

White Paper

MPC74XXBUSWP
Rev. 1.1, 11/2003

Memory Bus Throughput
of the MPC74xx



Richie David,
CPD Applications
risc10@email.mot.com

This document compares the memory bus data transfer rate of the MPC7450 in 60x and MPX bus modes. Devices in the MPC74xx family should perform similarly.

This document covers the following topics:

| Part | Page |
|----------------------------------|-------------|
| Section 1, "Overview" | 1 |
| Section 2, "Performance Factors" | 2 |
| Section 3, "Throughput Test" | 2 |
| Section 4, "Conclusion" | 4 |
| Section 5, "Revision History" | 5 |

1 Overview

In embedded applications, the I/O throughput is a common bottleneck in a processor's performance. *Throughput* is the amount of data transferred from one place to another in a specified amount of time. This white paper examines the system bus throughput of the MPC7450 PowerPC™ family of processors and how the bus mode, pipelining, order of transaction execution, and the system controller affect it.

A program that creates a repeatable data transfer to measure the bus throughput was run under Dink32, a flexible software tool that enables evaluation and debugging of the 32-bit microprocessors that implement the PowerPC architecture. This program consists of a simple loop of consecutive loads that represent a large data transfer from SDRAM. These tests were run on two platforms: an MVP reference board with a MPC7450 processor and a Marvell GT64260A bridge chip, and a Sandpoint reference board with a MPC7410 processor and a Tundra Tsi107™ PowerPC host bridge. Both systems run with a 100-MHz memory bus, 3/1/1/1 CAS latency, and enabled instruction and data caches.

The data transfer rates are calculated using the following equation:

$$\text{transfer rate} = \text{bytes/sec} = [(\text{bytes transferred})/(\text{cycles})] \times (\text{cycles/sec})$$

where

Memory bus speed = 100 MHz

Data bus width = 64 bits = 8 bytes

Data transferred in a burst = 32 bytes

2 Performance Factors

The following sections discuss performance factors.

2.1 MPX vs. 60x Bus Mode

One of the main factors that was tested was the performance difference between the MPX and 60x bus modes. A number of key features of the MPX bus improve performance over the 60x bus. The MPX bus supports address and data bus streaming. No requirement exists for a dead cycle between consecutive bus tenures. MPX bus mode allows data tenures to be executed out of order with respect to their corresponding address tenures to help hide the latency of slow peripherals by allowing the data from a fast device to avoid waiting for the slow device and return data as soon as it is ready. This re-ordering is not supported by the 60x bus. The actual effects of these enhancements on performance are measurable.

The 60x bus can transfer 8 bytes of data in 1 clock cycle. However, with the dead cycle between data tenures, each transaction, a 32-byte burst, takes 5 clock cycles. Therefore, the theoretical maximum throughput of the 60x bus with a 100-MHz memory bus speed is 640 MBytes/sec.

$$\text{Transfer rate} = (32 \text{ bytes}/5 \text{ clock cycles}) \times 100 \text{ MHz} = 640 \text{ MBytes/sec}$$

The key feature of the MPX bus protocol supports data streaming, which eliminates the dead cycle between data tenures for transfers from the same master to the same target. Therefore, the theoretical maximum throughput of the MPX bus with a 100 MHz memory bus speed is 800 Mbytes/sec.

$$\text{Transfer rate} = (32 \text{ bytes}/4 \text{ clocks}) \times 100 \text{ MHz} = 800 \text{ MBytes/sec}$$

2.2 Pipelining

Pipelining is a technique of separating transactions into address and data tenures so that a subsequent transaction's address tenure can begin before a previous data tenure is complete. The number of outstanding transactions allowed on the bus can have great impact on bus throughput. The MPC7450 processor can support up to 16 additional address tenures before the previous data tenure is complete. The MPC7410 supports up to seven outstanding transactions.

2.3 Bridge Chip Differences/Limitations

Differences between the GT64260A and the Tsi107 host bridge also have an impact on overall bus performance. Neither the GT64260A nor the Tsi107 host bridge can currently operate at a memory bus speed of 133 MHz, and are limited to a memory bus speed of 100 MHz. However, a Tsi107 host bridge with a memory bus speed of 133-MHz will be available soon. Although the MPC7450 can support up to 16 outstanding transactions, the GT64260A can pipeline only eight transactions deep and the Tsi107 host bridge can pipeline only one transaction deep. The eight outstanding transactions supported by the GT64260A are constrained to four read and four write transactions. Finally, the Tsi107 host bridge does not support MPX bus mode, and is unable to take advantage of data streaming or out-of-order transactions. However, the GT64260A can support both MPX and 60x bus modes.

3 Throughput Test

This section describes the actual transfer rates with respect to bus mode, pipelining, out-of-order data tenures, and the bridge chips as measured on a MPC7450 multiprocessing verification platform (MVP) and a MPC7410 Sandpoint system.

3.1 MPC7450 with the GT64260A Bridge Chip

The Marvel GT64260A supports 60x and MPX bus modes and an 8-level deep pipeline. While in MPX bus mode, the MVP platform results are described in the following sections.

