Application Note

AN1801
Rev. 0.2, 11/2003

Performance Differences between MPC8240 and the Tsi106™ Host Bridge

Top Changwatchai
Roy Jenevein
risc10@email.sps.mot.com
CPD Applications

This paper discusses some of the major performance differences between the MPC8240 memory and PCI interfaces and those of the Tundra Tsi106™ PowerPC™ host bridge. This document compares the MPC8240 (rev. 1.0) to the Tsi106 host bridge in the following areas:

- Section 1, “Architectural Differences,” discusses architectural differences between the two parts that impact performance.
  - — The MPC8240 offers faster memory and PCI buses and features PCI data-streaming capabilities.
  - — The Tsi106 host bridge supports an external L2 cache, which may cause lower memory bus speed.
- Section 2, “Simulation,” presents results from an analytical model for comparing various system configurations that highlight these architectural differences.

The MPC8240 integrated host processor takes advantage of a small yet powerful 32-bit, superscalar MPC603e processor core. This full-featured high-performance processor core is software-compatible with PowerPC microprocessors.

1 Architectural Differences

This section discusses the differences in memory, L2 cache, and the PCI bus interface between the MPC8240 and the Tsi106 host bridge.

1.1 Memory

Both the MPC8240 and the Tsi106 host bridge support three types of dynamic RAM (DRAM):

- fast page mode (FPM)
- extended data out (EDO)
- synchronous DRAM (SDRAM)

Of the three types, SDRAM typically provides the best bus bandwidth utilization. The Tsi106 host bridge supports an 83-MHz SDRAM bus (66 MHz if an L2 cache is present), and the MPC8240 supports a 100-MHz SDRAM bus.

Access latency can be reduced (thereby increasing bus efficiency) if more open pages of memory can be maintained. The Tsi106 host bridge can maintain two open pages, and the MPC8240 can maintain four open pages at once.

DRAM devices are accessed with a multiplexed address scheme. Each unit of data is accessed by first selecting its row address (also known as “opening a page”) and then selecting its column address. This unit of data is transferred on the memory bus in a “beat”. Transferring data in bursts (multiple-beat transfers) can increase bus efficiency and lower data access latency. During a burst transfer, the page and column are accessed for the first beat (just as for a single-beat transfer). For subsequent beats in the burst, the page is kept open, and only the column address changes. Because the page does not have to be reopened during a burst, these subsequent beats have much lower latency. For example, a typical SDRAM may have an 8-1-1-1 access latency, which essentially means eight bus clocks to transfer the first beat of data and one bus clock to transfer each subsequent beat in the burst. (Technically, these access latency numbers refer to the timing of the transfer start (TS_) and transfer acknowledge (TA_) signals on the 60x bus. TS_ signals the beginning of the bus transaction, and each TA_ signals that a data beat transfer completed successfully. The first number is the latency of the first data beat, and refers to the number of bus clocks (inclusive) from TS_ to the first TA_. Subsequent numbers refer to the number of bus clocks to transfer successive beats of the burst. Thus 8-1-1-1 means if TS_ is asserted in clock 1, then TA_s are asserted in clocks 8, 9, 10, and 11.)

Allowing pages to remain open between bursts can also result in lower latency. If a memory access hits in an already open page, its first beat is accessed more quickly. The Tsi106 host bridge can maintain two open pages at once. The MPC8240 can maintain 4. Typical access latencies are presented in Table 1.

Table 1. Typical SDRAM Access Latencies

	Access Closed Page (miss)	Access Open Page (hit)
“Fast” SDRAM	8-1-1-1	6-1-1-1
“Slow” SDRAM	10-1-1-1	7-1-1-1

On the other hand, access latencies may be higher due to several factors, such as closing an already open page or handling error correction code (ECC). If the maximum number of supported open pages are open when an access misses, one page must be closed before the new one is opened, causing the access time for the first beat to increase by about two bus cycles. Supporting more open pages, as the MPC8240 does, reduces the frequency of this occurrence.

The Tsi106 host bridge includes ECC support for FPM and EDO. It does not support ECC for SDRAM, but does allow an external device to provide this support. The MPC8240 supports ECC for FPM, EDO, and SDRAM. When ECC is enabled, access latency to memory is increased. For SDRAM ECC, the latency per beat is increased by one bus cycle.

The analytical model is conservative, does not include the effect of open pages, and assumes that ECC is not used.

1.2 L2 Cache

The Tsi106 host bridge provides support for an external, lookaside L2 cache, the MPC8240 does not. An L2 cache can reduce data access latency by providing faster access to cached data. However, because the L2 cache shares the processor bus with memory, adding an L2 can introduce timing constraints that may lower the supported bus speed. In addition, the L2 tag RAM cannot run faster than 66 MHz, and an Tsi106 host bridge with L2 can support a 66-MHz SDRAM bus, as opposed to the MPC8240 (with no L2), which can support a 100-MHz SDRAM bus.

In addition, cache maintenance overhead, such as L2 castouts, can lower the speed-up of having an L2 cache. The analytical model includes the effect of an L2 cache, but does not take the cache size as a parameter. Instead, it takes the cache miss rate directly as a parameter.

1.3 PCI Bus Interface

Both the MPC8240 and the Tsi106 host bridge implement a 32-bit PCI bus interface, which provides a bridge between the 60x processor bus, memory, and the PCI bus. The MPC8240 offers two improvements:

- faster PCI bus
- PCI data-streaming capabilities

1.3.1 Bus Speed

The Tsi106 host bridge supports a 33-MHz PCI bus, whereas the MPC8240 supports a 66-MHz PCI bus. If the PCI devices can support this speed, PCI accesses should proceed twice as fast.

1.3.2 Data Streaming

For identical PCI bus configurations, the MPC8240 streaming capabilities allow higher PCI data throughput than the Tsi106 host bridge. The MPC8240 has two main streaming features that are not present in the Tsi106 host bridge. As a PCI target, it can support data bursts larger than 32 bytes; as a PCI master, it can support transactions larger than 32 bytes through its on-chip dual-channel DMA engine. By streaming, the MPC8240 can utilize PCI bus bandwidth more efficiently than the Tsi106 host bridge with greater potential data throughput.

