

Anchor Chips Incorporated

# Application Note

Programming AN3042 (CO-MEM Lite)  
Bus Master DMA Transfers

Rev 0.95  
12/16/98  
Mike Davis  
Applications Engineering

## Contents

INTRODUCTION .....	3
DMA FEATURE SUMMARY .....	4
QUICK START .....	5
PC ARCHITECTURE 101 .....	6
<i>Which DMA Controller?</i> .....	6
<i>Software</i> .....	6
<i>Hardware</i> .....	7
<i>Other Important Firmware and Software Components</i> .....	9
THE SAMPLE CODE .....	9
PERFORMANCE .....	10
LIMITATIONS OF THIS SAMPLE CODE .....	10
<i>Interrupts</i> .....	10
<i>Data To/From the Local-Side Memory</i> .....	11
<i>Sharing the AN3042 DMA Controller</i> .....	11
RECOMMENDED TOOLS .....	15
GLOSSARY .....	15
REFERENCES .....	17

The information in this document is subject to change without notice and should not be construed as a commitment by Anchor Chips Incorporated. While reasonable precautions have been taken, Anchor Chips Incorporated assumes no responsibility for any errors that may appear in this document.

No part of this document may be copied or reproduced in any form or by any means without the prior written consent of Anchor Chips Incorporated.

Anchor Chips 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 Anchor Chips product could create a situation where personal injury or death may occur. Should Buyer purchase or use Anchor Chips products for any such unintended or unauthorized application, Buyer shall indemnify and hold Anchor Chips and its officers, employees, subsidiaries, affiliates and distributors harmless against all claims, costs, damages, 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 Anchor Chips was negligent regarding the design or manufacture of the part.

Product and company names herein may be the trademarks of their respective owners.

Anchor Chips' customers are granted a license to use the information herein, and the associated code on Anchor's web site, **only** in connection with developing products that utilize Anchor chips. All other uses are prohibited.

The acceptance of this document will be construed as an acceptance of the foregoing conditions.

Copyright 1998, Anchor Chips Incorporated  
All rights reserved.

## Introduction

This Application Note explains Anchor-supplied sample code that demonstrates how to use the Bus Mastering DMA controller built into the AN3042 (CO-MEM Lite) chip. All components, hardware and software, focus on one objective: to move data (quickly and with minimal loading on other parts of the PC) from Host Memory to the AN3042's Shared Memory, and visa versa. This is shown by the arrow labeled **TRANSFERS** in Figure 1.

Anchor Chips provides this sample software to help you get your AN3042 design done and into production. Since we've already done much of the system programming for you, you can concentrate on your application, not on low-level system software.

Typical **sustained** performance for this code is 75 MB/sec for Reads from Host Memory, and 100 MB/sec for Writes to Host Memory.

If you are pressed for time, you can use the code as provided, treating it as a black box, without delving into how it actually works. You can focus mostly on the code that *calls* doDMA, and be done with it. See the DMA Feature Summary and Quick Start.

However, we think you will find it useful to understand how the AN3042 fits into the PC architecture to help diagnose problems when things go wrong.

The code discussed here is part of the Visual C++ example *Test3042*. You can download Test3042.zip from the Anchor Chips web site at <http://www.anchorchips.com/pcidev/cmемlite/download.htm> . We suggest that you download and unzip Test3042 first, then peruse the code as you read this App Note.

If you are using the AN3042 in an embedded system that doesn't use Windows, you may have to make changes at the device driver level, but the application-layer code may be useful unchanged.

We'll cover:

- Features of the DMA Controller.
- How to get started quickly, ignoring most of the details.
- PC Architecture and the Pentium<sup>1</sup> as it applies to Bus Master DMA transfers.
- The functions testDMA and doDMA, components of Test3042.
- The way our Device Drivers help set up DMAs.
- Performance issues.
- The limitations of this sample code.

---

<sup>1</sup> When this App Note says *Host CPU* or *Host Processor*, think *Pentium*. *Pentium* means the Pentium itself and all its successors (Pentium, Pentium II and Pentium Pro). Windows uses all Pentium variants in modes which are identical for our purposes (4 KB pages, 32-bit flat model). Although Win9x and WinNT can theoretically run on older processors (386 and 486), even 486 PCs rarely have PCI buses.

## DMA Feature Summary

Here's a summary of the features of the AN3042's DMA controller (see the p.2, pp. 39-40, and pp. 51-52 of the AN3042 Technical Reference Manual [Ref. 1]).

- The DMA controller's registers can be accessed (initialized) from either the PCI or Local bus. This is sometimes referred to as *PCI-initiated* DMA or *Local-initiated* DMA respectively. In *Test3042*, we only show PCI-initiated DMA.
- Ownership bits (L and P) in the DMACTL register allow the Host side or Local side to claim the DMA controller (to prevent the other side from corrupting the DMA registers by accessing them at the same time).
- The DMA controller always moves data by making the AN3042 the PCI Bus Master.
- The DMA controller always moves data a DWORD at a time, using all 32-bits (4 byte-lanes) for each data phase. Hence, the size of a transfer can be 4, 8, 12, ...16 K bytes, never 1, 2, 3, 5, ... bytes.
- The AN3042's Shared Memory is either the source or destination in every DMA transfer. The other device (the destination or source) can be Host Memory, or any other PCI device that is mapped into the PCI Memory Space<sup>2</sup>.
- The maximum transfer length is 16 KB.
- Most data phases occur with zero wait states (depending on the bandwidth supported by the PCs chipset and host memory). TRDY# and IRDY# remain asserted, and a new datum is moved on each 33 MHz PCI clock.
- The DMA controller can transfer data to/from any address in the full 32-bit PCI Memory Address Space.
- When the AN3042 is the PCI bus master, it uses the PCI cache line commands MRM (Memory Read Multiple), MRL (Memory Read Line) and MWI (Memory Write and Invalidate) where possible to optimize throughput and minimize the load on the rest of the PC.

