures can occur if other instructions exist between steps 3 and 4, and failures will be readily seen in some situations such as when instructions that reference memory are inserted between steps 3 and 4 above while in System Management mode.

SPECIFICATION CHANGES

The Specification Changes listed in this section apply to the Intel® Mobile Celeron™ Processor (BGA) Datasheet (Order Number#245106-001), or the Intel® Mobile Celeron Processor in Mobile Module MMC-1 (Order Number #245101-001), or Intel® Mobile Celeron Processor in Mobile Module MMC-2 (Order Number #245102-001). All Specification Changes will be incorporated into a future version of the appropriate Intel® Celeron™ processor documentation.

H1. WC Buffer Eviction Data Ordering

The *Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide*, states in Section 9.3.1 that “a completely full WC [Write Combining] buffer will always be propagated as a single burst transaction with ascending data order.” This statement is incorrect and should be changed to “a completely full WC buffer will always be propagated as a single burst transaction using any of the valid chunk orders.”