1.3.3 No Forced Disconnects Every 32 Bytes

When acting as a non-DMA PCI bus master, both the Tsi106 host bridge and the MPC8240 limit transaction size to 32 bytes. (The PCI specification does not have this limit on transaction size. PCI transaction sizes are realistically limited by internal buffers of the PCI devices or by other device requests and the PCI master latency timer.) When acting as a target, the Tsi106 host bridge also limits transactions to 32 bytes by issuing a PCI disconnect after up to 32 bytes are transferred. To continue transferring data, the mastering device must start a new transaction. The MPC8240, on the other hand, does not have this 32-byte limit. It has the capability to transfer data as long as data can be supplied. (In practice, transaction size is limited by the buffer sizes of the PCI devices. To limit the risk of prefetching into improper memory address spaces, the MPC8240 initiates a disconnect when a transaction crosses a 4096-byte page boundary.)

As a target, the MPC8240 is able to avoid the 32-byte limit on transaction size because it has two internal 32-byte buffers for reads and two buffers for writes. (The Tsi106 host bridge also has two write buffers (only one read buffer), but the internal state machine still forces a disconnect every 32 bytes as a PCI target.) When one buffer is filled from the source (for example, memory), the other buffer can be written to the destination (for example, the requesting PCI device).

Each disconnect causes a delay before the next transaction starts. These delays can reduce data throughput. This disconnect penalty includes sending the disconnect signal, transmitting the new address (because the PCI multiplexes its address/data bus), allowing PCI bus turnaround and bus arbitration and initiating the new transaction. The penalty may be increased if contention occurs with other PCI devices or if the bus is not parked.

The minimum penalty for a disconnect for a PCI write is three PCI bus clocks—two before FRAME is asserted, plus one to transmit the new address. However, because this situation doesn't allow other devices to take the bus, the actual penalty is probably higher depending on the PCI device. Also, PCI reads incur a higher penalty than writes. A one-cycle turnaround is required because the address and data are transmitted in different directions. After the new address is transmitted to the Tsi106 host bridge, some latency occurs

while the address is passed on to the memory device and data is read into the Tsi106 host bridge. As data becomes available on the Tsi106 host bridge, it can then be transferred to the requesting PCI device.

NOTE

Results from the analytical model are presented for a minimum disconnect penalty of five clocks, and for a more typical disconnect penalty of eight clocks.

As a target, the Tsi106 host bridge does support fast back-to-back transactions by the same master, in which the master starts a new transaction immediately, without an idle state. In theory, fast back-to-back PCI writes can incur a single clock penalty between writes just to transmit the new address [1], if the Tsi106 host bridge does not issue a disconnect (transactions must be 32 bytes or smaller). If transactions are disconnected, a disconnect penalty applies.

1.3.4 DMA Controller

The MPC8240 has a software-programmed dual-channel DMA controller that supports the following four types of data transfer:

- PCI-to-PCI
- PCI-to-memory
- memory-to-PCI
- memory-to-memory

For PCI-to-memory and memory-to-PCI transactions, data can be streamed continuously. The MPC8240 does not interrupt data transfer if any data is left to transfer, no errors occur on the PCI or memory buses, and no other intervening transaction occurs, including CPU memory accesses and memory refresh events.

How efficiently these data streaming features can be utilized depends on factors such as the following:

- transaction size that each PCI device can support
- size of each transaction, and if the data can be supplied quickly enough to fill the bus
- whether PCI data accesses tend to be to contiguous addresses
- how well software makes use of the DMA controller

To exploit these features fully, write software with these factors in mind.

2 Simulation

An analytical model was used to explore expected performance of synthetic workloads on current and potential MPC8240 and MPC603e + Tsi106 host bridge system configurations. This section presents the results of this ongoing research and highlights several of the performance implications.

2.1 System Configurations

Several system configurations were compared using the analytical model, which takes system parameters as input and uses a set of formulas to calculate expected run time as output. Table 2 shows the modeled MPC8240 system configurations. Table 3 shows the MPC603e + Tsi106 host bridge configurations. These tables also present run times for the various configurations (lower run time means higher performance), as

[1] PCI reads incur a greater penalty.

well as the normalized rates of execution (higher rate means higher performance). The rates are normalized to configuration K-3a, an MPC8240 with 266-MHz processor clock, 66-MHz memory bus clock, and 33-MHz PCI bus clock (see Table 2).

NOTE

The MPC8240 configurations assume that software has been modified to take advantage of PCI data streaming and DMA features. Other than that, the only differences between the two types of configurations are noted in the tables.

The analytical model assumes that memory timing is 8-x-x-x, or 8 bus clocks to transfer the first (critical) data beat. The model does not take into account the effect of open pages and timing of subsequent data beats in the burst, because of an assumption that the timing of the rest of the system depends mainly on the critical data beat. However, this assumption may not be accurate, particularly for bulk memory transfers. The effect of open pages for some applications can be significant and should be explored further, especially if high data throughput is important.

Two PCI disconnect penalties are modeled: a minimum realistic penalty at five clocks, and a more typical one at eight clocks. However, both of these disconnect penalties are conservative and minimize the advantage of PCI streaming. A higher penalty would favor the MPC8240 and highlight its PCI data-streaming capabilities.

Several factors would increase the PCI disconnect penalty, such as non-parked bus, any delay due to internal signal latency, or any PCI bus contention. The actual disconnect penalty is highly dependent on the workload and transaction types, system configuration, and PCI devices. For example, a disconnect may trigger a device driver interrupt that could take many cycles to resolve.

