

# TMS320TCI6484 Memory Access Performance

Communication Infrastructure

Brighton Feng, Jane Lu

## Abstract

The TMS320TCI6484 has 32KB L1D SRAM, 32KB L1P SRAM and 2MB L2 SRAM. A 32-bit 667MHz DDR2 SDRAM interface is provided on the DSP to support up to 512MB of external memory.

Memory access performance is very critical for software running on the DSP. On the TCI6484 DSP, all the memories can be accessed by DSP cores and multiple DMA masters.

The DSP core is capable of performing up to 128 bits of load/store operations per cycle. When accessing L1D SRAM, the DSP core can access the memory at up to 16GB/second at a 1GHz core clock frequency.

The DMA switch fabric, which provides the interconnection between the C64x+ core (and their local memories), external memory, the EDMA controllers, and on-chip peripherals, has an access port to each end point, which is capable of sustaining up to 5.333GB/second at a 1GHz core clock frequency. There are six EDMA transfer controllers that can be programmed to move data concurrently between any memory endpoints on the device.

This document gives designers a basis for estimating memory access performance and provides measured performance data achieved under various operating conditions. Some factors affecting memory access performance are discussed.

## Contents

1	Introduction	3
2	DSP Core vs. EDMA3 and IDMA for Memory Copy	6
3	Performance of DSP Core Access Memory	8
3.1	Performance of DSP Core Access L2	8
3.2	Performance of DSP Core Access External DDR2 Memory	9
4	Performance of DMA Access Memory	15
4.1	DMA Transfer Overhead	15
4.2	EDMA Bandwidth vs. Transfer Flexibility	15
4.2.1	First Dimension Size (ACNT) Considerations, Burst Width	15
4.2.2	Two Dimension Considerations, Transfer Optimization	16
4.2.3	Index Consideration	17
4.2.4	Address Alignment	18
5	References	19

## Tables

Table 1	Theoretical Bus Bandwidth of DSP Core, IDMA and EDMA .....	6
Table 2	Maximum Throughput of Different Memory Endpoints on TCI6484 VDB .....	6
Table 3	Transfer Bandwidth Comparison Between DSP Core, EDMA and IDMA .....	6
Table 4	EDMA Transfer Overhead .....	15
Table 5	IDMA Transfer Overhead .....	15

## Figures

Figure 1	TMS320TCI6484 Memory System .....	4
Figure 2	DSP Core Access L2 .....	8
Figure 3	DSP Core Reads from DDR2 .....	10
Figure 4	DSP Core Writes to DDR2 .....	12
Figure 5	TCI6484 vs. TCI6488 for Cacheable DSP Core Write on DDR2 Memory .....	13
Figure 6	DSP Core Read vs. Writes to Non-Cacheable DDR2 .....	14
Figure 7	Effect of ACNT Size on EDMA Bandwidth .....	16
Figure 8	Linear 2D Transfer .....	17
Figure 9	Index Effect on EDMA Bandwidth .....	18

## 1 Introduction

The TMS320TCI6484 has:

- 32KB L1D (Level 1 Data) SRAM, which runs at the DSP Core speed, and can be used as normal data memory or cache.
- 32KB L1P (Level 1 Program) SRAM, which runs at the DSP Core speed, and can be used as normal program memory or cache.
- 2MB L2 (Level 2) unified SRAM, which runs at the DSP Core speed divided by two, and can be used as normal memory or cache for both data and program.

A 32-bit 667MHz DDR2 SDRAM interface is provided on the DSP to support up to 512MB of external memory, which can be used as data or program memory.

Memory access performance is critical for software running on the DSP. On the TCI6484 DSP, all the memories can be accessed by the DSP core and multiple DMA masters.

Each TMS320C64x+ core has the ability to sustain up to 128 bits of load/store operations per cycle to the level-one data memory (L1D), and is capable of handling up to 16GB/second. When accessing data in the level-two (L2) unified memory or external memory, the access rate depends on the memory access pattern and cache.

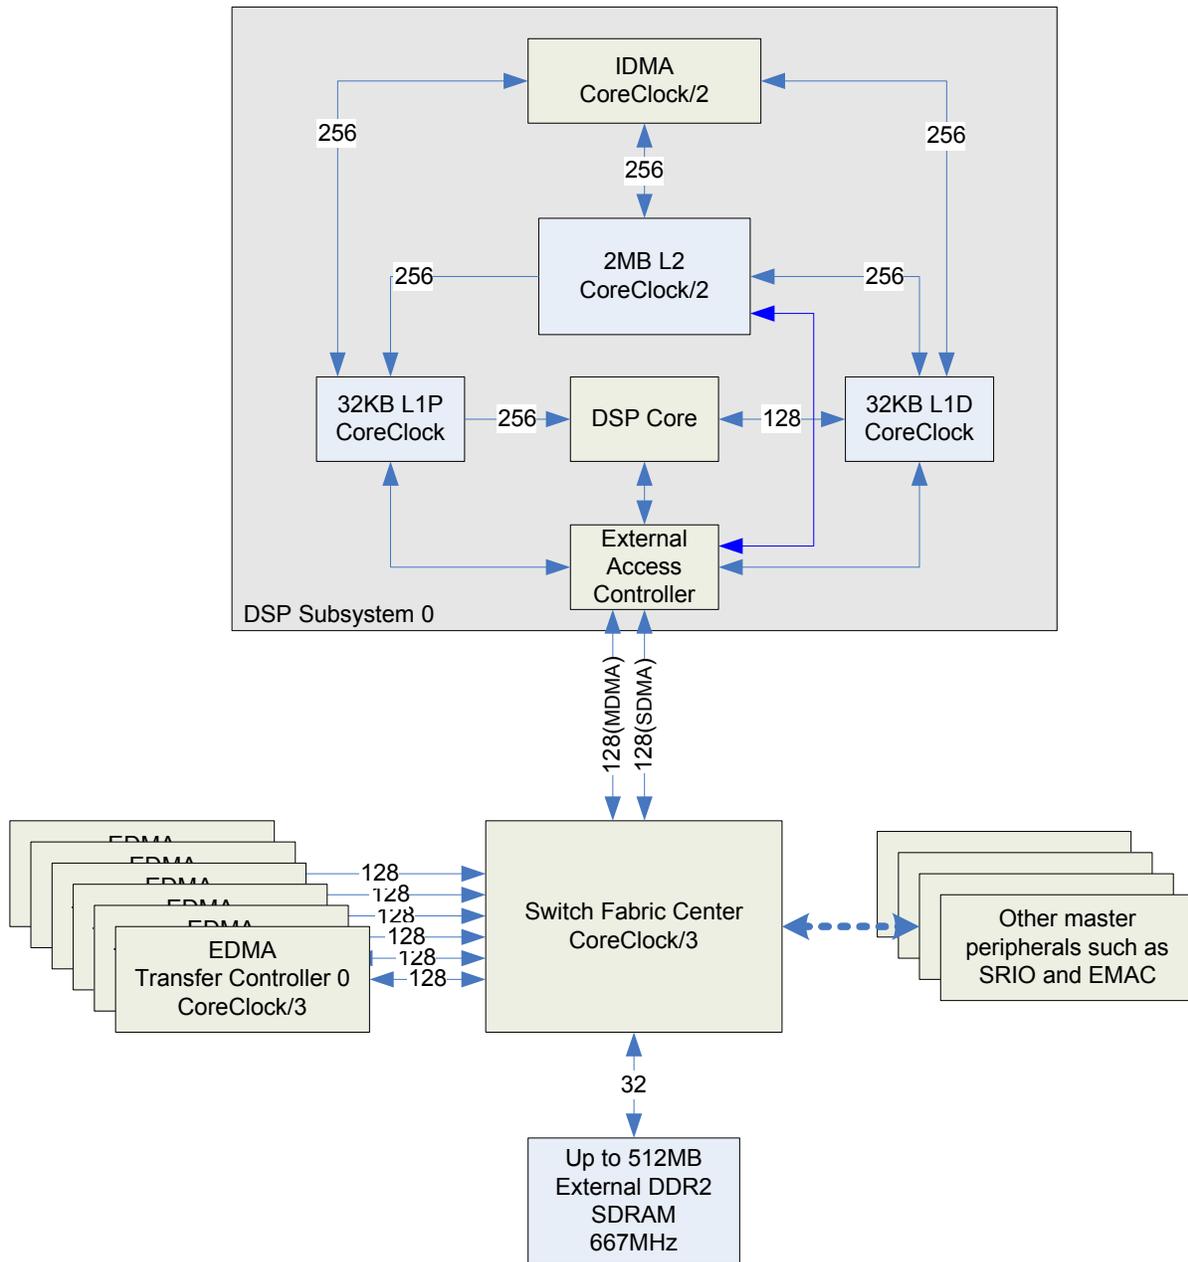
There is an internal DMA (IDMA) engine that can move data at a rate of the DSP Core speed divided by two, capable of handling up to 16GB/s, in the background of DSP core activity (i.e. data can be brought in to buffer A while the DSP core is accessing buffer B). The IDMA can only transfer data between level-one (L1), level-two (L2), and a peripheral configuration port — it cannot access external memory.