---

<sup>2</sup> See *Memory Space* in the Glossary.

## Quick Start

Here's all you need to do to use the AN3042's DMA controller.

First, there are a few things you need to do before calling doDMA :

Install the driver (comemL.vxd or comemL.sys) into c:\windows\system and reboot your PC.

In your program:

1. Open the driver by calling comemOpenDriverL (*not* shown in Listings 1 or 2). See the program file TestUtil.cpp that you downloaded with *Test3042*.
2. Call comemCopyBarPtrL to get a Linear pointer to the AN3042 so you can access its Op Regs (Operation Registers) from your program (see Listing 1).
3. Call comemAllocContigMemL to allocate a buffer which is contiguous in both physical and linear address space (see Listing 1).
4. Call doDMA (see Listing 1).

That's all there is to it.

Then, here's what happens in doDMA (see Listing 2). For most applications, you can use doDMA unchanged.

For simplicity of explanation, let's assume the direction of transfer is *towards* the Host Memory (the AN3042 controller will do one or more MWI [Memory **W**rite Invalidate] transactions):

**source** = a block of *n* DWORDs in AN3042 Shared Memory (previously loaded with data).

**destination**= a buffer in Host Memory (previously allocated by a call to comemAllocContigMemL).

1. Set the DMALBASE register to point to the **source**, the start of the data in Shared Memory (see line 19 of doDMA in Listing 2). A value of 0 points to the zeroth DWORD in Shared Memory.
2. Set the DMASIZE register to the number of DWORDs to be transferred (line 24).
3. Set the DMAHBASE register to the physical address of the **destination** (line 29).
4. Clear the DMA completion bit (HINT register, bit 5) (line 33).
5. Kick off the DMA by writing to just the low order byte of the DMACTL register (line 37). Here, we both set the direction and start the DMA.
6. Wait for the DMA complete bit to be set (HINT register, bit 5) (lines 50-67).

## PC Architecture 101

### Which DMA Controller?

One type of DMA in PCs is **Bus Master DMA**. In this App Note (and all other Anchor Chips' PCI literature) the term *DMA* always refers to Bus Master DMA — DMA driven by the DMA controller in the AN3042, *Bus Master* because the AN3042 becomes the PCI bus master.

From the perspective of a PCI device, for moves to (Writes) and moves from (Reads) Host Memory, the *Target* device is always the *Host to PCI Bridge*. This is called a *host* transfer.

However, the Target device *can* be another device on the PCI bus (including AN3042s). This is called a *peer-to-peer* transfer.

The other type of DMA in PCs called **System DMA** uses the DMA controller in the chipset. The Intel 8237 chip was the System DMA controller in the original IBM PCs. It was later integrated into PC chipsets, for example, the PIIX3 (PCI ISA IDE Xcelerator) part of Intel's chipsets. Every PC has at least two System DMA controllers. (see Hardware below regarding chipsets).

The System DMA controller is far more limited (it can only address the lowest 1 MB of physical memory) and slower than the AN3042 DMA controller, having been designed originally for the ISA bus (a bus that might be able to move 6 MB/sec. on a good day). Moreover, the System DMA controller can only move data between the ISA bus and Host Memory, never between PCI devices.

### Software

The code that we provide in Test3042 and the Device Driver has the following parts (see Figure 1):

① High-level Application Code

GUI (Graphical User Interface) routines (not discussed further in this App Note).

② Low-level Application Code

The primary code discussed in this App Note (see Listing 1 on p. 12 and Listing 2 on p. 14).

③ CO-MEM Lite Interface Library (comemLif.lib)

Major functions:

1. Allows the application to create objects of class CO-MEM Lite. This code is architected to easily support multiple AN3042 chips in one system.
2. Provides routines that interface to the Device Driver ④.

For the sample program Test3042, ①, ② and ③ are *statically* linked by the Visual C++ linker, producing Test3042.exe. This is indicated by the ①, ② and ③ boxes actually touching.

The short line connecting ③ and ④ indicates that ④ is *dynamically* (run time) linked to the Test3042 process when Test3042 opens the device driver. This dynamic linking occurs in the call to comemOpenDriverL in TestUtil.cpp (not shown in Listings 1 or 2).

④ Device Driver

A Ring 0 Device Driver is an essential part of our AN3042 support because only Ring 0 code is privileged to access the Paging Unit. Application code (layers ① and ②) running at Ring 3 cannot access the Paging Unit.