Behavior parameters for the analytical model were obtained using a logic analyzer on a 266-MHz system and time-stamping all PCI bus activity. The application was a host processor issuing PCI transactions to a 2-D accelerator card to generate 3-D graphics to provide the necessary data size and rates over the PCI bus along with computation load factors. (Based on sample run on test platform: 1,000,000,000 instructions caused 27,838 pages to be written (4096 bytes each), resulting in roughly 10% bus utilization (see later figures where this is varied). To take advantage of PCI streaming, software should set up the MPC8240 DMA processor or command a PCI device to master the transaction.

Table 2. Performance of MPC8240 System Configurations [1]

Name	Processor Clock [2] (MHz)	60x/Memory Bus Clock (MHz)	PCI Bus Clock (MHz)	Relative Performance [3] (Run Time in Seconds, Rate Normalized to Config. K-3a)			
				PCI Disc. [4] = 5 Clocks		PCI Disc. = 8 Clocks	
				Run Time	Rate	Run Time	Rate
K-1a [5]	300	100 (3:1)	33	8.92	1.290	8.96	1.289
K-1b [5]			66	8.46	1.361	8.48	1.362
K-2a	266	90 (3:1)	30	10.05	1.145	10.09	1.145
K-2b			60	9.53	1.208	9.58	1.206
K-3a	266	66 (4:1)	33	11.51	1.000	11.55	1.000
K-3b			66	11.05	1.042	11.07	1.043
K-4a	250	100 (5:2)	33	9.69	1.188	9.73	1.187
K-4b			66	9.23	1.247	9.25	1.249
K-5a	250	83 (3:1)	33	10.52	1.094	10.56	1.094
K-5b			66	10.06	1.144	10.08	1.146
K-6a	200	66 (3:1)	33	12.92	0.891	12.96	0.891
K-6b			66	12.46	0.924	12.48	0.925

1 MPC8240 configurations feature PCI data streaming, 32-bit PCI bus, and no L2 cache.

2 No 250 MHz parts are currently available, so the 250 MHz configurations are 266 MHz parts run at the lower frequency.

3 Lower run time means higher performance. Normalized rate is inverse of run time. Higher rate means higher performance.

4 PCI target disconnects occur every 256 bytes (representative buffer size for PCI device).

5 Configuration K-1 does not correspond to an announced part, and is used for comparative purposes only.

The MPC8240 configurations in Table 2 feature various processor speed/memory bus speed combinations. Each combination is modeled with 33-MHz PCI bus and 66-MHz PCI bus. For the synthetic workload, the resulting run time was computed for each of the two PCI disconnect penalties (five and eight clocks). The rate of execution was calculated by taking the inverse of the run time and normalizing to configuration K-3a. (For example, take K-6a. For PCI disconnect = 5, this configuration's run time is 12.92 seconds, as opposed to 11.51 seconds for K-3a. Dividing 1/12.92 by 1/11.51 yields 0.891. For PCI disconnect = 8, dividing 1/12.96 by 1/11.55 yields 0.891.)

Table 3. Performance of MPC603e + Tsi106 Host Bridge System Configurations [1]

Name	Processor Clock [2]	60x/Memory Bus Clock	L2 Cache [3] Present?	Relative Performance [4](Run time in seconds, Rate Normalized to Config. K-3a)			
	(MHz)	(MHz)		PCI disc. [5] = 5 clocks		PCI disc. = 8 clocks	
				Run time	Rate	Run time	Rate
G-1 [6, 7]	300	100 (3:1)	-	9.39	1.226	9.71	1.189
G-2a	300	66 (9:2)	-	11.47	1.003	11.79	0.980
G-2b			Yes	8.60	1.338	8.92	1.295
G-3a	266	66 (4:1)	-	11.98	0.961	12.30	0.939
G-3b			Yes	9.07	1.269	9.39	1.230
G-4 [6, 7]	250	100 (5:2)	-	10.16	1.133	10.48	1.102
G-5 [7]	250	83 (3:1)	-	10.99	1.047	11.31	1.021
G-6a	250	63 (4:1)	-	12.65	0.910	12.97	0.891
G-6b			Yes	9.57	1.203	9.89	1.168
G-7a	200	66 (3:1)	-	13.39	0.860	13.71	0.842
G-7b			Yes	10.42	1.105	10.75	1.074

1 MPC603e + Tsi106 host bridge configurations feature 33 MHz, 32-bit PCI bus.

2 No 250 MHz parts are currently available, so the 250 MHz configurations are 266 MHz parts run at the lower frequency.

3 L2 cache timing is 3-1-1-1.

4 Lower run time means higher performance. Normalized rate is inverse of run time; higher rate means higher performance.

5 PCI target disconnects occur every 32 bytes.

6 Configurations G-1 and G-4 do not correspond to announced parts and are used for comparative purposes only.

7 Configurations G-1, G-4, and G-5 will not support an L2 cache because 60x bus is greater than 66 MHz.

The MPC603e + Tsi106 host bridge configurations in Table 3 are similar to the MPC8240 configurations in Table 2, except that the PCI bus speed is fixed at 33 MHz. In addition, the configurations with 66 MHz or lower-memory bus speed are modeled with an optional L2 cache. Again, the run time was computed for the two PCI disconnect penalties and the rate of execution was calculated by taking the inverse of the run time and normalizing to configuration K-3a (from Table 2.). (For example, take G-3a. For PCI disconnect = 5, its run time is 11.98 seconds (for K-3a it is 11.51). Dividing 1/11.98 by 1/11.51 yields 0.961. For PCI disconnect = 8, dividing 1/12.30 by 1/11.55 yields 0.939.)

The results in Table 2 and Table 3 are presented graphically in Figure 1 and Figure 2. Figure 1 presents the normalized performance (rates of execution) for PCI disconnect = 5, and Figure 2 presents the normalized performance for PCI disconnect = 8. For each of these figures, the left half represents the MPC8240 configurations (Table 2), and the right half represents the MPC603e + Tsi106 host bridge configurations (Table 3). Each MPC8240 configuration (K-1 through K-6) features two bars representing the 33-MHz PCI (a) and 66-MHz PCI (b) versions (66 MHz is longer). Each MPC603e + Tsi106 host bridge configuration (G-1 through G-7) presents the no-L2 cache ('a') and, where possible, the with-L2 cache ('b') versions.

