Freescale Semiconductor
Application Note

Document Number: AN3433
Rev. 1, 12/2007

MSC8144 PCI Controller Performance

by *Barbara Johnson*
NCSD DSP Applications
Freescale Semiconductor, Inc.
Austin, TX

The MSC8144 peripheral components interconnect (PCI) controller complies with the *PCI Local Bus Specification*, Rev. 2.2. It uses a 32-bit multiplexed address/data bus that can operate at frequencies up to 66 MHz. It is 3.3 V compatible.

The PCI controller functions as a PCI agent-only and can act as a PCI initiator or a PCI target device. When it functions as an initiator, the four SC3400 cores, the DMA controller, the QUICC Engine™ subsystem, the serial Rapid I/O® controller, and the TDM modules can initiate transactions to the PCI. When it functions as a target, the PCI controller can transfer data to/from the M2, M3, DDR memories or the configuration registers (CCSR).

This application note describes the performance of the MSC8144 PCI controller. It presents results for cases in which the MSC8144 functions as the initiator or as the target device. The document also presents techniques to achieve optimum PCI performance. CodeWarrior™ projects are also provided for the user to perform the measurements.

Contents

1 Performance Testing 2
2 Outbound Performance 2
2.1 PCI Settings 3
2.2 DMA Settings 3
2.3 Write Performance 3
2.4 Outbound Write Considerations 4
2.5 Read Performance 4
3 Inbound Performance 4
3.1 PCI Settings 5
3.2 DMA Settings 5
3.3 Write Performance 5
3.4 Read Performance 6
3.5 Inbound Read Consideration 6
4 Summary 7
4.1 General Recommendations 7
4.2 Recommendations for the MSC8144 as the Initiator 7
4.3 Recommendations for the MSC8144 as the Target 7

freescale™ semiconductor

1 Performance Testing

All MSC8144 PCI controller tests were performed on the MSC8144 application development system (ADS). The MSC8144ADS consists of an MPC8560 PCI host processor that connects to the MSC8144 PCI agent through the 32-bit PCI bus. The MSC8144 is configured so that the core operates at 1 GHz and the internal PCI controller operates at 200 MHz. The external PCI interface operates at 66 MHz.

Performance testing involves two types of data transfers between the MSC8144 and MPC8560:

- *Outbound test*. The MSC8144 initiates the write or read access to/from the MPC8560. The MSC8144 acts as the master device and the MPC8560 acts as the target device.
- *Inbound test*. The MPC8560 initiates the write or the read access to/from the MSC8144. The MPC8560 acts as the master device and the MSC8144 acts as the target device.

Both outbound and inbound tests use the DMA controller for data transfers between memory and the PCI interface. The MSC8144 DMA controller transfers data for the outbound tests, and the MPC8560 DMA controller transfers data for the inbound tests. All tests use a block size of 64 Kbyte data to be transferred between PCI and memory.

The MSC8144 timer is configured to count the number of PCI clock cycles that a data transfer requires to complete. For the inbound tests, the MPC8560 accesses the MSC8144 timer registers through the inbound window that maps to the CCSR register space. The timer increments at a rate of 66 MHz.

With a 66 MHz PCI clock, the maximum theoretical bandwidth for a 32-bit PCI bus is 264 Mbps as shown in Equation 1. Actual performance is measured in terms of the number of bytes transferred at a given time as shown in Equation 2. From these two values, we can derive the PCI usage, which is the ratio of the actual performance to the maximum bandwidth as shown in Equation 3.

Eqn. 1

$$Max\ PCI\ Bandwidth = 66\ MHz \times 32\text{-}bits = 264\ MBps$$

Eqn. 2

$$Actual\ PCI\ Performance = Number\ of\ Bytes/Number\ of\ 66\ MHz\ Cycles$$

Eqn. 3

$$PCI\ Usage = Actual\ PCI\ Performance/Max\ PCI\ Bandwidth$$

2 Outbound Performance

The outbound performance test sets up the MSC8144 as the initiator and the MPC8560 as the target. The test measures both outbound writes to the PCI and outbound reads from the PCI. This section discusses the usage test results presented in Table 1.

Table 1. MSC8144 PCI Outbound Performance

Transaction Type	Internal Pipeline	
	Disabled	Enabled
Outbound Write - Streaming Enabled	0.99	0.99
Outbound Write - Streaming Disabled	0.80	0.80
Outbound Read Line or Read Multiple[1]	0.24	0.27

Note:
[1] Outbound read performance is dependent on the PCI target.