By accessing the Paging Unit, the Device Driver provides two vital services to the Ring 3 application:

1. The Device Driver converts the value in the AN3042's BAR0 (Base Address Register 0) from a Physical to a Linear address.  
BAR0 is initialized by the BIOS at boot time. In Physical Address Space, BAR0 points to the base of all AN3042 functions that are visible from the PCI side (operation registers and Shared Memory). The Device Driver returns the Linear address of the contents of BAR0 (Base Address Register 0) to the application so the application can directly access the AN3042.
2. When called upon (through the IOCTL\_COMEM\_ALLOCCONTIGMEM entry point), the Device Driver allocates 16 KB blocks (four 4 KB Pages) that become the source or destination for AN3042 DMA *TRANSFERS*. These 16 KB blocks *must* be contiguous in both Physical and Linear memory space. Contiguity can only be guaranteed by Ring 0 code, because only it has access to the Page Tables. 16 KB is the size of AN3042 Shared Memory, hence 16 KB is the maximum size of a DMA transfer.

The actual Device Driver file in your system (normally installed in c:\windows\system) will be one of two types:

Win 9x	comemL.vxd
Win NT	comemL.sys

As shown by the arrows which go to the host side of ⑥, it is common for each software layer (①, ②, ③ and ④) to directly access the AN3042 (through the Host to PCI bridge ⑦).

## Hardware

Test3042 and other host processor software accesses the hardware components by generating memory accesses that are converted by the Paging Unit<sup>3</sup> from Linear to Physical addresses.

⑤ Host memory – Is in the same Physical Address Space as PCI devices (but of course at a different location within that space).

⑥ Chipset – These are most commonly Intel Chipsets, such as the Intel® 440BX AGPset and Intel® 440ZX AGPset. See <http://www.intel.com/design/pcisets>.

⑦ Host to PCI Bridge – a part of the Chipset that interfaces the Pentium, Host Memory, etc., to the PCI bus.

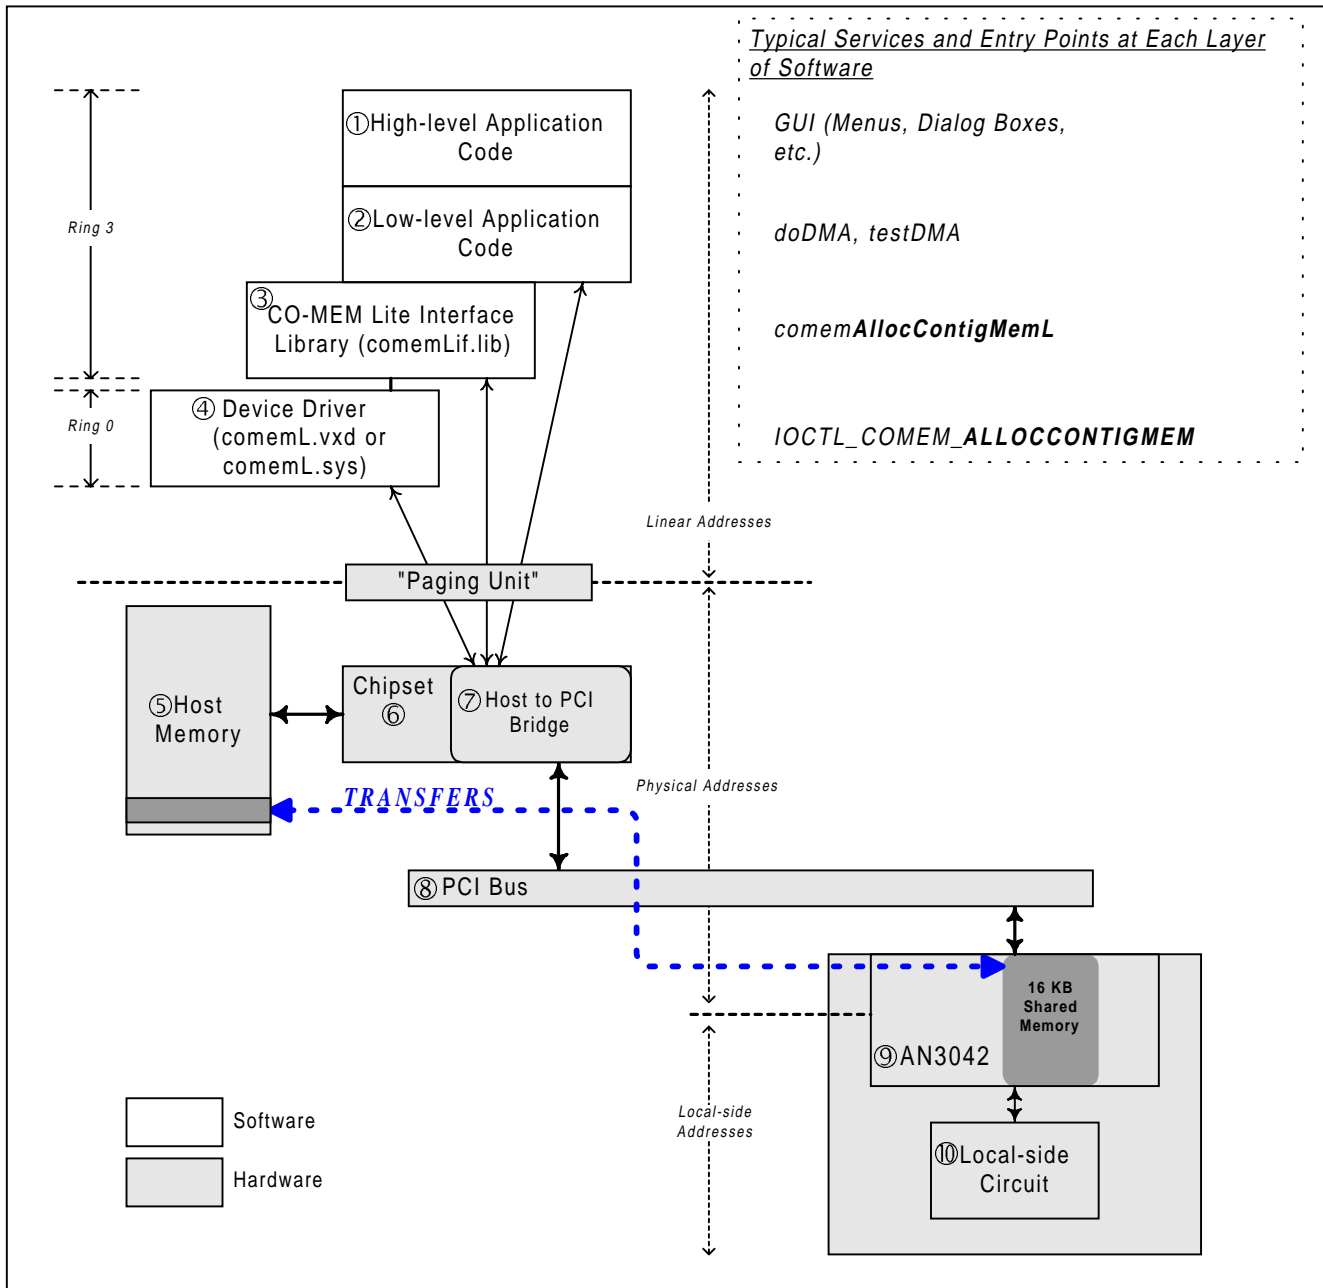
⑧ PCI Bus – 32 bit, 33 MHz.