For example, in Figure 1, the configuration K-3a features a 266-MHz processor clock, 66-MHz memory bus clock, and 33-MHz PCI clock, and the bar length is 1.000. Configuration K-3b is the same except for a 66-MHz PCI clock and bar length of 1.042.

Longer bars mean higher performance. Note that the Y-axis (performance improvement) is centered around the normalized configuration (value = 1). A bar that is twice as long in these figures does not mean the configuration runs twice as fast. For example, in Figure 1., G-7b runs about 10% faster than K-3a, and G-7a runs about 15% slower than K-3a. G-7b, by adding an L2 cache, runs about 30% faster than G-7a.

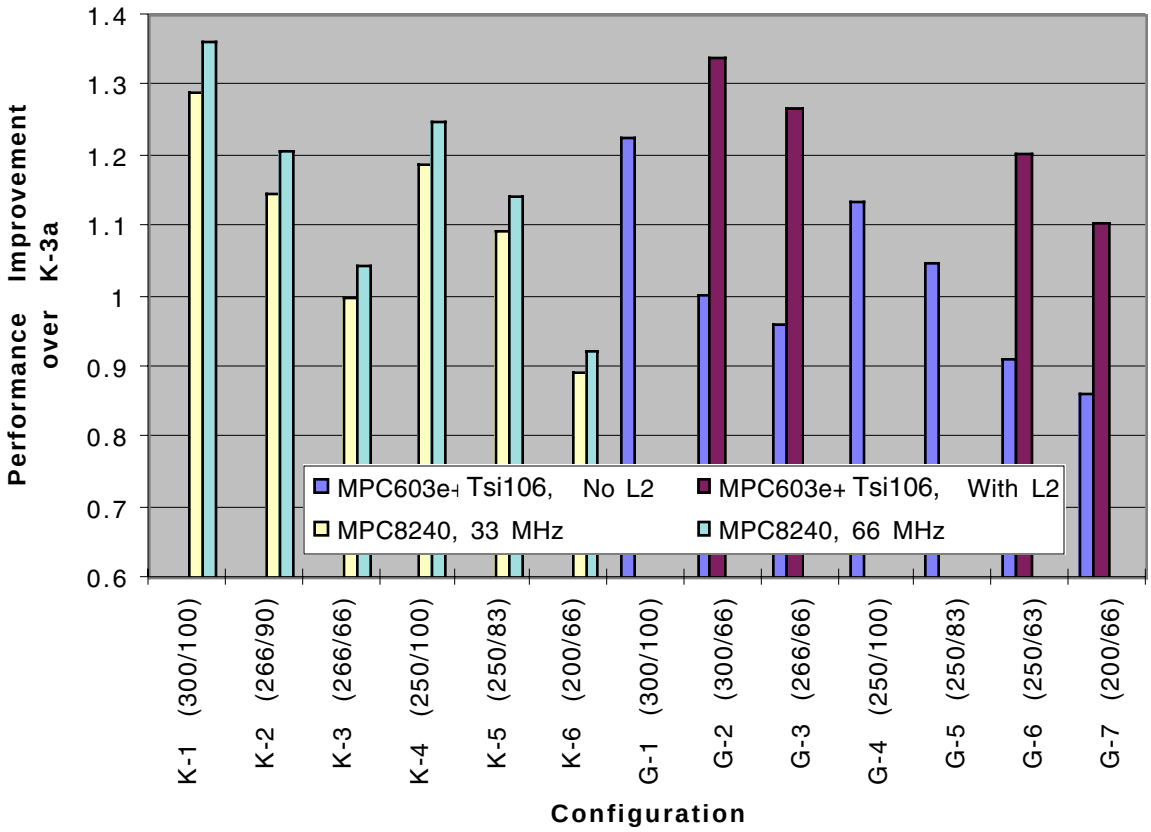

Figure 1. Relative System Performance for PCI Disconnect = 5 (Normalized to Configuration K-3a)