2.1 PCI Settings

The target MPC8560 inbound window is mapped to its DDR memory space, which is configured to be prefetchable memory. Streaming is enabled in the initiator MSC8144 outbound window so that consecutive cache lines are transferred into one PCI transaction. The MSC8144 PCI latency timer is set to the maximum value to guarantee the MSC8144 a minimum amount of time on the bus. When the MSC8144 gains ownership of the bus, the latency timer counts down at the rate of the PCI clock. When it reaches zero, it must release the bus. Therefore, if the latency timer constantly expires before the transfer completes, the MSC8144 must rearbitrate for the bus, which delays completion of the transaction and negatively affects performance. Table 2 shows the PCI settings for both the initiator and target devices.

Table 2. Outbound Test Settings

Device	Initiator/Target	Register	Value	Description
MPC8560	Target	PITAR1	0x00080000	Inbound 0 local base address (DDR)
		PIWAR1[EN] PIWAR1[PF]	1 1	Enable inbound window Enable prefetching
MSC8144	Initiator	POBAR0	0x000E0000	Outbound 0 local base address
		POCMR0[EN] POCMR0[SE]	1 1	Enable outbound window Enable streaming
		LTCR	0xF8	Maximum latency timer value

2.2 DMA Settings

The MSC8144 DMA controller is programmed to transfer a burst size of 64 Kbyte data between the PCI and MSC8144 M2, M3, or DDR memory. For increased performance, the DMA controller should be configured to read through one CLASS port and write through a different CLASS port to minimize contention.

2.3 Write Performance

When streaming is enabled, the MSC8144 achieves over 99 percent usage or approximately 261 Mbps when performing outbound write transactions to the MPC8560. For 100 percent usage, each data phase of a burst transaction is transferred in a single clock cycle.

2.4 Outbound Write Considerations

Errata PCI5 for the MSC8144 mask 0M72C and M3 mask 0M18D states that when the MSC8144 processor acts as the initiator, the PCI controller can lock up when the following scenario occurs:

- Target asserts $\overline{\text{STOP}}$ and $\overline{\text{TRDY}}$ to indicate to the MSC8144 that it will take the last data phase but will not take any more data from the MSC8144.
- MSC8144 as the initiator asserts FRAME and keeps $\overline{\text{IRDY}}$ asserted to indicate that will send one more data phase.
- The last data phase is the first 4 bytes of a 32-byte cache line and the preceding data phase is the last 4 bytes of the preceding cache line.

However, if it can be guaranteed that the target will not assert $\overline{\text{STOP}}$ during the last data phase that crosses a cache line boundary, the problem with the PCI controller locking up can be avoided. Otherwise, streaming should be disabled by setting PCIFCR[MSD] to avoid the problem. When streaming is disabled, the outbound write performance is reduced to 80 percent usage or about 211 Mbps.

2.5 Read Performance

Outbound read performance is approximately 27 percent usage or 71 Mbps when the PCI internal pipeline is enabled by setting C2GPR[PPE]. No difference in outbound read performance is observed when the read line or read multiple command is issued to the target MPC8560. Theoretically, it is possible to achieve full usage depending on the PCI target. However, in this case, the target is the bottleneck. When the target takes more than eight PCI clocks to complete a subsequent data phase, it forces the initiator to terminate the transaction. Restarting the transaction yields a decreased usage.

When the MSC8144 PCI controller functions as an initiator-only, enabling the PCI internal pipeline improves usage. When the PCI internal pipeline is disabled, outbound read performance slightly drops from 27 percent to 24 percent usage or 63 Mbps. Note that the PCI internal pipeline should be enabled only when there is no inbound PCI traffic, which means the MSC8144 PCI controller behaves as an initiator only (PCICCR[BMST] = 1) and does not respond to memory space accesses (PCICCR[MEM] = 0).

Outbound read streaming also improves performance when the MSC8144 PCI controller acts as an initiator-only. The master-delayed read transactions should be disabled so that read transactions are cascaded into a single stream. If the master-delayed read transactions are enabled, the PCI issues a new read transaction after a previous read transaction completes. To enable outbound read streaming, two bits must be set: C2GPR[PMDRD] to disable delayed read transactions and C*n*ACR[PME] to enable priority mask.