3.2 60x Bus Mode

With the MVP operating in a 60x bus mode, an initial memory latency of 19 cycles is observed. *Memory latency* is the time from assertion of \overline{TS} to the assertion of the first \overline{TA} for a non-pipelined transaction. For pipelined transactions, memory latency is measured from the cycle after the last \overline{TA} of the previous data tenure to the first \overline{TA} of the current tenure. During the test loop, a sustained four-transaction-deep pipeline allows a sustained maximum throughput of 640 MBytes/sec. The idle cycle between data tenures accounts for the pipelined memory latency of 2 cycles. The theoretical maximum sustainable throughput of 640 MBytes/sec is obtainable in a real system, as shown in Figure 1.

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---------|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|---|
| Cycles | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | |
| Address | 1 | | | | | 2 | | | | | 3 | | | | | 4 | | | | | 5 | | | | | 6 | | | | |
| Data | | | | | | | | | | | | | | | | | | | | | 1 | 1 | 1 | 1 | X | 2 | 2 | 2 | 2 | X |

Transfer rate = (32 bytes/5 clock cycles) x 100 MHz = 640 MBytes/sec

Figure 1. 60x Bus Transactions

3.3 MPX Bus Mode

In MPX bus mode, data streaming is the main performance advantage over a 60x bus mode. A four-beat burst transfer takes only 4 clock cycles. Data can be returned on every clock cycle. During the data transfer, an initial memory latency of 20 cycles was observed. With out-of-order data tenures disabled through GT64260A configuration register settings, the initial memory latency is equal to 60x bus mode at 19 cycles. An additional one-cycle penalty with out-of-order data tenures is enabled because \overline{DBG} is not parked in this case. As with 60x bus mode, during the test loop a sustained four-transaction deep pipeline exists. However, the GT64260A is not able to stream back-to-back transactions; a dead cycle in between data tenures still exists. Instead of seeing one cycle for the first beat due to data streaming, two cycles are required, as in 60x. This equates to an observed 640 MBytes/sec data transfer rate, as shown in Figure 2.

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---------|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|---|
| Cycles | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | |
| Address | 1 | | | | | 2 | | | | | 3 | | | | | 4 | | | | | 5 | | | | | 6 | | | | |
| Data | | | | | | | | | | | | | | | | | | | | | 1 | 1 | 1 | 1 | X | 2 | 2 | 2 | 2 | X |

Transfer rate = 100 MHz x 8 bytes per beat x (4 beats/5 clock cycles) = 640 MBytes/sec

Figure 2. MPX Bus Transactions

3.4 Pipelining Disabled

With bus pipelining disabled, only one outstanding transaction can be on the bus at any particular time. This feature is enabled or disabled through configuration registers. Therefore, the long initial memory latency is seen for each transaction instead of being masked by pipelined transactions, which has a major effect on bus throughput. In 60x bus mode, it takes 14 cycles for the first beat of data and 1 cycle for each of the three remaining beats of the burst. Therefore, a four-beat burst transfer takes 17 clock cycles, which equates to a data transfer rate of 188 MBytes/sec. In MPX bus mode, the first beat of data took 13 cycles, which causes

a data transfer rate of 200 MBytes/sec. This test is performed with reads from SDRAM. When accessing a slower device, the effects of not pipelining on performance are even greater. From a performance standpoint, pipelining should always be enabled.

$$\text{Transfer rate (60x bus mode)} = (32 \text{ bytes}/17 \text{ cycles}) \times 100 \text{ MHz} = 188 \text{ MBytes/sec}$$

$$\text{Transfer rate (MPX bus mode)} = (32 \text{ bytes}/16 \text{ cycles}) \times 100 \text{ MHz} = 200 \text{ MBytes/sec}$$

3.5 Out-Of-Order Data Tenures

MPX bus mode allows data tenures to be performed out of order with respect to their corresponding address tenures, which allows a read from a fast device to return data when it is ready, instead of waiting for a slower device that is addressed first. To test out-of-order data tenures, a load from a PCI device is included along with the loads from SDRAM. This case is not intended to test maximum throughput because a read from a slow device has a negative effect on performance. However, the ability of GT64260A to perform out-of-order data tenures is verified. With the same series of SDRAM reads with a read from a slower PCI device, transactions are pipelined. Subsequent transactions return data when ready and need not wait for the PCI device to return its data. In these tests, it takes more than 200 cycles for the PCI device to return data. In the meantime, several SDRAM transactions can complete out-of-order data tenures, which is clearly an advantage when using multiple devices of varying speeds. However, when moving a large amount of data from one device (SDRAM, for example), this feature does not show a significant performance benefit.