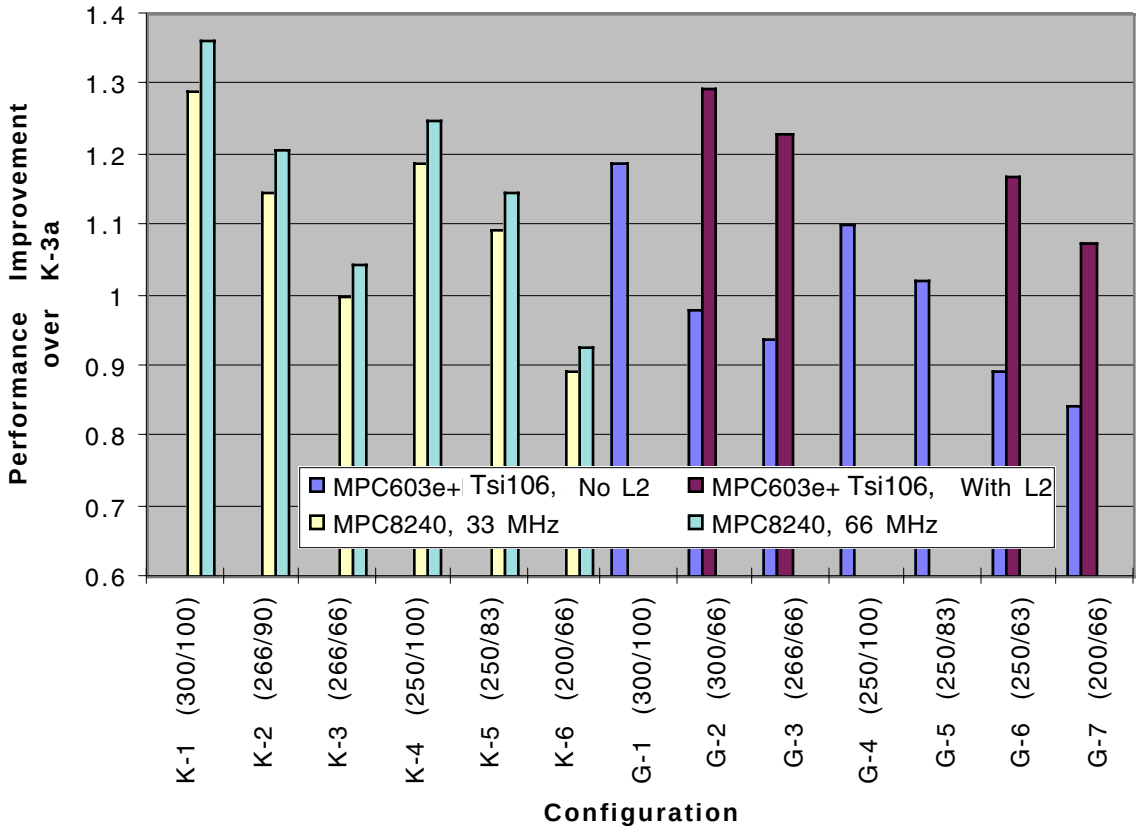

Figure 2. Relative System Performance for PCI Disconnect = 8 (Normalized to Configuration K-3a)

Note several items from this data:

- The workload modeled is memory-bound, not compute-bound. Changing memory bus speed, PCI bus speed, and L2 configuration has more effect than changing core processor speed. For example, moving from K-1 to K-4 (processor 300 MHz -> 250 MHz) does not have as much impact as moving from K-2 to K-3 (memory 90 MHz -> 66 MHz). Similarly, moving from G-2 to G-7 (processor 300 MHz -> 200 MHz) has less effect than moving from G-1 to G-2 (memory 100 MHz -> 66 MHz).
- The workload modeled utilizes the memory bus more than the PCI bus. Varying the memory speed has more effect than varying the PCI bus speed (for the MPC8240 configurations). Adding an L2 cache for the MPC603e + Tsi106 host bridge configurations has a much larger impact than moving from 33 MHz to 66 MHz PCI for the MPC8240 configurations.
- On the other hand, the benefit of adding an L2 cache is reduced, considering that it limits memory bus speed. For example, compare K-4 with G-6b. Both are the fastest configurations modeled for a 250 MHz processor. Even with an L2 cache, G-6b is just comparable in speed to K-4 with PCI disconnect = 5. The MPC8240 with 66 MHz PCI (K-4b) is slightly faster. And with PCI disconnect = 8, even the MPC8240 with 33 MHz PCI (K-4a) beats G-6b. Comparing the fastest 266 MHz processor configurations (K-2 and G-3b) does show the MPC603e + Tsi106 host bridge+L2 configuration leading. The cost and increased board space of adding an L2, however, must be balanced against simply using an MPC8240.
- For the MPC8240 configurations, changing PCI disconnect penalty has little effect on relative performance because PCI data streaming minimizes the effect of the disconnect penalty. The MPC8240 bars are very similar between Figure 1 and Figure 2. For the MPC603e + Tsi106 host

bridge configurations, changing PCI disconnect penalty has a greater effect, with higher disconnect penalty favoring MPC8240 configurations (so the Figure 2 MPC603e + Tsi106 host bridge bars are shorter than those in Figure 1).

The remaining figures present data only for PCI disconnect = 5, the most conservative case.

2.2 PCI Transaction Size

One factor influencing the effectiveness of PCI data streaming for the MPC8240 configurations is the typical transaction size, which the PCI devices and the PCI address access pattern determine. Recall that Table 2 assumed a disconnect every 256 bytes. Figure 3 shows how varying the MPC8240 PCI transaction size affects the performance improvement of the MPC8240 over the MPC603e + Tsi106 host bridge combination for PCI disconnect = 5.

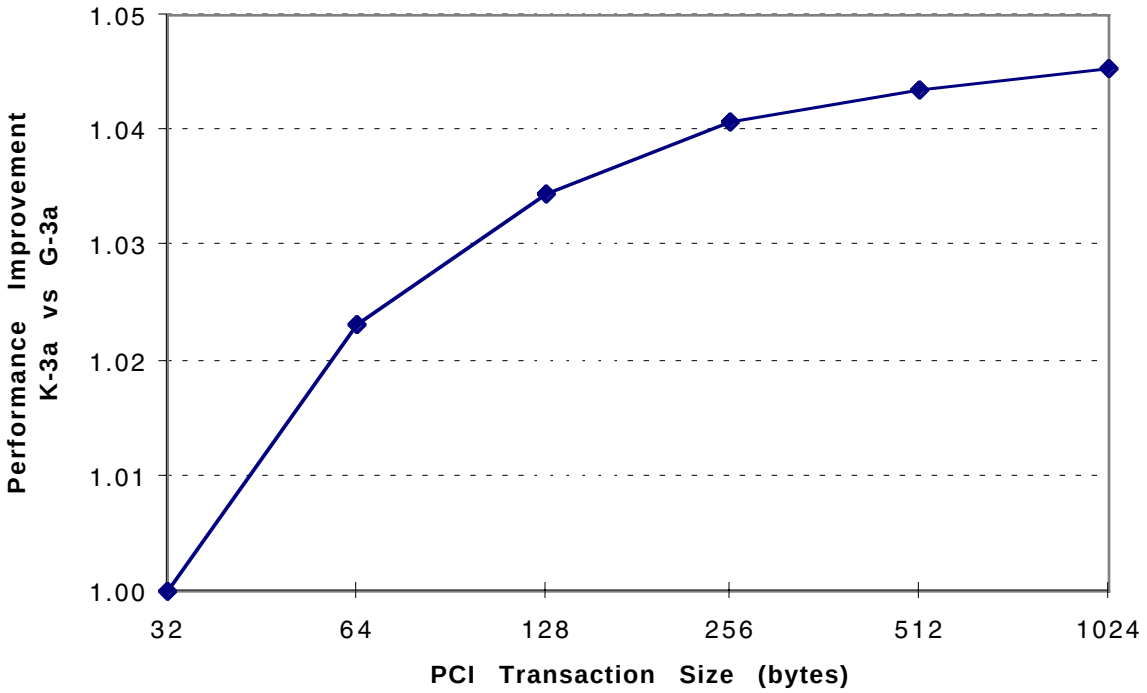

Figure 3. Effect of Varying PCI Transaction Size on Performance Improvement (Configuration K-3a over Configuration G-3a)

Note that most of the speedup occurs by the 256-byte mark, which is close to full streaming (no disconnects at all). Very little additional performance improvement is achieved for larger PCI device buffer sizes.