The DMA switch fabric, which provides the interconnection between the C64x+ cores (and their local memories), external memory, the enhanced DMA v3 (EDMA3) controllers, and on-chip peripherals, has a 128-bit access bus to each end point, runs at the DSP core frequency divided by three. Therefore, in theory, it is capable of sustaining up to 5.333GB/second at a 1GHz core clock frequency.

There are six EDMA TC (Transfer Controllers) that can be programmed to move data concurrently in the background of DSP core activity, between the on-chip level-one (L1) memory, level-two (L2) memory, external memory, and the peripherals on the device. Each has a 128-bit data bus, and each transfer engine is capable of handling up to 5.333GB/second of data throughput at a core rate of 1GHz. The EDMA3 architecture has many features designed to facilitate simultaneous multiple high-speed data transfers. With a working knowledge of this architecture and the way in which data transfers interact and are performed, it is possible to create an efficient system and maximize the bandwidth utilization of the EDMA3.

[Figure 1](#) shows the memory system of the TMS320TCI6484. The number on the line is the bus width. Most modules run at CoreClock/n. The DDR2 can run up to 667MHz.

**Figure 1** TMS320TCI6484 Memory System



This document gives designers a basis for estimating memory access performance, and provides measured performance data achieved under various operating conditions. Most of the tests operate under best-case situations to estimate maximum throughput that can be obtained. The transfers described in this document serve as a sample of interesting or typical performance conditions.

Some factors affecting memory access performance are discussed in this document, such as access stride, index, conflict, etc.

This document is helpful for analyzing the following common questions:

- Should I use the DSP core or DMA for data copy?
- How many cycles will be consumed for my function with many memory accesses?

Most of the performance data in this document is examined on the 1GHz TCI6484 VDB (Verification and Debug Board) with 32-bit 667MHz DDR2 memory.

## 2 DSP Core vs. EDMA3 and IDMA for Memory Copy

The bandwidth of memory copy is limited by the worst of following three factors:

- Bus bandwidth
- Throughput of source endpoint
- Throughput of destination endpoint

[Table 1](#) summarizes the theoretical bandwidth of the C64x+ core, IDMA, and EDMA on a 1GHz TCI6484.

**Table 1 Theoretical Bus Bandwidth of DSP Core, IDMA and EDMA**

Master	Maximum bandwidth MB/s	Comments
C64x+ core	16000	(128 bits)/(8 bit/byte)*1000M= 16000MB/s
IDMA	16000	(256 bits)/(8 bit/byte)*(1000M/2) = 16000MB/s
EDMA TC0~5	5333	(128 bits)/(8 bit/byte)*(1000M/3)= 5333MB/s

[Table 2](#) summarizes the theoretical throughput of different memory endpoints on a 1GHz TCI6484 VDB with 32-bit 667MHz DDR2 external memory.

**Table 2 Maximum Throughput of Different Memory Endpoints on TCI6484 VDB**

Memory	Maximum Bandwidth MB/s	Comments
L1D	32000	(256 bits)/(8 bit/byte)*(1000M) = 32000MB/s
L1P	32000	(256 bits)/(8 bit/byte)*(1000M) = 32000MB/s
L2	16000	(256 bits)/(8 bit/byte)*(1000M/2) = 16000MB/s
DDR2	2666.7	(32 bits)/(8 bit/byte)*(666.7M)=2666.7MB/s

[Table 3](#) shows the transfer bandwidth measured for linear memory block copy with EDMA, IDMA, and DSP Core for different scenarios. The copied block size for this test for L1D is 8KB, for L2 is 256KB, and for DDR2 is 128MB. The bandwidth is measured by taking the total bytes transferred and dividing by the time used.

**Table 3 Transfer Bandwidth Comparison Between DSP Core, EDMA and IDMA**

Bandwidth (MB/s) for Src-> Dst	DSP core	EDMA	IDMA
L1D-> L1D (16KB L1D cache)	7606	3406	3762
L2-> L1D (16KB L1D cache)	3163	3406	7160
L1D-> L2 (16KB L1D cache)	7798	3773	9082
L2-> L1P (16KB L1P cache)	N/A	3406	7167
L2 -> L2 (32KB L1D cache)	2284	3542	3416
L2 -> DDR2 (non-cacheable)	431	2563	N/A
L2 -> DDR2 (32KB L1D cache)	431		N/A
L2 -> DDR2 (32KB L1D, 256KB L2 cache)	455		N/A
DDR2 -> L2 (non-cacheable)	94	2395	N/A
DDR2 -> L2 (32KB L1D cache)	545		N/A
DDR2 -> L2 (32KB L1D, 256KB L2 cache)	740		N/A

**Table 3 Transfer Bandwidth Comparison Between DSP Core, EDMA and IDMA**

Bandwidth (MB/s) for Src-> Dst	DSP core	EDMA	IDMA
DDR2 -> DDR2 (non-cacheable)	67	990	N/A
DDR2 -> DDR2 (32KB L1D cache)	196		N/A
DDR2 -> DDR2 (32KB L1D, 256KB L2 cache)	291		N/A
DDR2 -> DDR2 (non-cacheable, different DDR2 bank)	73	990	N/A
DDR2 -> DDR2 (32KB L1D cache, different bank)	225		N/A
DDR2 -> DDR2 (32KB L1D, 256KB L2 cache, diff bank)	315		N/A

Generally speaking, the DSP core accesses internal memory efficiently, while using the DSP core to access external data is a bad use of resources and should be avoided. The IDMA is good at linearly moving a block of data in internal memory (L1D, L1P, L2), but it cannot access external memory. The EDMA3 should be given the task of transferring data to/from external memory.