⑨ An AN3042 and Local-side Circuit on an add-in board plugged into the PCI Bus.

⑩ The Local Circuit. Can consist of:

- a complex state machine, or nearly any microprocessor
- a memory subsystem, which can consist of a mix of memories
- any application circuitry.

<sup>3</sup> Intel never uses the terms Paging Unit per se in their documentation. They talk about all the low-level components of the paging mechanism, such as Page Tables, Page Directory, etc., but never specifically Paging Unit. We suspect the reason for this is that the paging unit is quite distributed into components on the Pentium chip, and in Page Tables in Host Memory. See the Glossary for the Anchor Chips' definition of Paging Unit.



**Figure 1: PC Components Involved in Bus Master DMA Transfers**

Notes about Figure 1:

- In general, each software layer calls services in lower layers. However, after a layer has obtained a Linear Address (a 'C' pointer) to the AN3042, it is free to read or write directly to the AN3042 without 'calling through' lower layers. See the three bi-directional arrows that point to the *Host to PCI Bridge*, which originate at ② and ③, as well as ④. Also see code in Listing 2. Statements such as line 24  

```
g_regp->dmasize = length - 4;
```

access the AN3042 directly.
- All Software lives in the domain of Linear Addresses, whereas all hardware lives in the domain of Physical Addresses.



- Application programs in Win9x and WinNT use a flat model (32-bit protected mode). There are exceptions, which are buried in Windows internals and normally don't concern us. For example, Win9x still uses *16-bit* protected mode for much of the code in User.exe (one of several key Windows executable files). Test3042 and all new applications that you write should be compiled into flat-model 32-bit code.

## Other Important Firmware and Software Components

- The Plug and Play (PnP) BIOS (a part of the System BIOS).

At boot time, during the first second after Reset is released, the PnP BIOS scans the PCI bus and assigns a unique Physical Address range ('Region') to each PCI BAR (Base Address Register). This process is called *enumeration*.

Since the PnP BIOS is a fixed component on the motherboard, it 'knows' which resources are required by devices on the motherboard (such as Host Memory), and can reserve those, assigning the remaining Physical Addresses to add-in devices.

See the PCI 2.1 Specification (Ref. 3) for details about the device scan, BAR operation, etc.

When Windows boots, it may decide to redo what the PnP BIOS did. For example, if Windows doesn't think that the device has a valid device driver (Windows can't find a match in the Device Driver Database), it will do the 'safe' thing, disabling the device by writing zero to all its BARs.

- Windows – Called on by the application and the Device Driver for *many* services.
- comemL.inf – Is the major file to tell the Windows installation code how to match the AN3042 to the correct Device Driver.
- The Registry.
- The Device Driver Database, which consists all .bin files in the directory c:\windows\inf. These files are a compiled database of all the .inf files in the c:\windows\inf tree. You can delete them to force the Plug and Play system to rebuild this database during the next reboot.

## The Sample Code

See Listings 1 and 2. Important code is in **bold type**. Since the code is heavily commented, please see the code for details.

The code consists of two C++ functions: doDMA and testDMA. (If you don't know C++, don't panic. We've used very few of the C++ extensions to C.)

testDMA is intended to be called by higher level code after the device driver has been opened, then testDMA calls doDMA.

