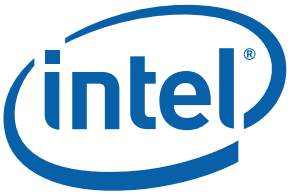

Intel® Ethernet Controller I210 Specification Update

January 2013
Revision 1.1

Revision History

Date	Revision	Description
January 2013	1.1	Added Errata #8, #9, #10 and #11.
October 2012	1.0	Initial release (Intel public). Added Specification Clarification #1 through #3. Added Errata #1 through #7. Added Software Clarification #1 and #2.

Note: This page intentionally left blank.

1.1 Introduction and Scope

This document applies to the Intel® Ethernet Controller I210.

This document is an update to a published specification, the *Intel® Ethernet Controller I210 Datasheet*. It is intended for use by system manufacturers and software developers. All product documents are subject to frequent revision, and new order numbers will apply. New documents may be added. Be sure you have the latest information before finalizing your design.

1.2 Product Code and Device Identification

Product Codes: WGI210AT (Commerical Temperature Range); WGI210IT and WGI210IS (Industrial Temperature Range).

The following tables and component drawings describe the various identifying markings on each device package:

Table 1. Markings

Device	Stepping	Top Marking	Description
I210	A2	WGI210AT	Production (Commerical Copper)
I210	A2	WGI210IT	Production (Industrial Temperature Copper)
I210	A2	WGI210IS	Production (Industrial Temperature Fiber)

Table 2. Device IDs

I210 Device ID Code	Vendor ID	Device ID	Revision ID
WGI210AT (not programmed/factory default)	0x8086	0x1531	0x3
WGI210AT (copper only)	0x8086	0x1533	0x3
WGI210IS (fiber, industrial temperature)	0x8086	0x1536	0x3
WGI210IS (100BASE-KX/BXbackplane, industrial temperature)	0x8086	0x1537	0x3
WGI210IS (external SGMII, industrial temperature)	0x8086	0x1538	0x3

Table 3. MM Numbers

Product	MM Number	Spec	Media
WGI210AT - Production (Commerical Copper)	925131	SLJXQ	Tape and Reel
WGI210AT - Production (Commerical Copper)	925132	SLJXR	Tray
WGI210IT - Production (Copper and Industrial Temperature Range)	925133	SLJXS	Tape and Reel
WGI210IT - Production (Copper and Industrial Temperature Range)	925138	SLJXT	Tray
WGI210IS - Production (Fiber and Industrial Temperature Range)	925142	SLJXW	Tape and Reel
WGI210IS - Production (Fiber and Industrial Temperature Range)	925143	SLJXX	Tray

1.3 I210 Production Marking Diagram

Figure 1. I210 Production Top Marking Example (Commerical Temperature Copper)

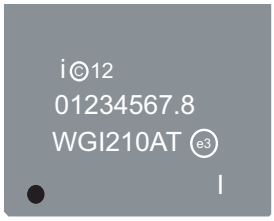

Figure 2. I210 Production Top Marking Example (Industrial Temperature Copper)

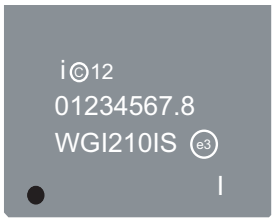

Figure 3. I210 Production Top Marking Example (Industrial Temperature Fiber)

Notes:

Line 1: With no spaces, "i"©YY
Line 2: Fab Lot Trace Code 0123456.78 (10-char max)
Line 3: S-Spec Code and Pb-free mark (e3 or e1)
Line 4: "I" in lower-right corner for industrial temperature rated devices

1.4 Nomenclature Used In This Document

This document uses specific terms, codes, and abbreviations to describe changes, errata and/or clarifications that apply to silicon/steppings. See Table 4 for a description.