The cache configurations dramatically affect the DSP core performance, but do not affect EDMA and IDMA performance. All test data for the DSP core in this application note are based on cold cache, i.e., all the caches are flushed before the test.

The DDR2 bank architecture affects the performance slightly for the DDR2->DDR2 case. If the source and destination are in the same bank, the DDR2 row switch happens frequently. Every row switch introduces extra delay cycles; if the source and destination are in a different bank, the frequency of row switch is reduced, thus improving the performance.

The above EDMA throughput data is measured on TC0 (Transfer Controller 0). The other five transfer controllers have the same performance on the TCI6484.

### 3 Performance of DSP Core Access Memory

L1 runs at the same speed as the DSP core, so the DSP core can access L1 memory one time per cycle. For special applications that require accessing a small data block very quickly, part of the L1 can be used as normal RAM to store the small data block.

Normally, L1 is used as cached. If a cache hit happens, the DSP core can access data in one cycle. If a cache miss happens, the DSP core stalls until the data comes into the cache.

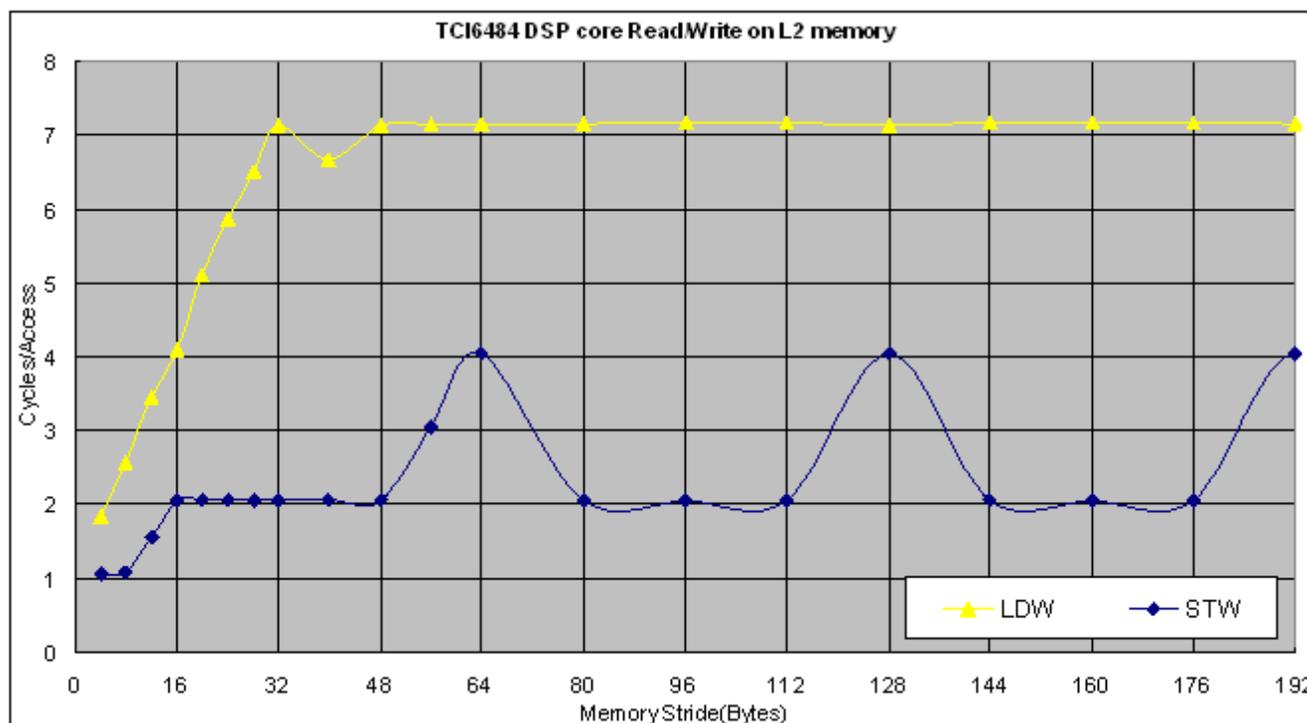
The following sections examine the access performance for DSP Core accesses of L2 and external DDR2 memory. The pseudo code for this test appears as follows:

```
flushCache();
preCycle= getTimeStampCount();
for(i=0; i< accessTimes; i++)
{
    Access Memory at address;
    address+= stride;
}
cycles = getTimeStampCount()-preCycle;
cycles/Access= cycles/accessTimes;
```

#### 3.1 Performance of DSP Core Access L2

Figure 2 shows data collected from 1GHz TCI6484 VDB. The time required for 1024 consecutive LDW (LoaD Word) or STW (STore Word) instructions was measured, and the average time for each instruction is reported. The cycles for LDB/STB and LDDW/STDW are the same as for LDW/STW.

Figure 2 DSP Core Access L2



Since the L1D is a read-allocate cache, a DSP core read of L2 should always go through L1D cache. So, a DSP core access of L2 highly depends on the cache. The address increment (or memory stride) affects cache utilization. Contiguous accesses utilize cache to the fullest. A memory stride of 64 bytes or more causes every access to miss in the L1 cache because the L1D cache line size is 64 bytes.

Since the L1D is not a write-allocate cache, and the cache is flushed before the test, any write to the L2 goes through the L1D write buffer (4x16bytes). For a write operation, if stride is less than 16 bytes, several writes may be merged into one write to the L2 in the L1D write buffer, thus achieving an efficiency close to 1 cycle/write. When the stride is a multiple of 64 bytes, every write always access to the same bank of L2 (because the L2 is organized as 4 x 16 byte banks), which requires 4 cycles. For other strides, the consecutive writes access to different banks of L2, they can be overlapped with pipeline, which requires 2 cycles only.

### 3.2 Performance of DSP Core Access External DDR2 Memory

DSP core access of external DDR2 memory highly depends on the cache. When the DSP core accesses external memory spaces, a TR (transfer request) may be generated (depending on whether the data are cached) to the Switch Fabric Center. The TR will be for one of the following:

- a single element - if the memory space is non-cacheable
- an L1 cache line - if the memory space is cacheable and the L2 cache is disabled
- an L2 cache line - if the memory space is cacheable and L2 cache is enabled

No transfer request is generated in the case of an L1 or L2 cache hit.