## Performance

Typical **sustained** performance for this code is 75 MB/sec. for Reads from Host Memory, and 100 MB/sec. for Writes to Host Memory.

By taking certain shortcuts (not shown in Listings 1 and 2), we have been able to push Write performance to over 120 MB/sec., remarkably close to the 132 MB/sec. theoretical limit of a 32-bit 33 MHz PCI bus. The overhead to set up the next transaction, and pauses during a transfer, is extremely low for this code and the AN3042.

Large transfers (16 KB) are fastest because the key factor in overall speed is the time to set up the next transfer.

Most of the time to set up the next transfer is consumed by transactions on the PCI bus itself, not execution of the code in the Pentium core. When you reach the optimization phase of your project, try to minimize the actual DMA setup and polling transactions on the bus. They are *much* slower than execution of instructions in the Pentium core.

For example, a write to an AN3042 Op Reg takes about **6** PCI clocks = 180 ns. In the same time (assume a 300 MHz Pentium = 3ns instruction execution time, and DMA code running from the Pentium's instruction cache), the processor core will execute about **60** instructions. If we can eliminate one PCI transaction, we can save the time it takes to run about 60 instructions!

## Limitations of This Sample Code

### Interrupts

doDMA (Listing 2) polls the AN3042 for DMA completion. At the expense of considerable complexity, an interrupt-driven DMA device driver could be written.

There is a possibility (though not a certainty) that *overall system performance* would be improved using interrupts, because interrupts would allow processor cycles that are consumed in the spin loop (waiting for DMA completion) to be used to by other threads to improve overall system responsiveness.

Throughput across the PCI bus, however, would probably *not* be improved, because PCI throughput is mostly limited by:

- DMA setup time
- PCI bandwidth consumed by other PCI devices
- The ability of the Target device (the Host to PCI Bridge and the Host Memory system) to source or sink the data. Newer chipsets can be more efficient than older ones, because they use a variety of techniques to sustain zero wait-state PCI bursts.
- Loading on Host Memory by other agents, for example, the AGP bus, and instruction fetches that that must actually fetch from Host Memory (they missed in both Level 1 and Level 2 caches).

Polling is reasonably efficient because DMAs are done quickly — 16 KB transfers complete in about 150  $\mu$ sec. Interrupts may not improve performance much because of the overhead involved in interrupt processing:

- saving and restoring processor state
- a small amount of cache thrashing when the context is switched.

## Data To/From the Local-Side Memory

In some applications, it may be OK to leave the data in Shared Memory and access it from the Local-side Circuit without moving to other memory on the Local Side. In most applications, however, an extra step would be required: moving the data from the Shared Memory to a Local-side memory so the Shared Memory could be used for the next load of data.

The code in doDMA does nothing to prod the Local-side Circuit to move the data from 3042 Shared Memory to its own local memory.

Here's a possible double-buffered scheme to move the data from Shared Memory to Local-side memory, designed to maximize sustained throughput. (For simplicity of explanation, assume data is moving from Host Memory to the Local-side memory.)

The 16 KB Shared Memory is split into two halves. A single-DWORD guard band at address 0 and 0x2000 is required because of Errata #4 (Local bus Read access to the AN3042 Shared Memory will block PCI Write access to Shared Memory in some cases).

1. AN3042 DMA controller moves data from Host Memory to 1<sup>st</sup> half of Shared Memory.

Simultaneously, the Local-side Circuit moves data from 2<sup>nd</sup> half of Shared Memory to the Local-side Memory (data put there from the Host Memory the last time the PCI side 'owned' the 2<sup>nd</sup> half).

2. When both halves are transferred, the processors synchronize through the mailbox registers (HLDATA or LHDATA).
3. The programs on each side change buffer pointers.  
Host points to 2<sup>nd</sup> half, Local points to 1<sup>st</sup> half.
4. Both processors would then start their transfers to/from their new half.  
Host to 2<sup>nd</sup> half, Local from 1<sup>st</sup> half.

AN3042 DMA can be set up and initiated by the Pentium ('Host-initiated DMA'), or by the Local circuit ('Local-initiated DMA').

## Sharing the AN3042 DMA Controller

For the code in Listing 1 and 2, we've assumed that only one thread will access the DMA controller. If more than one thread on the host side wants to access the DMA Controller, or if the Local-side processor might access the DMA controller, code to perform software arbitration must be added.

Ref. 5 recommends using the VDMAD (Virtual DMA Driver, a VxD) to virtualize<sup>4</sup> any DMA controller.

Additionally, you may want to use the L and P bits in the DMACTL register to arbitrate between Host-side and Local-side ownership of the AN3042's DMA controller.

---

<sup>4</sup> To *virtualize* a device is to provide an intermediary VxD between the application layers and the device. (The computer science community loves the word *virtual*.) A VxD, a Virtual Anything (x) Driver, is then called by any thread that wants to access the device. The VxD maintains a separate state table for each thread, arbitrates for the device, and manages any conflicts that arise.