Table 4. Terms, Codes, Abbreviations

Name	Description
Specification Changes	Modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.
Errata	Design defects or errors. Errata may cause device behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.
Specification Clarifications	Greater detail or further highlights concerning a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.
Documentation Changes	Typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.
Yes or No	If the errata applies to a stepping, "Yes" is indicated for the stepping (for example: "A0=Yes" indicates errata applies to stepping A0). If the errata does not apply to stepping, "No" is indicated (for example: "A0=No" indicates the errata does not apply to stepping A0).
Doc	Document change or update that will be implemented.
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.
Eval	Plans to fix this erratum are under evaluation.
(No mark) or (Blank box)	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Red Change Bar/or Bold	This Item is either new or modified from the previous version of the document.
DS	Data Sheet
DG	Design Guide
SDM	Software Developer's Manual
EDS	External Data Specification
AP	Application Note

1.5 Hardware Clarifications, Changes, Updates and Errata

See Section 1.4 for an explanation of terms, codes, and abbreviations.

Table 5. Summary of Hardware Clarifications, Changes and Errata; Errata Include Steppings

Specification Changes	**Status**
None.	N/A
Specification Clarifications	**Status**
1. PCIe Completion Timeout Mechanism Compliance	N/A
2. Padding on Transmitted SCTP Packets	N/A
3. Dynamic LED Modes Can Only Be Used in an Active Low Configuration	N/A
Documentation Updates	**Status**
None.	
Errata	**Status**
1. I^2C Data Out Hold Time Violation	A2 NoFix
2. NC-SI Hardware Arbitration Issues	A2 NoFix
3. SGMII: Counters Incorrectly Increment on Collision	A2 NoFix
4. BMC-only Packets Not Counted as Host Sent/Received Packets	A2 NoFix
5. Device Off Deadlock	A2 NoFix
6. Marginal Low 10 Mb Amplitude	A2 NoFix
7. Non-monotonic Integrated SVR Ramp	A2 NoFix
8. Protocol Offload: Incorrect Response to MLDv2 Queries	A2 NoFix
9. Writes to the VPD RW Area are Not Reliable	A2 Fixed
10. NC-SI: Get NC-SI Pass-through Statistics Response Might Contain Incorrect Packet Counts	A2 NoFix
11. MCTP Commands From SMBus are Dropped	A2 Fixed

1.5.1 Specification Changes

None at this time.

1.5.2 Specification Clarifications

1. PCIe Completion Timeout Mechanism Compliance

Clarification:

The I210 Completion Timeout Value[3:0] must be properly set by the system BIOS in the I210 PCIe Configuration Space Device Control 2 register (0xC8; W). Failure to do so can cause unexpected completion timeouts.

The I210 complies with the PCIe 2.0 specification for the completion timeout mechanism and programmable timeout values. The PCIe 2.0 specification provides programmable timeout ranges between 50 µs to 64 s with a default time range of 50 µs - 50 ms. The I210 defaults to a range of 16 ms - 32 ms.

Workaround:

The completion timeout value must be programmed correctly in PCIe configuration space (in Device Control 2 register); the value must be set above the expected maximum latency for completions in the system in which the I210 is installed. This ensures that the I210 receives the completions for the requests it sends out, avoiding a completion timeout scenario. Failure to properly set the completion timeout value can result in the device timing out prior to a completion returning.

The I210 can be programmed to resend a completion request after a completion timeout (the original completion request is assumed to be lost). But if the original completion arrives after a resend request, two completions may arrive for the same request; this can cause unpredictable behavior. Intel NVM images set the resend feature to off. Intel recommends that you do not change this setting.

2. Padding on Transmitted SCTP Packets

Clarification:

When using the I210 to offload the CRC calculation for transmitted SCTP packets, software should not add Ethernet padding bytes to short packets (less than 64 bytes). Instead, the TCTL.PSP bit should be set so that the I210 pads the packets after performing the CRC calculation.

3. Dynamic LED Modes Can Only Be Used in an Active Low Configuration

Clarification:

In any of the dynamic LED modes (FILTER_ACTIVITY, LINK/ACTIVITY, COLLISION, ACTIVITY, PAUSED), LED blinking should only be enabled if the LED signal is configured as an active low output.

1.5.3 Documentation Changes

None.

1.5.4 Errata

1. I^{2}C Data Out Hold Time Violation

Problem:

The I210 should provide a data out hold time of 50 ns on the SFP_I2C_Data pin. The actual hold time is about 16 ns.