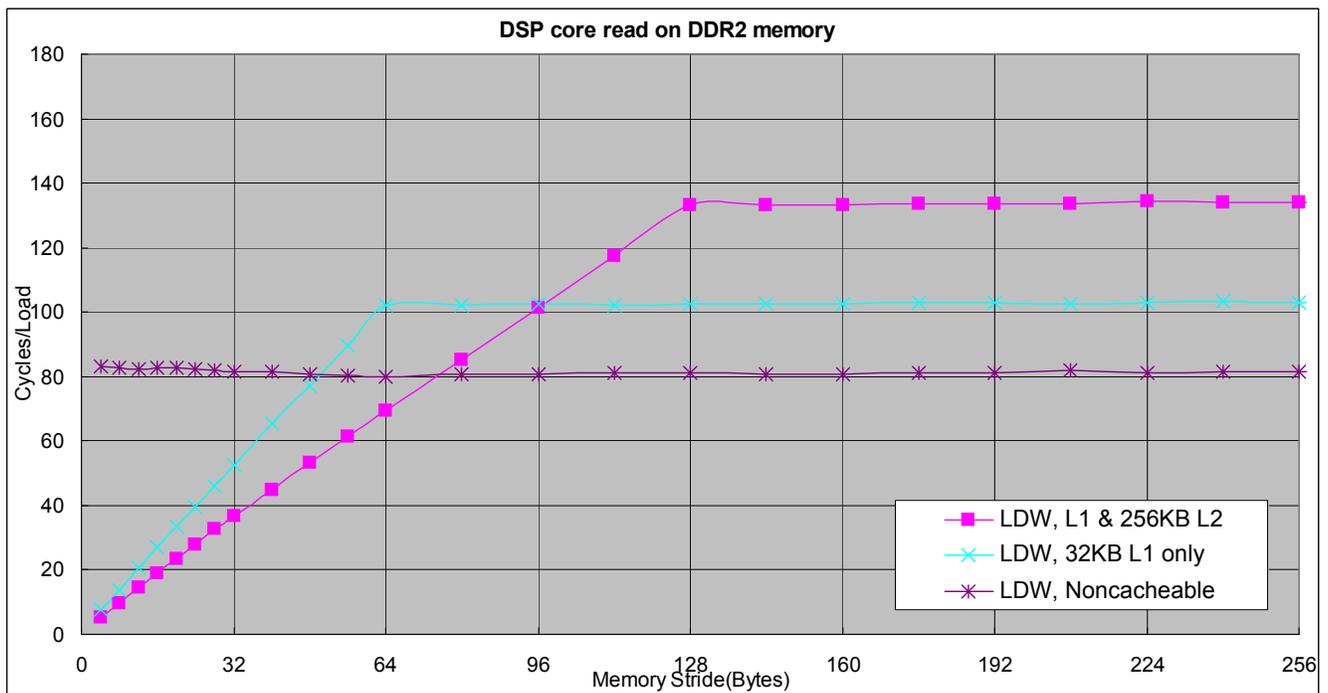
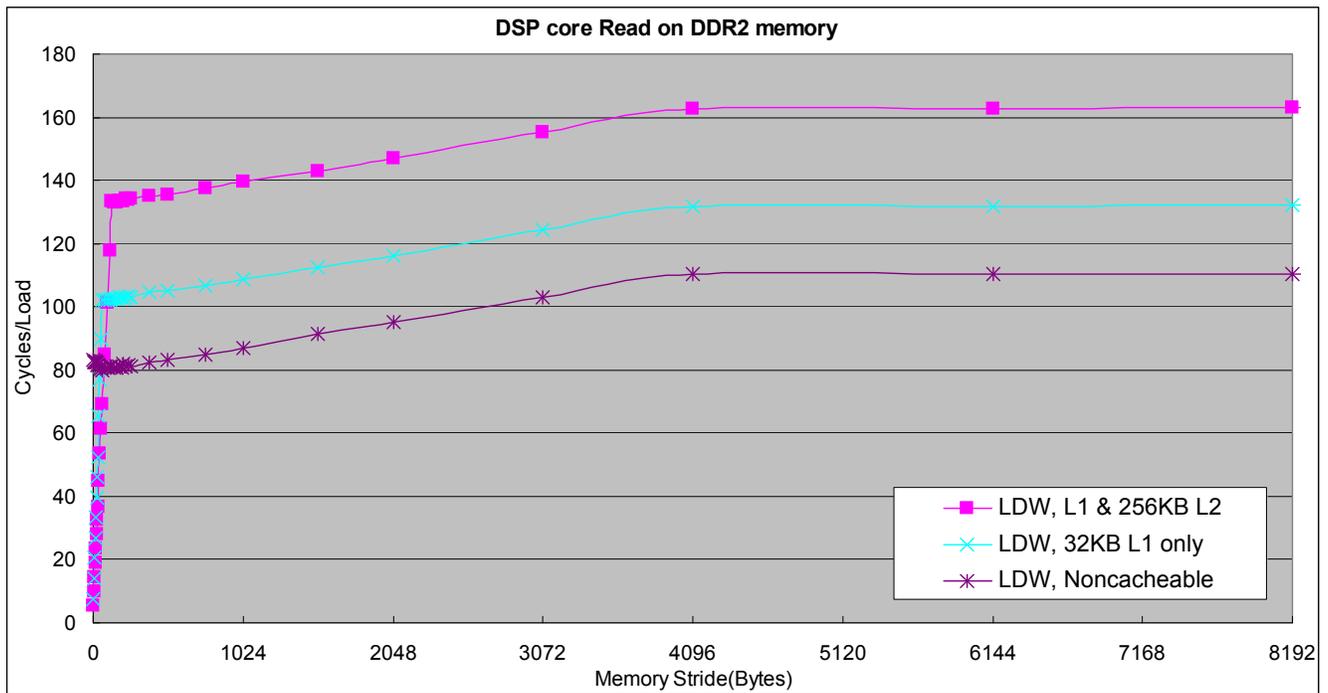
An external memory can be cached by L1 cache, L2 cache, or neither. If the appropriate MAR bit for a memory space is not set, it is not cacheable. If the MAR bit is set and L2 cache size is zero (all L2 is defined as SRAM), the external memory space is cached by L1. If the MAR bit is set and L2 cache size is greater than 0, the external memory space is cached by L2 or L1.

The address increment (or memory stride) affects cache utilization. Contiguous accesses utilize cache memory to the fullest. A memory stride of 64 bytes or more causes every access to miss in the L1 cache because the L1 line size is 64 bytes. A memory stride of 128 bytes causes every access to miss in L2 because the L2 line size is 128 bytes.

If cache miss happens, the DSP core will stall, waiting for the return data.

[Figure 3](#) shows data collected from 1GHz TCI6484 VDB with 32-bit 667MHz DDR2. The time required for 1024 LDW (LoaD Word) instructions was measured, and the average time for each instruction is reported. The cycles for LDB, and LDDW are the same as LDW. Please note that the second figure enlarges the left part of the first figure.