2.3 PCI Bus Utilization

Another factor influencing the effectiveness of PCI streaming is the amount of traffic that is streamed on the PCI bus. Raising the number of page writes relative to the number of executed workload instructions increases bus utilization, thereby increasing the ratio of PCI traffic to computation. Figure 4 shows how increasing bus utilization increases performance improvement using streaming for K-3a and K-3b over G-3a, and PCI disconnect = 5.

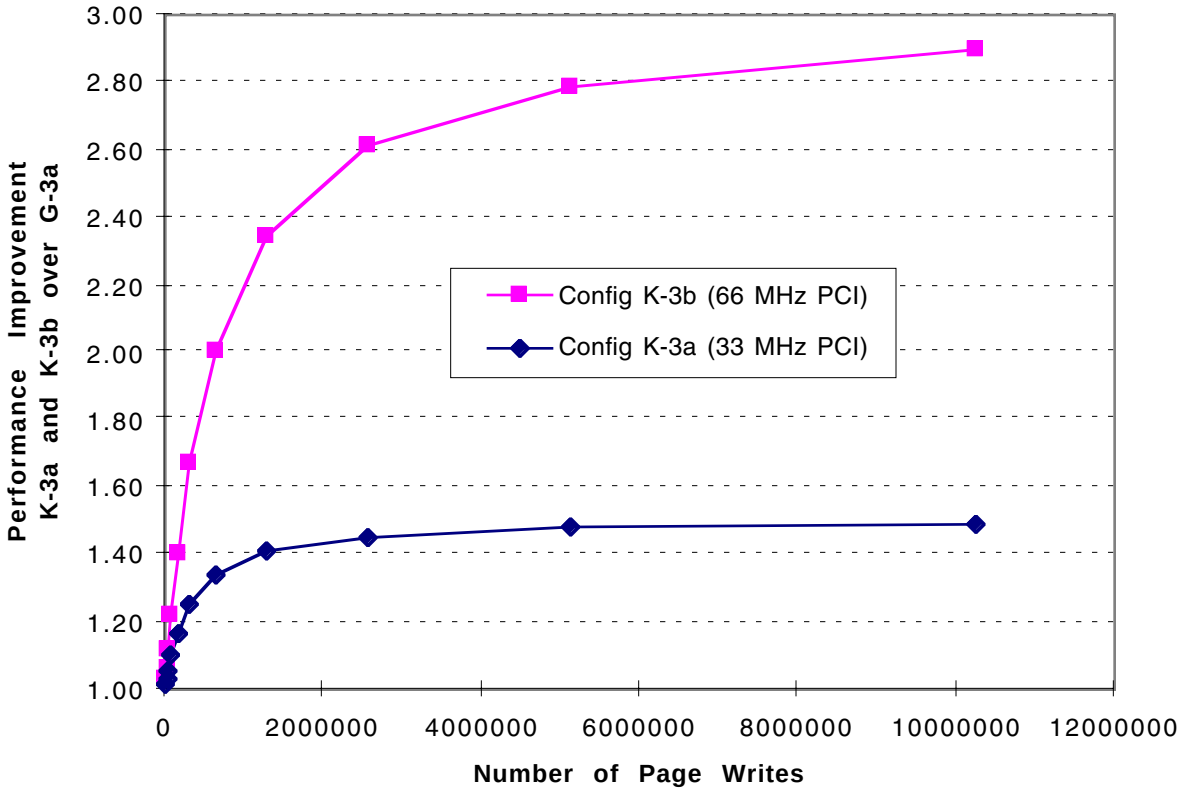

Figure 4. Effect of Varying PCI Bus Utilization (Number of Page Writes) on Performance Improvement (Configurations K-3a and K-3b over Configuration G-3a)

As the PCI bus traffic increases, performance improvement caused by PCI data streaming increases too. Note that as the PCI bus becomes saturated, performance improvement levels off.

Figure 5 presents the same data in a different way. It is identical to Figure 4, except that the X-axis is scaled differently. The number of page writes is converted to the bus utilization for configuration G-3a. Note that these bus utilization numbers do not represent bus utilization figures for configurations K-3a or K-3b.