## Listing 1: testDMA

```

1  DWORD g_linBAR[COMEM_MAX_BARS]    // Global array which will contain the Linear BAR0 and BAR1 addresses.
2
3  // This structure MUST be designated 'volatile' for the VC++ 'Release' version (as opposed to the 'Debug' version).
4  // If this structure is not designated 'volatile', optimizers destroy the intent of inner-loop code in doDMA below.
5  volatile struct op_regs_struct_42_t * g_regp;
6
7  //
8  // testDMA - An example function that tests the 3042's DMA controller that:
9  //   - allocates and deallocates the Host Mem physical pages
10 //   - writes patterns to the source and destination memory
11 //   - calls doDMA to initiate the transfer
12 //   - checks the result of the 3042 DMA transfer.
13 //
14 DWORD testDMA(DWORD direction, DWORD comemID)
15 {
16     DWORD lin, phys, i, returnCode;
17     int errorCnt = 0;
18
19     printf ("\n");
20
21     if (direction == directionToHost)
22         printf ("DMA data from 3042 Shared Mem to Host Mem...\n");
23     else
24         printf ("DMA data from Host Mem to 3042 Shared Mem...\n");
25
26     // Get the Linear BAR0 and BAR1 pointers (we only use BAR0 here) so we can access the 3042's registers.
27     // You may want to move this out of testDMA, doing it only once at the start of your application.
28     //
29     // Throughout this code, 'L' at the end of function names means Lite, as in CO-MEM Lite.
30     returnCode = comemCopyBarPtrL(g_linBAR, comemID);
31     if (returnCode != NO_ERROR)
32     {
33         reportErrorCode(returnCode, "creating BAR pointers.");
34         errorCnt++;
35     }
36
37     // Setup the structure template we use to access the 3042 op regs.
38     g_regp = (struct op_regs_struct_42_t *) (g_linBAR[0] + OP_REGS_BASE_42);
39
40
41     // Allocate 4 Pages (4 KB each, 16 KB total = size of 3042 Shared Mem) which are (must be!)
42     // contiguous in Physical (and Linear) memory space.
43     //
44     // We pass the Physical address to the 3042's DMA controller DMAHBASE register.
45     // We need the Linear address to reference the Pages from this program, which runs on the 'Linear side' of
46     // the Pentium's Paging Unit.
47     //
48     // This allocation may be done once when the application is launched. However, here we chose to
49     // allocate and deallocate for each invocation of testDMA. Remember to pair allocations with
50     // deallocations.
51     returnCode = comemAllocContigMemL (4,          // In:  Number of 4 KB Pages. comemAllocContigMemL can allocate 1 to 4 Pages,
52                                         //          but must allocate 4 for this test.
53                                         &lin,      // Out: Returns the Linear address of the allocated Pages.
54                                         &phys,     // Out: Returns the Physical address of the allocated Pages.
55                                         comemID);  // In:  The 3042 that 'owns' the pages.
56
57     if (returnCode != NO_ERROR)
58     {
59         printf("Error in comemAllocContigMemL.");
60         return (++errorCnt);
61     }
62
63     DWORD * hostLin  = (DWORD *) lin;
64     DWORD * hostPhys = (DWORD *) phys;
65
66     // For this test move the max block size. Start at the beginning of Shared Mem and use the maximum DMAsize.
67     DWORD SMstart = 0;
68     DWORD DMAsize = 0x4000;    // 16 KB
69
70     if (direction == directionToHost)
71     {
72         // Initialize source buffer (3042 Shared Mem).
73         //
74         // initSharedMem writes to the 3042's Shared Mem as a PCI target, which is about 10 times
75         // slower than the DMA transfer.
76         initSharedMem(SMstart, DMAsize, ADDR_PATTERN, comemID);
77
78         // Write pattern to destination, so this test can't pass unless 3042 DMA controller actually moves the data.
79         // printf ("Initializing destination at Linear 0x%08x  pattern=0xFEEDBEEF\n", hostLin);
80         for (i = 0 ; i < DMAsize/4 ; i++)
81             hostLin[i] = 0xFEEDBEEF;
82     }
83     else