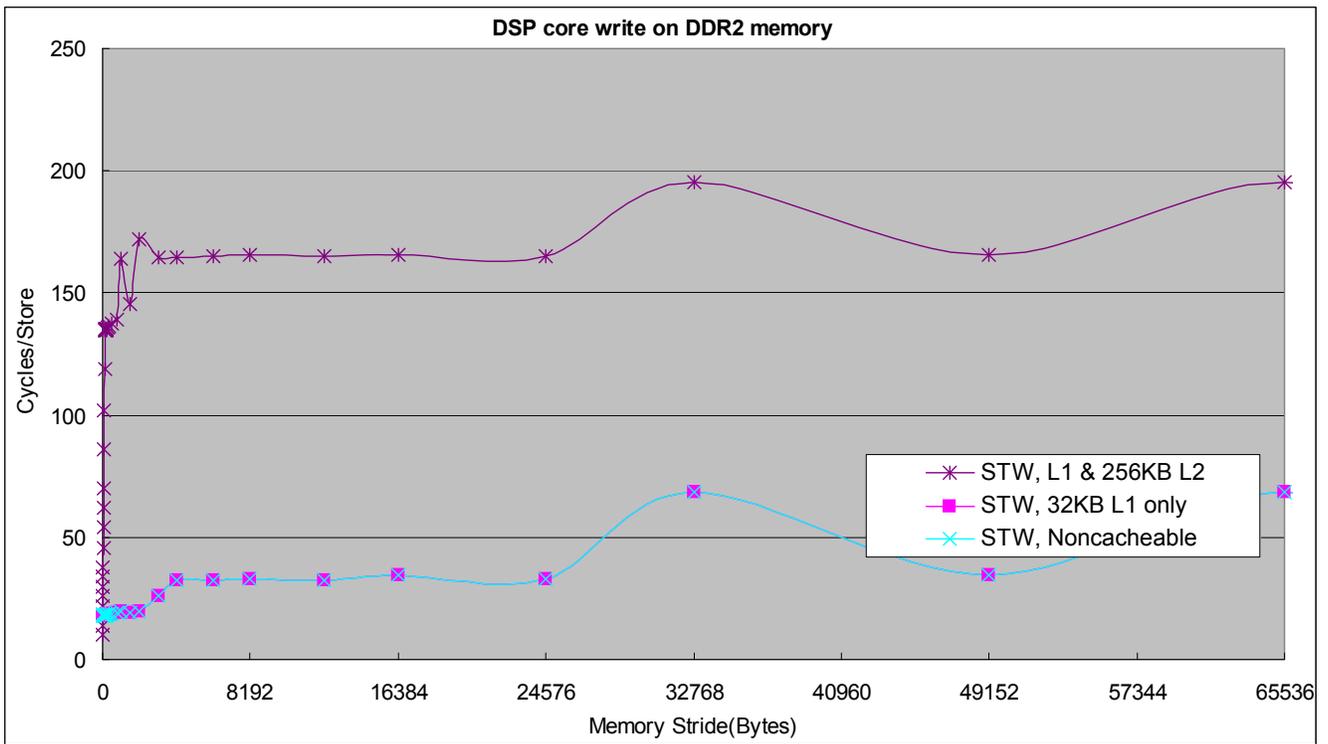
**Figure 3 DSP Core Reads from DDR2**



For a memory stride of less than 128 bytes, the performance is dominated by cache as discussed above. For memory stride larger than 128, the performance becomes worse because of DDR2 SDRAM row switch. The row size or bank width on the TCI6484 VDB is 4096 bytes, so for a stride larger than 4096, and every read access to a new row, the row switch results in about 30 extra cycles. Please note that the DDR2 SDRAM row switch overhead may be different for different DDR2 SDRAM.

[Figure 4](#) shows data collected from a 1GHz C64x+ writing to 32-bit, 667MHz external DDR2 memories. Please note that the second figure enlarges the left part of the first figure. The time required for 1024 STW (STore Word) instructions was measured, and the average time for each instruction is reported. The cycles for STB, and STDW are the same as for STW.

Figure 4 DSP Core Writes to DDR2



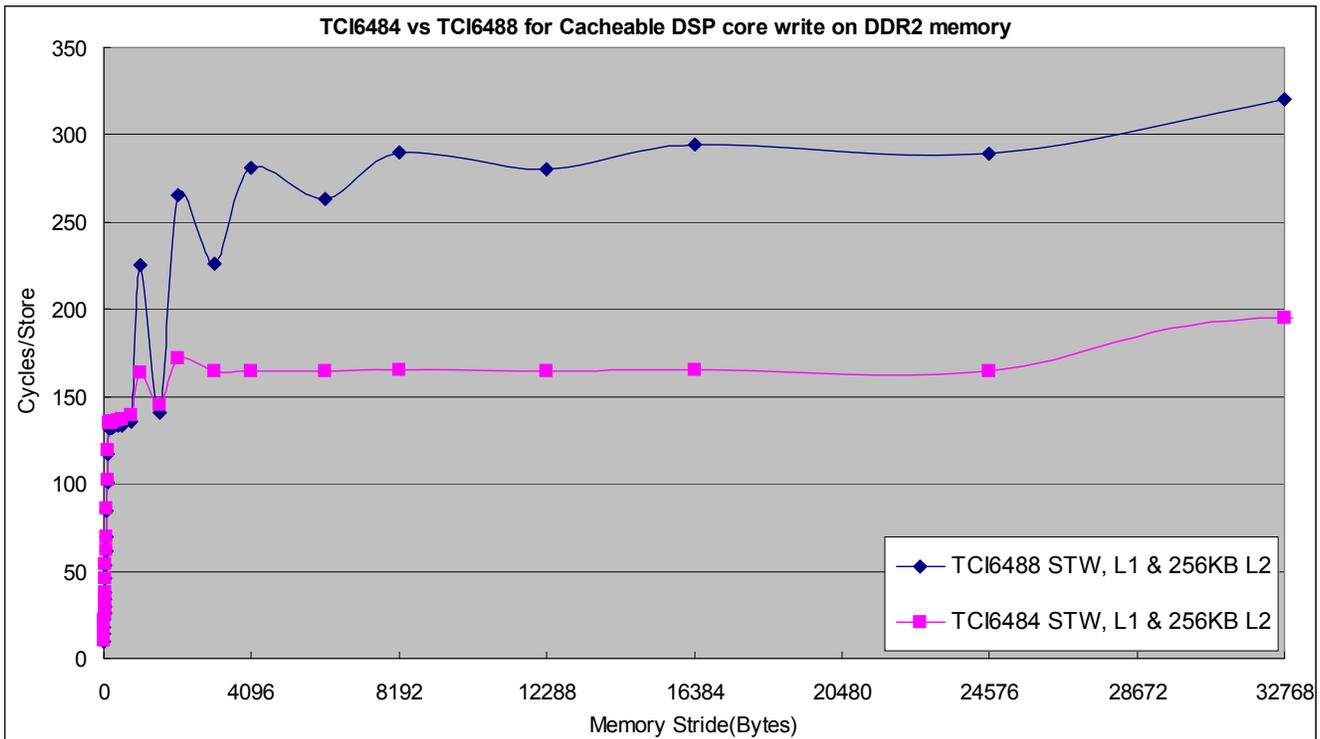
Since the L1D cache is not a write-allocate cache, and the cache is flushed before this test, it has no effect on the write operation. Therefore, the “32KB L1 only” case and the “Noncacheable” case are overlapped in the above figure.

Similar to the cacheable read operation, for the cacheable write operation with a memory stride of less than 128 bytes, the performance is dominated by cache as discussed above. For a memory stride larger than 4096 bytes, the performance becomes worse because of DDR2 SDRAM row switch.

L2 cache is a write-allocate cache. For any write operation, it always reads the 128 bytes including the accessed data into a cache line first, and then modifies the data in the L2 cache. This data will be written back to real external memory if a cache conflict happens or by manual writeback. When the memory stride is equal to or larger than 1024 bytes, the cycles needed for a write operation increases dramatically, because the conflict happens frequently for a big memory stride. Thus, every write operation may result in a cache line write back (for conflict) and a cache line read (for write-allocate).

The dirty cache line write back performance of a TCI6484 is improved compared to the TCI6482, TCI6486 and TCI6488. Figure 5 compares the cacheable write performance of the TCI6484 and the TCI6488.