Implication:

Timing specification violation. There have been no reports of failures resulting from this timing. Note that the data input hold time required is zero, so the provided output hold time should be more than enough as long as the I2C CLK and DATA signals are reasonably matched on the board.

Workaround:

None.

Status:

A2 NoFix

2. NC-SI Hardware Arbitration Issues

Problem:

1. During normal operation, the I210 might get FLUSH commands with a smaller ID then the device ID. The I210 should pass on the received FLUSH; but it sends it's own ID for ~2 µs and then passes on the lower ID FLUSHes.
2. The time from received-idle (while in a wait idle state) until the I210 sends IDLE on ARB_OUT is 1.7 µs; the maximum time allowed (by the specification) is T9 = 640 ns.
3. If a token timeout occurs while the I210 waits to send an XON packet, the internal state machine is reset and the XON is never sent.
4. Hardware arbitration timeout mechanism stops upon receiving pause packets from the MC. The timer stops counting until the pause indication drops.
5. The I210 sends XON opcode after the end of the Master Assignment process, even if the XOFF time (~300 ms) has expired.

   If the I210 exits the congested mode during the Master Assignment process it sends XON opcode at the end of the Master Assignment process even if the XOFF has expired.

   The I210 doesn't consider the Master Assignment duration in the XOFF expiration time.
6. When the I210 enters congestion mode it sends XOFF opcode and also makes a request for TOKEN in order to send a XOFF frame.

   When the I210 enters congestion mode it should send a XOFF frame if it has the TOKEN or XOFF opcode even if it hasn’t received the TOKEN.

   The I210 shouldn’t send both of them (opcode and frame) in any case.

Implication:

1. No implication in actual operation. Eventually, the lower IDs pass and arbitration succeeds.
2. The issue is not expected to cause problems because the timeout period is longer. Minor NC-SI compliance violation related to hardware arbitration.
3. The MC is released by the XOFF timer expiration. Minor NC-SI compliance violation related to hardware arbitration.
4. Longer than expected timeout (no specification violation).
5. No implication.
6. Slight delay in traffic coming from the MC but no platform implication.

Workaround:

None.

Status:

A2 NoFix

3. SGMII: Counters Incorrectly Increment on Collision

Problem:

In SGMII mode/half duplex, the following statistics counters incorrectly increment when a collision occurs:

Name	Definition	Location
RLEC	Length error counter.	0x4040
CRCERRS	CRC error counter.	0x4000
RFC	Receive frame counter.	0x40A8

Implication:

Error counters might not be accurate.

Workaround:

None.

Status:

A2 NoFix

4. BMC-only Packets Not Counted as Host Sent/Received Packets

Problem:

When OS2BMC is enabled, packets that do not reach the LAN are not counted as packets sent by the host (HGPTC register). Similarly, packets received from the MC are not counted as packets received by the host (RPTHC register)

Implication:

HGPTC and RPTHC counts are not accurate.

Workaround:

Add the O2BGPTC count to the HGPTC count to get the accurate number of packets sent by the host. Add the B2OGPRC count to the RPTHC count to get the accurate number of packets received by the host.

Status:

A2 NoFix

5. Device Off Deadlock

Problem:

If firmware resets (such as due to a parity error) after entering device off, the I210 does not detect the error and should enter device off but not shut the device down.

This happens only after a firmware reset.

Implication:

The chances of such an event happening while moving to device off are minimal.

Workaround:

None.

Status:

A2 NoFix

6. Marginal Low 10 Mb Amplitude

Problem:

1. 10BASE-T amplitude
   On some designs, the I210 might not meet the IEEE specification that states that the 10 Mb peak differential amplitude be between 2.2V and 2.8V for all data sequences.
2. 10BASE-T TP_IDLE mask failures.
   Some designs might have mask failures on the 10BASE-T TP_IDLE with TPM load.
3. 10 BASE-Te (802.3az section 14.10) amplitude.
   On some designs, the I210 might not meet the IEEE specification that states that when 10BASE-Te is enabled the 10 Mb peak differential amplitude be between 1.54V and 1.96V for all data sequences.
4. 10BASE-Te TP_IDLE and link test pulse waveform mask failures.
   Some designs might have mask failures on the 10BASE-Te TP_IDLE and link test pulse with and without TPM load.