3 Inbound Performance

The inbound performance test sets up the MPC8560 as the initiator and the MSC8144 as the target. The test measures both inbound write to the PCI and inbound read from the PCI. This section discusses the usage results presented in Table 3.

Table 3. MSC8144 PCI Inbound Performance

Transaction Type	Internal Pipeline Disabled		
	DDR	M3	M2
Inbound write	0.99		
Inbound read/read line/read multiple. Prefetch disabled.	0.03		
Inbound read multiple. Prefetch enabled.	0.53	0.58	0.74

3.1 PCI Settings

The target MSC8144 inbound windows are mapped to its M2, M3, and DDR memory spaces, which are configured to be prefetchable memory. The MPC8560 PCI latency timer is set to the maximum value to guarantee the MPC8560 a minimum amount of time on the bus. Table 4 shows the PCI settings for both the initiator and target devices.

Table 4. Inbound Test Settings

Device	Initiator/Target	Register	Value	Description
MPC8560	Initiator	POWBAR0	0x00080000	Outbound 0 local base address
		POWBAR1	0x00090000	Outbound 1 local base address
		POWBAR2	0x00091000	Outbound 2 local base address
		POWAR0/1/2[EN]	1	Enable outbound windows 0, 1, 2
		LTCR	0xF8	Maximum latency timer value
MSC8144	Target	PIBAR0	0x000C0000	Inbound 0 local base address (M2)
		PIBAR1	0x000D0000	Inbound 1 local base address (M3)
		PIBAR2	0x00040000	Inbound 2 local base address (DDR)
		PIWAR0/1/2[EN] PIWAR0/1/2[PF]	1 1	Enable inbound windows 0, 1, 2 Enable prefetching

3.2 DMA Settings

In both inbound write and inbound read tests, the MPC8560 DMA controller transfers a burst size of 64 Kbyte data between PCI and local memory.

3.3 Write Performance

The MSC8144 achieves over 99 percent usage or approximately 261 Mbps when the MSC8144 M2, M3, or DDR is the target of an inbound write transaction from the MPC8560. The PCI internal pipeline must be disabled as discussed in Section 2.5.

3.4 Read Performance

With the MSC8144 PCI internal pipeline disabled and prefetching enabled, the MS8144 achieves an inbound read usage of 53 percent from DDR, 58 percent from M3, and 74 percent from M2 memory. Clearly, M2 memory has the lowest latency and should be the preferred memory for inbound reads. The read performance is directly related to the memory windows being prefetchable. When prefetching is disabled, read performance significantly drops to 3 percent usage. Inbound read performance is the same whether the read, read line, or read multiple command is used when prefetching is disabled.

3.5 Inbound Read Consideration

Errata PCI8 for MSC8144 mask 0M72C and M3 mask 0M18D states that coherency problems can result when the following situation occurs:

- The initiator issues a read burst from an MSC8144 inbound window defined as prefetchable.
- The initiator disconnects on the last byte of a 32-byte cache line.
- The MSC8144 PCI controller prefetches the next cache line.
- The initiator issues another read burst transaction from the address of the prefetched data.
- Data is read from the prefetch buffer instead of memory. This data can potentially be dirty data.

Because unused prefetched data is not discarded, any future reads from the address of the prefetched data can result in invalid data if new data is written to memory. One way to avoid this problem is to disable the corresponding inbound window to be non-prefetchable by clearing PIWAR*x*[PF] when the MSC8144 is the target device. With the prefetch disabled, data streaming does not occur because the PCI controller disconnects after the first data phase. As a result, the inbound read usage is only 3 percent or 8 Mbps. The preferred solution is to implement a software work around to guarantee coherency. With this method, the inbound window can be configured as prefetchable to improve usage. Because a read transaction that crosses the half cache line boundary (bytes 16–31) may result in the prefetch of the next line, a speculative read line from the first 4 bytes of the next cache line discards the prefetch buffer, therefore avoiding any potential coherency problems. The suggested software work around is shown in Figure 1.