```

```

83     {
84         // Initialize source buffer in Host Mem
85         // printf ("Initializing source at Linear 0x%08x  DMAsize=0x%x  pattern=ADDR_PATTERN\n", hostLin, DMAsize);
86         for (DWORD i = 0 ; i < DMAsize/4 ; i++)
87             hostLin[i] = i;
88
89         // Write pattern to destination, so this test can't pass unless 3042 DMA controller actually moves the data.
90         initSharedMem (SMstart, DMAsize, FILL_PATTERN, comemID);
91     }
92
93     // Set 3042's Latency Timer to the max (at the expense of potentially causing latency issues
94     // for other PCI devices).
95     // You'll definitely want to move this to an initialization routine to do it only once.
96     PCI_CONFIG_HEADER_0 temp;
97     returnCode = comemGetPCIInfoL(&temp, comemID);
98     temp.LatencyTimer = 0xff;
99     returnCode = comemSetPCIInfoL(&temp, comemID);
100
101     // Set up to calculate the time and MB/sec ('bandwidth') for the DMA transfer.  Use the built-in
102     // Pentium high-performance timer.
103     LARGE_INTEGER llmHPTimerFreq;    // High Performance Timer: Frequency
104     LARGE_INTEGER llmHPT1;           // High Performance Timer: Time 1
105     LARGE_INTEGER llmHPT2;           // High Performance Timer: Time 2
106     LARGE_INTEGER llmT_uSec;         // Time in microseconds
107     QueryPerformanceFrequency(&llmHPTimerFreq);
108     QueryPerformanceCounter(&llmHPT1);
109
110     // Note: 'Host Mem' can be PCI Memory (physical) addresses other than those assigned to Host Memory.
111     // In this example, we use Host Mem as the source and destination.  But the 3042 does
112     // support peer-to-peer DMA transfers.  Get the physical addresses of the peer PCI devices from their BARs.
113     errorCnt += doDMA ( hostPhys,      // host-side physical address
114                       SMstart,        // Shared Mem start (first byte of shared mem = address 0)
115                       DMAsize,        // size of DMA transfer (in bytes)
116                       direction,      // the direction of the transfer
117                       comemID);       // which 3042 are we driving
118
119     if (QueryPerformanceCounter(&llmHPT2))
120     {
121         llmT_uSec.QuadPart = ( (llmHPT2.QuadPart - llmHPT1.QuadPart) *          // Delta time in ticks.
122                               ((LONGLONG)1E9/llmHPTimerFreq.QuadPart)          // Adjust for the frequency.
123                             ) / (LONGLONG)1E3;
124
125         double bandwidth = (double)1E6 * (double)DMAsize / (double)llmT_uSec.LowPart;
126
127         printf("DMA completed in  %3d sec  %03d msec  %03d usec: DMA Rate = %11.2le bytes/sec\n",
128               llmT_uSec.LowPart/(LONG)1E6,
129               (llmT_uSec.LowPart/(LONG)1E3) % 1000,
130               llmT_uSec.LowPart          % 1000,
131               bandwidth
132         );
133     }
134
135     if (errorCnt > 0)
136     {
137         printf ("Error: doDMA call failed hostPhys=%08x SMstart=%08x DMAsize=%08x direction=%d comemID=%d\n",
138               hostPhys, SMstart, DMAsize, direction, comemID);
139         // If an error, is probably fatal, but we'll attempt to forge ahead anyway.
140     }
141
142     // Check the destination.
143     if (direction == directionToHost)
144     {
145         for (i = 0 ; i < DMAsize/4 ; i++)
146         {
147             if (hostLin[i] != i)
148             {
149                 if (errorCnt++ < 8) // Report first 8 errors only.
150                     printf ("Error: PCIAddr=%08x LocalAddr=%08x Wrote=%08x Read=%08x\n",
151                           hostLin+i, i, i, hostLin[i]);
152             }
153         }
154         if (errorCnt > 0)
155             printf ("%d errors total\n", errorCnt);
156     }
157     else
158     {
159         errorCnt += checkSharedMem (SMstart, DMAsize, ADDR_PATTERN, comemID);
160     }
161
162     // Deallocate the 4 Pages we allocated above.
163     comemDeAllocContigMemL(&lin, &phys, comemID);
164
165     return(errorCnt);
166 } // end testDMA//

```

## Listing 2: doDMA

```

1 // doDMA
2 //
3 // Transfer up to 16KB via the 3042's DMA controller.
4 //
5 // Although in this example, we transfer between Host Mem and the 3042's Shared Mem, we could
6 // use doDMA *unchanged* to do peer-to-peer transfers between devices on the PCI bus.
7 //
8 DWORD doDMA (DWORD * PCIphysMem, // The PCI (physical) Memory space address (as opposed to I/O or Config space addresses).
9             // The 3042 needs this because it knows nothing (directly) about Linear addresses.
10            DWORD SMstart,        // Shared Mem start. Is 0 to point to the zeroth DWORD of 3042 Shared Mem.
11            DWORD length,         // Length in bytes. Gets converted to DWORDs below.
12            DWORD direction,      // To or From Host Mem.
13            DWORD comemID)        // Which 3042 we're driving.
14 {
15     DWORD errorCnt = 0;
16
17     // Set the starting Shared Mem address in DMALBASE register.
18     g_regp->dmalbase = SMstart; // Give it a byte address, but since DMALBASE bits 0 and 1 are dead,
19                                // this loads a DWORD address. The L in DMALBASE means Local, as in the
20                                // Local-side interface.
21
22     // Set the DMASIZE.
23     g_regp->dmasize = length - 4; // Give it a byte size, but since DMASIZE bits 0 and 1 are dead,
24                                // actually loads a DWORD size. Also, the value written to the DMASIZE register
25                                // must be reduced by one DWORD (4 bytes).
26
27     // Set the DMAHBASE to the PCI (physical) Memory address
28     g_regp->dmahbase = (DWORD) PCIphysMem; // Once again, we give it a byte address, but because bits 0 and 1 are dead
29                                // loads a DWORD address.
30
31     // Clear the DMA complete bit.
32     g_regp->hint &= DMA_COMPLETE_BIT;
33
34     // Write low-order byte of DMACTL to set the direction, and to KICK OFF the DMA.
35     // The only operations that MUST be done on each DMA transfer are the kick off and the wait for DMA_COMPLETE_BIT.
36     g_regp->dmactl_v.dmactl_bytes[0] = (unsigned char) direction;
37
38     // Now the 3042 DMA masters the PCI bus and moves the data.
39     // Looking at the PCI bus, you'd see mostly bursts of data moved by the DMA, with occasional polls of HINT.
40     // The length of the bursts is determined by the value in the 3042's Latency Timer,
41     // and other PCI bus activity (including our HINT polling).
42
43     // Wait for the DMA complete bit to be set.
44     //
45     // We also want to check for a timeout here, because we don't want our thread to hang with no notification
46     // if something goes wrong.
47     const DWORD loopsPerByte = 0x100; // Number of times through this 'while' loop per byte of DMA.
48     int timeout = length * loopsPerByte; // Give lots of time to complete.
49     while ( ((g_regp->hint & DMA_COMPLETE_BIT) != DMA_COMPLETE_BIT)
50            && timeout
51            )
52     {
53         // Reduce the frequency of HINT polling to improve performance, but poll often enough
54         // that a poll occurs soon after the end of the DMA transfer (ideally just after the end).
55         //
56         // The optimal delay to be inserted here is dependent on:
57         // - your processor speed
58         // - the length of DMA transfer
59         // - characteristics of the PCI bus arbiter, the Latency Timer setting, and other PCI bus traffic
60         // - etc., etc.
61         // A loop count of 150 works well for the 16 KB blocks, 200 MHz Pentium Pro and 440FX Chipset used to tune this example.
62         for (volatile DWORD i = 0 ; i < 150 ; i++)
63             ;
64
65         timeout--;
66     }
67
68     if (timeout <= 0)
69     {
70         errorCnt++;
71         printf ("Error: timed out before DMA_COMPLETE_BIT set. errorCnt=%d\n", errorCnt);
72     }
73
74     return (errorCnt);
75 } // end doDMA
76
77