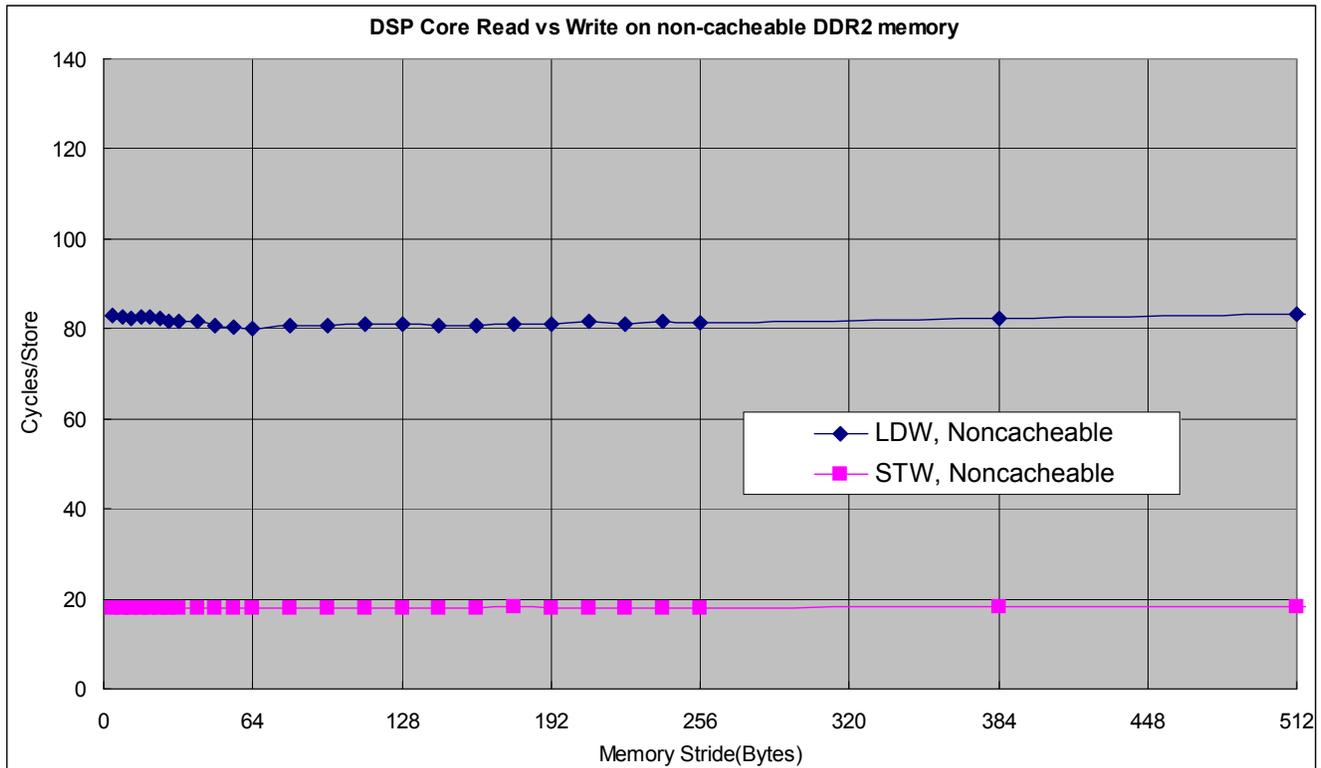
Figure 5 TCI6484 vs. TCI6488 for Cacheable DSP Core Write on DDR2 Memory



According to the above figure, for a big memory stride the TCI6484 cacheable write performance is much better than that of the TCI6488 because of the improvement on dirty cache line writeback.

For non-cacheable access, cycles for write is about 1/4 of the cycles for read as compared in Figure 6. This is because four write operations are combined into one transfer request in the L1D write buffer, while four read operations submit four transfer requests to the switch fabric center.

Figure 6 DSP Core Read vs. Writes to Non-Cacheable DDR2



## 4 Performance of DMA Access Memory

The EDMA3 architecture includes many features designed to facilitate simultaneous, multiple high-speed data transfers. Its performance is affected by the memory type and many other factors discussed in the following sections.

### 4.1 DMA Transfer Overhead

Initial latency is defined as the time between when a DMA event occurs and when real data transfer begins. Since initial latency is hard to measure, we measured transfer overhead instead. This value is defined as the sum of the latency and the time to transfer the smallest element. The values vary based on the type of source/destination peripheral and readiness of the source/destination ports and peripherals. [Table 4](#) and [Table 5](#) show the average cycles measured on 1GHz TCI6484 VDB for the smallest transfer (1 word) between different ports.

**Table 4 EDMA Transfer Overhead**

	Destination	L1D	L2	DDR2
Source				
L1D		1166	167	167
L2		167	166	182
DDR2		213	227	287

**Table 5 IDMA Transfer Overhead**

	Destination	L1D	L2
Source			
L1D		96	95
L2		95	95

In conclusion, transfer overhead is a big concern for short transfers and needs to be included when scheduling DMA traffic in a system. Single-element transfer performance will be latency-dominated. So, for small transfers you should make the trade-off between DMA and DSP core.

### 4.2 EDMA Bandwidth vs. Transfer Flexibility

EDMA3 channel parameters allow for many different transfer configurations. Most typical transfers burst properly, and memory bandwidth is fully utilized. However, in some less common configurations transfers are unable to burst, reducing performance. To properly design a system, it is important to know which configurations offer the best performance for high speed operations, and which must trade throughput for flexibility.

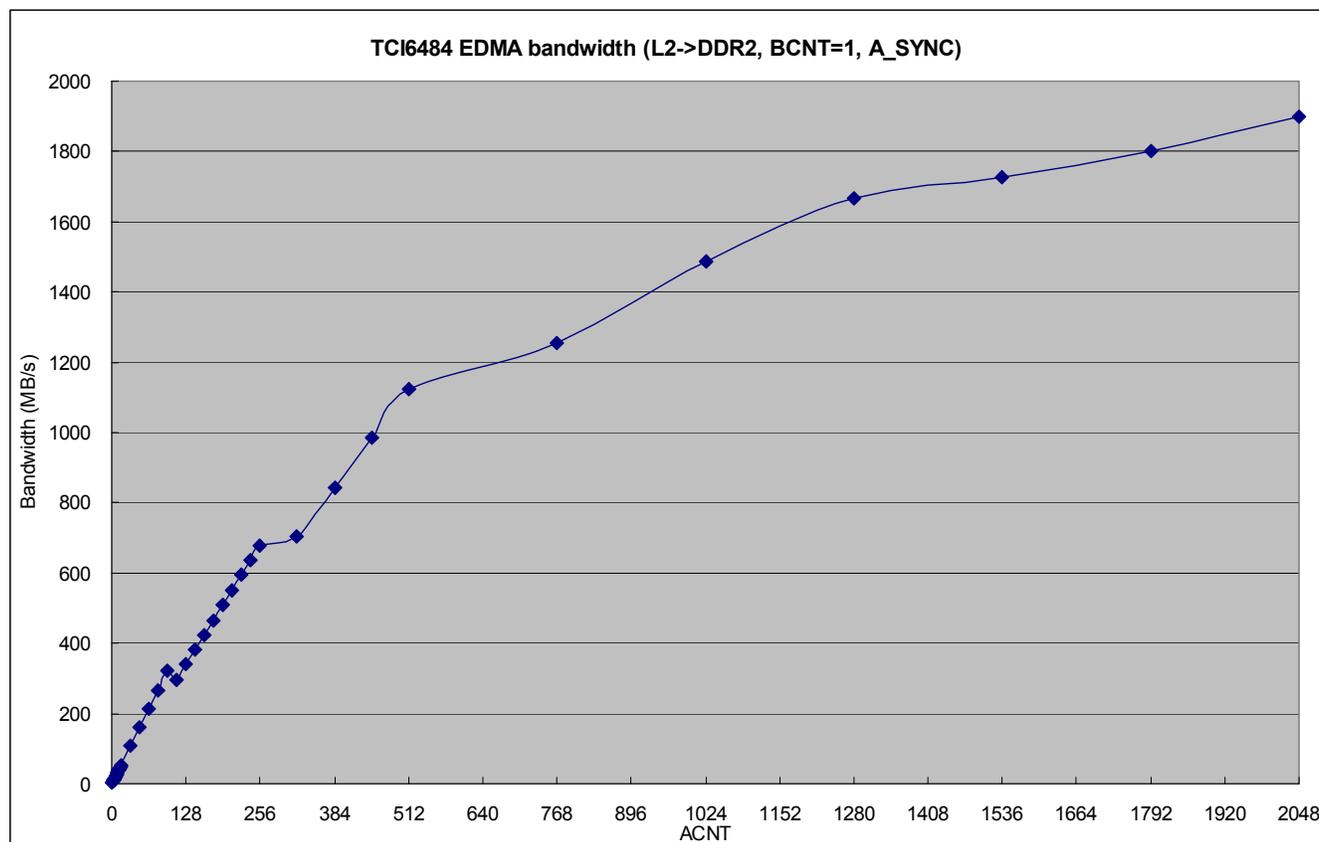
#### 4.2.1 First Dimension Size (ACNT) Considerations, Burst Width