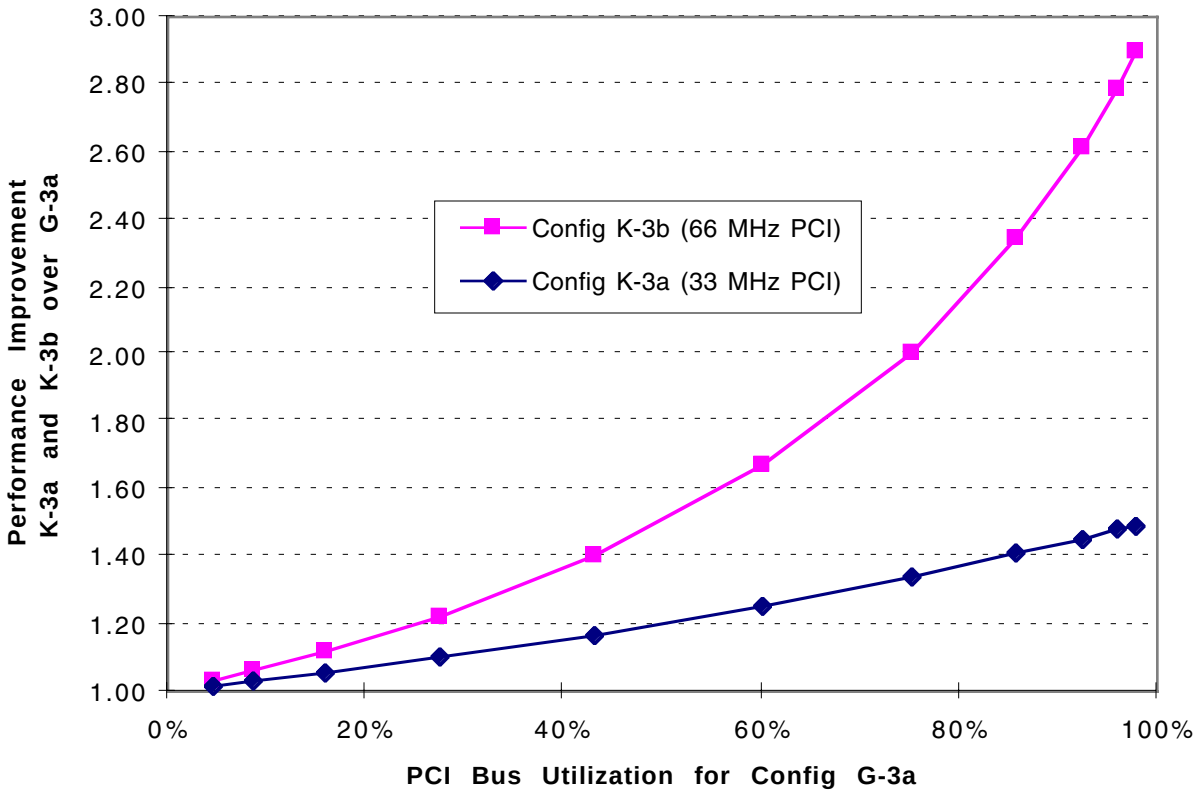

Figure 5. Effect of Varying PCI Bus Utilization on Performance Improvement (Configurations K-3a and K-3b over Configuration G-3a)

As shown in Figure 5, the workloads simulated in Table 2 and Table 3, which are based on observed tests, have a low average bus utilization, and therefore show little improvement using streaming. However, as the number of page writes increases, performance improvement also increases. For 33-MHz PCI, performance improvement increases linearly, while performance improvement increases even faster for 66-MHz PCI.

Figure 6 and Figure 7 show the data transfer rate for comparable configurations K-3a, K-3b, and G-3a, as the number of PCI page writes varies, for PCI disconnect = 5.

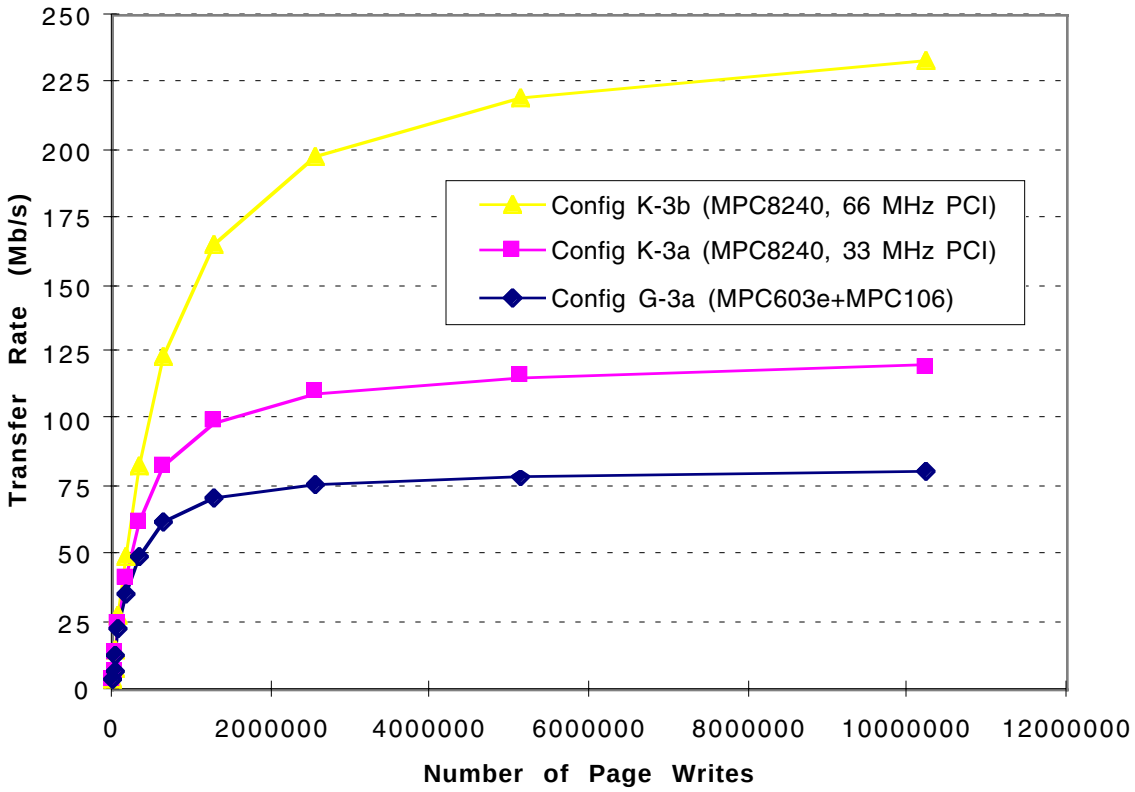

Figure 6. Effect of Varying PCI Bus Utilization (Number of Page Writes) on Data Transfer Rate (Configurations K-3a and K-3b over Configuration G-3a)

Figure 6 shows the maximum transfer rate approached asymptotically as the number of page writes increases. For an MPC603e + Tsi106 host bridge system, up to 81 Mb/s can be transferred, whereas for the MPC8240 the maximum transfer rate is 122 Mbytes per second for 33 MHz PCI, and 244 Mb/s for 66 MHz PCI.