3.6 MPC7410 with the Tsi107 Host Bridge

This section describes the observed bus performance of the MPC7410 when using the Tsi107 host bridge. The Tsi107 host bridge does not support MPX bus mode, and supports a one-level deep pipeline: only one outstanding and one pending transaction. Thus, the Tsi107 host bridge cannot take advantage of the benefits of data streaming or out-of-order data tenures. During the large data transfer, the initial memory latency is 9 cycles, but due to the one-level-deep pipeline, the overall throughput suffers. For pipelined transactions, the memory latency is 2 cycles, a rate that cannot be sustained because further pipelining is lacking and the transactions are similar to the initial memory latency. This pattern is repeated for the entire data transfer. Therefore, it takes 14 cycles for every two bursts of data, which results in a throughput of 457 MBytes/sec. Clearly, the data transfer rate suffers due to the limited pipelining ability of the Tsi107 host bridge, as shown in Figure 3.

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---------|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|---|
| Cycles | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | |
| Address | 1 | | | | | 2 | | | | | | | | | 3 | | | | | | 4 | | | | | | | | | 5 |
| Data | | | | | | | | | 1 | 1 | 1 | 1 | X | 2 | 2 | 2 | 2 | X | X | X | X | X | X | 3 | 3 | 3 | 3 | X | 4 | 4 |

$$\text{Transfer rate} = (64 \text{ bytes}/14 \text{ cycles}) \times 100 \text{ MHz} = 457 \text{ MBytes/sec}$$

Figure 3. Tsi107 Host Bridge Bus Transactions

4 Conclusion

The maximum sustainable throughput of a MPC7450 system using the Marvel GT64260A bridge chip with pipelining enabled is 640 MBytes/sec. For a large data transfer from one device, using MPX bus over 60x bus gives no advantage. However, when the transfers are from devices of varying speeds, the MPX bus mode's support for out-of-order data tenures gives better performance. To achieve the theoretical maximum

throughput of MPX bus mode of 800 MBytes/sec, the bridge chip must stream back-to-back data tenures to allow data to be transferred on every clock cycle. Currently, Marvell, the manufacturer of the GT64260A device, states that the best chance to see data streaming is with out-of-order transactions. Data streaming is being addressed in their future device. Pipelining and out-of-order data tenures clearly increase performance. However, further performance can be realized on MPC7450 systems when a bridge chip that can run the memory bus at higher frequencies and handle data streaming in all cases is available.

Table 1 shows maximum sustainable memory bus throughput at 100-MHz bus speed.

Table 1. Maximum Sustainable Memory Bus Throughput at 100-MHz Bus Speed

| Platform | Description | Theoretical Throughput (MBytes/sec) | Measured Throughput (MBytes/sec) |
|-----------|---|-------------------------------------|----------------------------------|
| MVP | 60x bus with pipelining | 640 | 640 |
| MVP | MPX bus with pipelining | 800 | 640 |
| MVP | MPX bus non-pipelined | — | 200 |
| Sandpoint | Tsi107 host bridge 60x bus with 1-deep pipeline | 457 | 457 |

5 Revision History

Table 2 shows the revision history of this document.

Table 2. Revision History

| Revision Number | Changes |
|-----------------|---------------------------|
| 0.0 | Initial release |
| 1.1 | Nontechnical reformatting |



THIS PAGE INTENTIONALLY LEFT BLANK



THIS PAGE INTENTIONALLY LEFT BLANK

Freescale Semiconductor, Inc.



Freescale Semiconductor, Inc.

HOW TO REACH US:

USA/EUROPE/LOCATIONS NOT LISTED:

Motorola Literature Distribution
P.O. Box 5405, Denver, Colorado 80217
1-480-768-2130
(800) 521-6274

JAPAN:

Motorola Japan Ltd.
SPS, Technical Information Center
3-20-1, Minami-Azabu Minato-ku
Tokyo 106-8573 Japan
81-3-3440-3569

ASIA/PACIFIC:

Motorola Semiconductors H.K. Ltd.
Silicon Harbour Centre, 2 Dai King Street
Tai Po Industrial Estate, Tai Po, N.T., Hong Kong
852-26668334

TECHNICAL INFORMATION CENTER:

(800) 521-6274

HOME PAGE:

www.motorola.com/semiconductors

Information in this document is provided solely to enable system and software implementers to use Motorola products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Motorola reserves the right to make changes without further notice to any products herein.

Motorola makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Motorola assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Motorola data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Motorola does not convey any license under its patent rights nor the rights of others. Motorola products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Motorola product could create a situation where personal injury or death may occur. Should Buyer purchase or use Motorola products for any such unintended or unauthorized application, Buyer shall indemnify and hold Motorola and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Motorola was negligent regarding the design or manufacture of the part.



Motorola and the Stylized M Logo are registered in the U.S. Patent and Trademark Office. digital dna is a trademark of Motorola, Inc. The described product is a PowerPC microprocessor. The PowerPC name is a trademark of IBM Corp. and used under license. TUNDRA, the Tundra logo, Tsi106, Tsi107, and Silicon Behind the Network are all trademarks of Tundra Semiconductor Corporation. All other product or service names are the property of their respective owners. Motorola, Inc. is an Equal Opportunity/Affirmative Action Employer.

© Motorola, Inc. 2003

MPC74XXBUSWP

**For More Information On This Product,
Go to: www.freescale.com**