To make full utilization of bandwidth in the transfer engine, it is important to fully utilize the bus width available and allow for data bursting.

Each TC has a 128 bit bus, 256bytes FIFO, and operates at 64 bytes of burst. ACNT size should be a multiple of 16 bytes to fully utilize the 128-bit bus width. ACNT should be a multiple of 64 bytes to fully utilize the 64-byte default burst width. ACNT should be a multiple of 256 bytes to fully utilize the 256-byte FIFO.

Figure 7 shows performance data from a TMS320TCI6484 running at 1GHz, transferring 1~2048 bytes from L2 to DDR2 using an EDMA3 channel.

**Figure 7** Effect of ACNT Size on EDMA Bandwidth



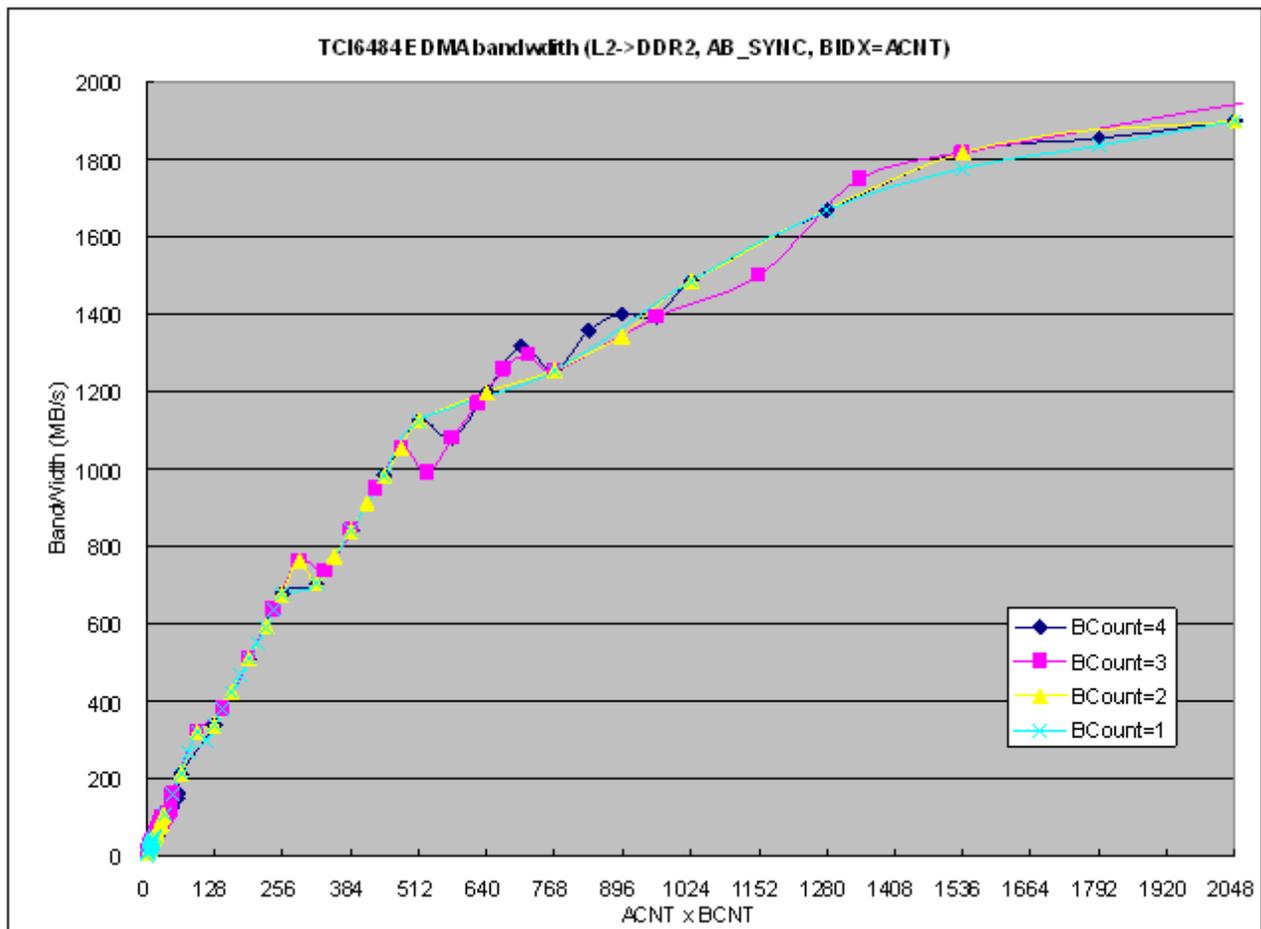
In conclusion, the bigger the ACNT, the more bandwidth can be achieved.

#### 4.2.2 Two Dimension Considerations, Transfer Optimization

If 2D transfer (AB\_Sync) is linear (BIDX=ACNT), the 2D transfer will be optimized as a 1D transfer.

Various ACNT and BCNT combinations were investigated; however, the overall transfer size (ACNT \* BCNT) was proven to have more bearing than the particular combination settings. Figure 8 shows the linear 2D transfer test results. The figure shows that no matter what the BCNT, the bandwidths are similar as long as ACNTxBCNT are the same.