Implication:

No implication on system level performance or interoperability, conformance test only impact.

Workaround:

None.

Status:

A2 NoFix

7. Non-monotonic Integrated SVR Ramp

Problem:

On some designs, both the 0.9V and 1.5V SVR show a non-monotonic start up.

Implication:

No functional impact for systems using an internal SVR, because the system is not vulnerable at the specific time that this non-monotonicity occurs.

Workaround:

None.

Status:

A2 NoFix

8. Protocol Offload: Incorrect Response to MLDv2 Queries

Problem:

When offloading the Neighbor Solicitation protocol, the I210 also implements the IPv6 Multicast Listener Discovery Version 2 (MLDv2) protocol to ensure that Neighbor Solicitation packets are received.

The I210 always responds to MLDv2 Multicast Address (and source specific) query packets with a Current State Record indicating EXCLUDE mode and an empty source list. In some circumstances the correct response should indicate INCLUDE mode with a source list.

Implication:

Incomplete implementation of the MLDv2 protocol. Possible intermittent failures of the Neighbor Solicitation protocol offload in a network configuration where the I210 is connected to a router that implements MLDv2 through an access switch that snoops the MLD traffic and forwards multicast packets accordingly.

Workaround:

None.

Status:

A2 NoFix

9. Writes to the VPD RW Area are Not Reliable

Problem:

VPD write accesses via the PCIe VPD Capability Structure are not always stored in the Flash.

Implication:

VPD writes are not reliable.

Workaround:

Do not use a RW area in the VPD structure. RO areas function correctly.

Status:

Fixed in NVM image release 3.16.

10. NC-SI: Get NC-SI Pass-through Statistics Response Might Contain Incorrect Packet Counts

Problem:

The I210 maintains packet counters that are used in the Get NC-SI Pass-through Statistics Response. These counters are cleared by any reset of the port, including the port reset generated by a PCIe reset.

Implication:

If a PCIe reset or port reset has occurred since the previous Get NC-SI Pass-through Statistics Response, the packet count values could be lower than the actual packet counts because the counters were cleared.

Workaround:

The packet counts in the Get NC-SI Pass-through Statistics Response can be used for debug purposes, but they should not be used for maintaining reliable statistics.

Status:

A2 NoFix

11. MCTP Commands From SMBus are Dropped

Problem:

The DMTF MCTP SMBus/I^2C Transport Binding Specification requires that the LSB of the 4th byte of an MCTP over SMBus packet be 1b. Such a packet is dropped by the I210.

Implication:

MCTP over SMBus is not functional since all commands are dropped.

Workaround:

None.

Status:

Fixed in NVM image release 3.16.

2. Software Clarifications

Table 6. Summary or Software Clarifications

Software Clarifications	Status
1. While In TCP Segmentation Offload, Each Buffer is Limited to 64 KB	N/A
2. Serial Interfaces Programmed By Bit Banging	N/A

1. While In TCP Segmentation Offload, Each Buffer is Limited to 64 KB

Clarification:

The I210 supports 256 KB TCP packets; however, each buffer is limited to 64 KB since the data length field in the transmit descriptor is only 16 bits. This restriction increases driver implementation complexity if the operating system passes down a scatter/gather element greater than 64 KB in length. This can be avoided by limiting the offload size to 64 KB.

Investigation has concluded that the increase in data transfer size does not provide any noticeable improvements in LAN performance. As a result, Intel network software drivers limit the data transfer size in all drivers to 64 KB.

Please note that Linux operating systems only support 64 KB data transfers.

For further details about how Intel network software drivers address this issue, refer to Technical Advisory TA-191.

2. Serial Interfaces Programmed By Bit Banging

Clarification:

When bit-banging on a serial interface (such as SPI, I2C, or MDIO), it is often necessary to perform consecutive register writes with a minimum delay between them. However, simply inserting a software delay between the writes can be unreliable due to hardware delays on the CPU and PCIe interfaces. The delay at the final hardware interface might be less than intended if the first write is delayed by hardware more than the second write. To prevent such problems, a register read should be inserted between the first register write and the software delay, i.e. “write”, “read”, “software delay”, “write”.

Note: This page intentionally left blank.