The theoretical maximum transfer rate using the analytical model, as PCI bus utilization approaches 100%, is:

Bus speed x Bus size x t/(t+d)

where

Bus speed = 33 MHz

Bus size = 32 bits (4 bytes)

t is the number of bus clocks between disconnects

d is the disconnect penalty = 5 PCI bus clocks

The MPC8240 has t = 64 (256-byte transaction size÷4 bytes/beat). Without PCI streaming, the Tsi106 host bridge has t = 8 (32-byte transaction size). The theoretical maximum transfer rates are therefore:

MPC8240 - 6666 x 4 x 64/69 = 244 Mb/s

MPC8240 - 3333 x 4 x 64/69 = 122 Mb/s

Tsi10633 x 4 x 8/13 = 81 Mb/s

This calculation assumes a lack of contention from other PCI devices.

These numbers do not include any other system considerations such as concurrent memory accesses from the CPU or other internal delays.

Figure 7 is identical to Figure 6, except that the X-axis is scaled differently. The number of page writes is converted to the bus utilization for configuration G-3a. Note that these bus utilization numbers do not represent bus utilization figures for configurations K-3a or K-3b.

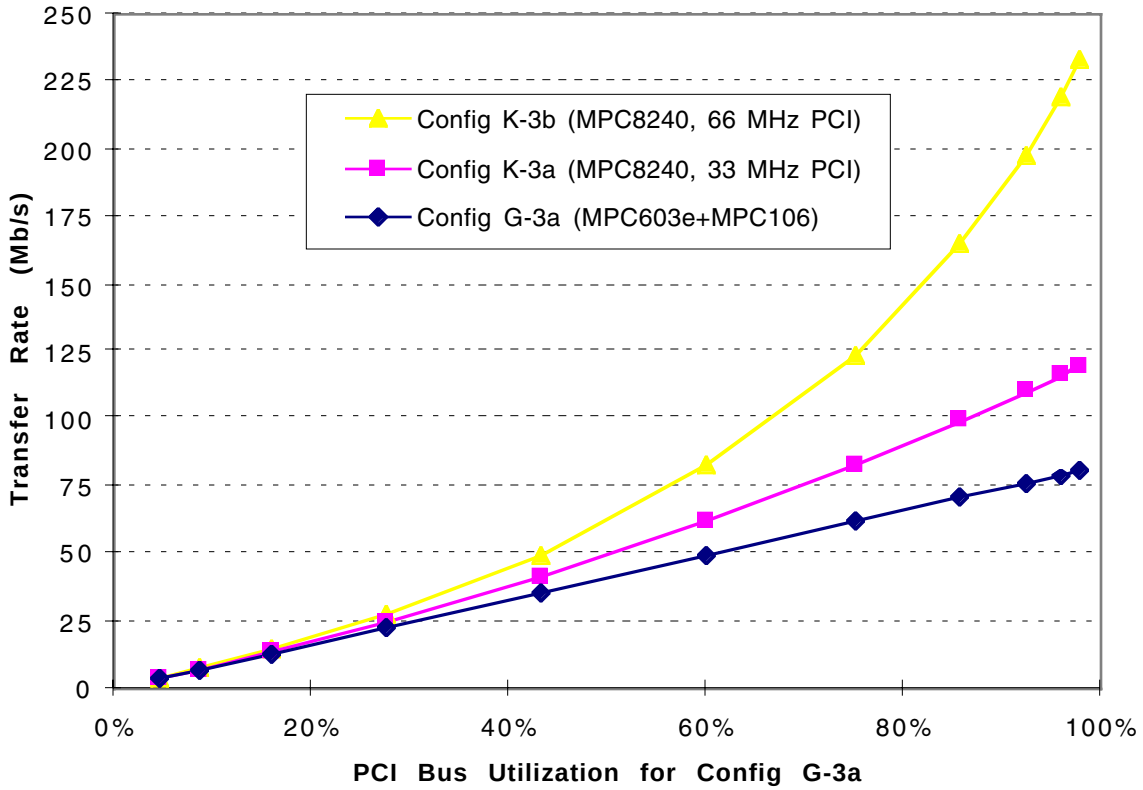

Figure 7. Effect of Varying PCI Bus Utilization on Data Transfer Rate (Configurations K-3a and K-3b over Configuration G-3a)

Figure 7 also shows the maximum transfer rate as PCI bus utilization for G-3a approaches 100%. For periods of low PCI bus activity, PCI streaming may give little improvement. However, for periods with high PCI bus activity, marked improvement is evident in the data transfer rate (as shown by the widening gap between the K-3a and K-3b configurations and the G-3a configuration), because the disconnect penalty begins to have more of an impact on throughput.

3 Conclusion

Although the MPC8240 and an MPC603e + Tsi106 host bridge combination are similar in many respects, architectural differences between the two can lead to measurable performance differences. The MPC8240 supports higher memory and PCI bus speeds, as well as PCI data streaming, while the Tsi106 host bridge supports an L2 cache.

For the workloads modeled, the MPC8240 in general can provide better performance for comparable system configurations, particularly during heavy PCI bus traffic. Adding an L2 cache can significantly improve the

performance of an Tsi106 host bridge system, making it comparable in speed to an MPC8240 system, though at a higher cost.

4 Revision History

Table 4 shows the revision history of this document.

Table 4. Revision History

Revision Number	Changes
0.0	Initial release
0.1	Nontechnical reformatting
0.2	Nontechnical reformatting

How to Reach Us:

Home Page:
www.freescale.com

E-mail:
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 or +1-480-768-2130
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 or +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 or 303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

freescale™ semiconductor