Figure 8 Linear 2D Transfer



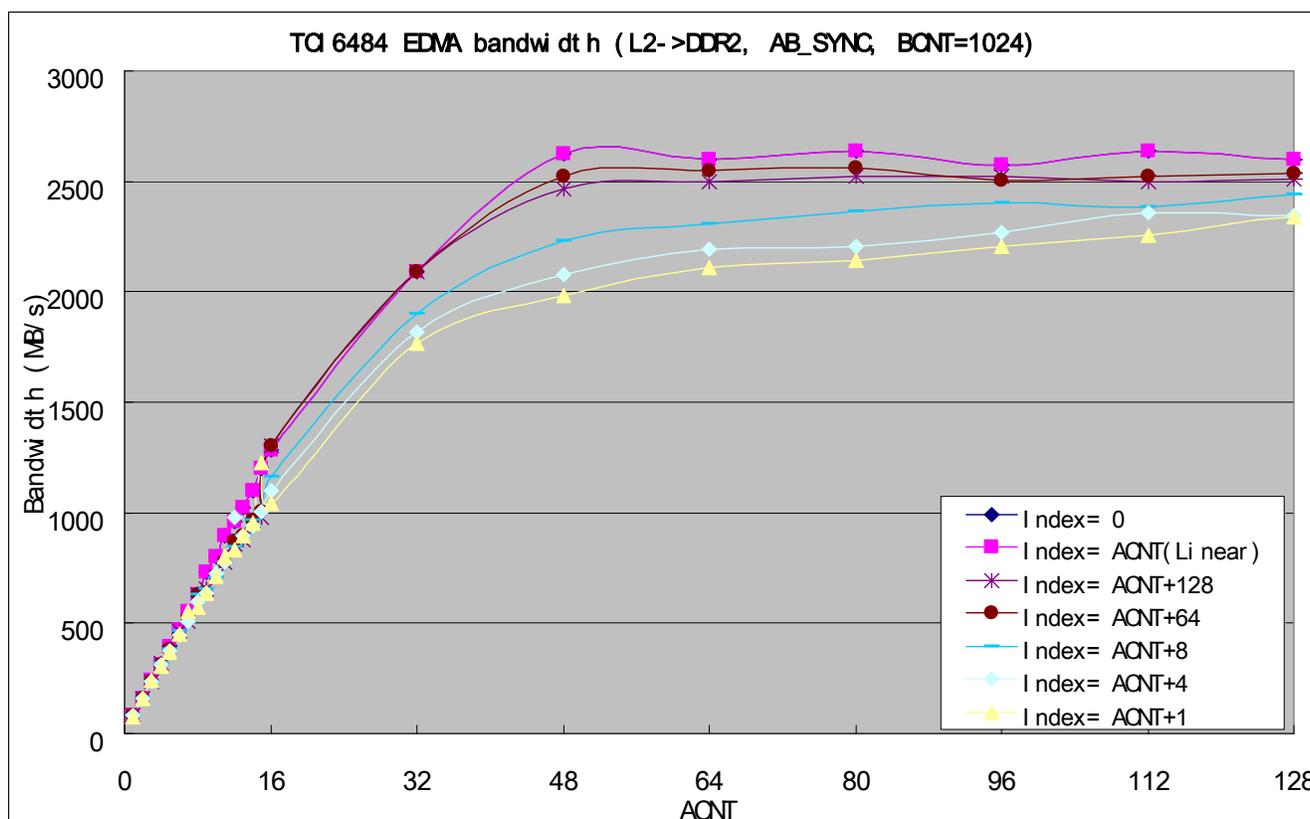
If 2D transfer is not linear, the bandwidth utilization is determined by the ACNT as shown in Figure 7.

#### 4.2.3 Index Consideration

Index dramatically affects the EDMA throughput. Linear transfer (Index= ACNT) fully utilizes bandwidth. Fixed Index (Index= 0) can utilize the same bandwidth as linear transfer. Other index modes will lower the EDMA performance. Odd index has the worst performance. If index is a power of 2, and it is larger than 8, the performance degradation is very small.

Figure 9 shows the index effect on EDMA throughput, transferring 1024 rows (BCNT= 1024) of 2D data from L2 to DDR2 with a different index.

Figure 9 Index Effect on EDMA Bandwidth



Without special note, all performance data in this application report are measured with Index= 0 or Index= ACNT.

#### 4.2.4 Address Alignment

Address alignment may slightly impact the performance. The default burst size of EDMA3 is 64bytes. If the transfer is across the 64-byte boundary, then the EDMA3 TC breaks the ACNT array into 64-bytes burst to the source/destination addresses. So, if the source or destination address is not aligned to the 64-byte boundary, and the transfer is across the 64-byte boundary, an extra burst will be generated to handle the unaligned head and tail data.

For big transfers, this overhead may be ignored. All data presented in this document are based on the address-aligned transfer.

## 5 References

1. TMS320TCI648x DSP Enhanced DMA (EDMA3) Controller User's Guide ([SPRUEE9](#))
2. TMS320C64x+ Megamodule Reference Guide ([SPRU871](#))
3. TMS320TCI6484 Datasheet ([SPRS438](#))
4. TMS320TCI6482 EDMA3 Performance ([SPRAAG8](#))



## IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

TI products are not authorized for use in safety-critical applications (such as life support) where a failure of the TI product would reasonably be expected to cause severe personal injury or death, unless officers of the parties have executed an agreement specifically governing such use. Buyers represent that they have all necessary expertise in the safety and regulatory ramifications of their applications, and acknowledge and agree that they are solely responsible for all legal, regulatory and safety-related requirements concerning their products and any use of TI products in such safety-critical applications, notwithstanding any applications-related information or support that may be provided by TI. Further, Buyers must fully indemnify TI and its representatives against any damages arising out of the use of TI products in such safety-critical applications.

TI products are neither designed nor intended for use in military/aerospace applications or environments unless the TI products are specifically designated by TI as military-grade or "enhanced plastic." Only products designated by TI as military-grade meet military specifications. Buyers acknowledge and agree that any such use of TI products which TI has not designated as military-grade is solely at the Buyer's risk, and that they are solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI products are neither designed nor intended for use in automotive applications or environments unless the specific TI products are designated by TI as compliant with ISO/TS 16949 requirements. Buyers acknowledge and agree that, if they use any non-designated products in automotive applications, TI will not be responsible for any failure to meet such requirements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

### Products

Amplifiers	<a href="http://amplifier.ti.com">amplifier.ti.com</a>
Data Converters	<a href="http://dataconverter.ti.com">dataconverter.ti.com</a>
DLP® Products	<a href="http://www.dlp.com">www.dlp.com</a>
DSP	<a href="http://dsp.ti.com">dsp.ti.com</a>
Clocks and Timers	<a href="http://www.ti.com/clocks">www.ti.com/clocks</a>
Interface	<a href="http://interface.ti.com">interface.ti.com</a>
Logic	<a href="http://logic.ti.com">logic.ti.com</a>
Power Mgmt	<a href="http://power.ti.com">power.ti.com</a>
Microcontrollers	<a href="http://microcontroller.ti.com">microcontroller.ti.com</a>
RFID	<a href="http://www.ti-rfid.com">www.ti-rfid.com</a>
RF/IF and ZigBee® Solutions	<a href="http://www.ti.com/lprf">www.ti.com/lprf</a>

### Applications

Audio	<a href="http://www.ti.com/audio">www.ti.com/audio</a>
Automotive	<a href="http://www.ti.com/automotive">www.ti.com/automotive</a>
Broadband	<a href="http://www.ti.com/broadband">www.ti.com/broadband</a>
Digital Control	<a href="http://www.ti.com/digitalcontrol">www.ti.com/digitalcontrol</a>
Medical	<a href="http://www.ti.com/medical">www.ti.com/medical</a>
Military	<a href="http://www.ti.com/military">www.ti.com/military</a>
Optical Networking	<a href="http://www.ti.com/opticalnetwork">www.ti.com/opticalnetwork</a>
Security	<a href="http://www.ti.com/security">www.ti.com/security</a>
Telephony	<a href="http://www.ti.com/telephony">www.ti.com/telephony</a>
Video & Imaging	<a href="http://www.ti.com/video">www.ti.com/video</a>
Wireless	<a href="http://www.ti.com/wireless">www.ti.com/wireless</a>

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265  
Copyright © 2009, Texas Instruments Incorporated