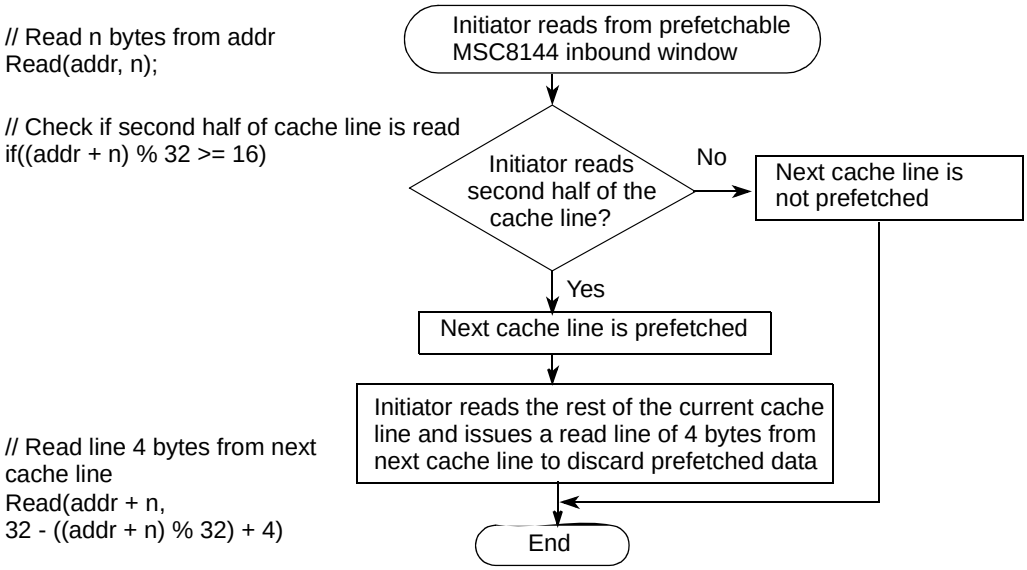

Figure 1. PCI8 Software Work Around

4 Summary

This section summarizes various ways to improve PCI performance.

4.1 General Recommendations

- The MSC8144 clock mode should be selected so that the internal frequency of the PCI controller is set at maximum. For example, a 1 GHz MSC8144 can set its PCI controller frequency up to 200 MHz.
- When the MSC8144 DMA controller is used to transfer data between memory and the PCI interface, it should be configured to read through one CLASS port and write through another CLASS port to minimize contention.
- Because PCI write performance is better than read performance, it is preferable to generate write transactions through the PCI bus. For example, suppose two MSC8144 devices are connected through the PCI bus and the first MSC8144 device needs to read data from second device. Instead of reading data directly using its DMA controller, the first device should configure the DMA controller in the second device to move data from the second device to it. In general, higher performance can be achieved by pushing data out.

4.2 Recommendations for the MSC8144 as the Initiator

- The MSC8144 should use the read multiple command by setting C2GPR[PRCS] when reading consecutive cache lines from prefetchable memory.
- The MSC8144 latency timer should be configured to minimize timeouts during the transaction. This value should also be balanced so that the MSC8144 initiator does not monopolize the bus, preventing other devices from mastering the PCI bus.
- When the MSC8144 is configured as an initiator-only and no inbound traffic occurs, the MSC8144 should enable the PCI internal pipeline by setting C2GPR[PPE] and enable outbound read streaming by setting C2GPR[PMDRD] and C*n*ACR[PME].
- For MSC8144 devices in which errata PCI5 for the MSC8144 mask 0M72C and M3 mask 0M18D applies, the MSC8144 should disable streaming by setting PCIFCR[MSD] when it cannot be guaranteed that the target will not assert $\overline{\text{STOP}}$ during the last data phase of a write transaction that crosses a cache line boundary. Otherwise, streaming should be enabled to improve performance.

4.3 Recommendations for the MSC8144 as the Target

- MSC8144 should set its inbound windows mapped to DDR, M3 and M2 memory to be prefetchable by setting PIWAR*x*[PF] when memory coherency can be guaranteed.
- For MSC8144 devices to which errata PCI8 for the MSC8144 mask 0M72C and M3 mask 0M18D applies, the software work around discussed in this document should be implemented to discard unused data.
- M2 provides lower latency than M3 and DDR memory, so it is the preferred memory for fast PCI read transfers.

How to Reach Us:

Home Page:
www.freescale.com

email:
support@freescale.com

USA/Europe or Locations Not Listed:
Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
1-800-521-6274
480-768-2130
support@freescale.com

Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064, Japan
0120 191014
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:
Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor
@hibbertgroup.com

Document Number: AN3433
Rev. 1
12/2007

freescale™
semiconductor