```

## Recommended Tools

- Microsoft Visual C++ 5.x. We are not presently supporting VC++ 6.x. We do not intend to support older versions of VC++.
- SoftICE by Compuware® NuMega at [www.numega.com](http://www.numega.com) .  
We've found this tool useful, especially for the following commands:
  - pci – dumps Config Regs for all PCI devices. Helps you find out if the system even sees your card (do the BARs have reasonable values?, is Bus Mastering enabled?). Many other uses.
  - peekd, poked – reads and writes to PCI devices (Physical addresses). This is a good way to bypass software that is under development and talk directly to the AN3042.
  - phys – displays all Linear Addresses for a given Physical Address. Is a good sanity check on the Linear Addresses returned by routines such as comemCopyBarPtrL.
- A logic analyzer to look at PCI bus transactions, preferably with PCI transaction formatting software. We currently use an HP 1660C with FuturePlus FSPCI64E support hardware and software. Although there is some elegance in dedicated PCI analyzers (VMetro, HP and others), we have found the HP 1660C and FuturePlus cost effective and functionally adequate.
- A reasonable library. See References.

## Glossary

Linear address – addresses on the processor (hence program) side of the Paging Unit. See References 2, 5 and 6. Also known as User Space or Virtual Address. For the converse, see Physical address.

Memory Space – A transaction on the PCI bus can access one of three independent address spaces:

- Configuration (Config) Space – accesses the Config Regs in any device
- Memory Space – used for most transactions after the system is configured. Supports bursting.
- I/O Space – provided only for backward compatibility with x86 processor I/O space.

Because most PCI devices do not support bursting for transactions in I/O Space, it should be used only where absolutely necessary – to support legacy devices. Legacy devices (such as VGA cards) perform many operations using registers mapped into I/O space.

Note that it is possible for a Memory Space transaction to not access Host Memory. A Memory Space transaction can access Host Memory, or can access any other memory or register on the PCI bus that is mapped via the BARs into Memory Space.

A Master tells the Target which Space it wishes to access via the Command emitted by the bus master during the Address Phase (see p. 21 of Reference 3), and by the LS bit (the 'space indicator' bit) of each BAR (0 = Memory, 1 = I/O).

See Physical Address.

Paging Unit – The collection of Pentium components that perform Linear to Physical and Physical to Linear address translation. The 'Paging Unit' consists of:

- the Page Directory
- Page Tables
- assorted bits in Pentium Control Registers (CRs) that determine various paging modes
- Windows kernel code that manages the assorted tables and bits.

The CRs reside in the core of the Pentium. The Tables reside in Host Memory, but the most recently used entries are cached in the core of the Pentium (in the TLBs – Translation Lookaside Buffers).

The Paging Unit also handles:

- Demand Paging, which allows code and data to be swapped into and out of Host Memory from a disk swap file as needed (on demand). See Chapter 13 of Reference 6.
- Protection, which allows a Process to own a set of pages, prohibiting other code from accessing those pages.

Pages are often locked down by the Paging Unit, so they do not move in Physical Space or get swapped to disk. For example, pages associated with PCI devices and the 16 KB buffers allocated by comemL.vxd or comemL.sys get locked down.

**Physical address** – addresses as seen by hardware, on the hardware side of the Paging Unit. For our purposes this hardware is all PCI devices and host memory. On the PCI bus, a Physical Address is identical to a PCI Memory Space Address (as opposed to an I/O Space Address or a Configuration Space Address). For the converse, see Linear address.

**Process** – a program (application) that is protected in its own memory space. Normally when you launch an application program, Windows creates a new Process for that program. The new Process may only access memory in its own memory space, or in the operating system through calls to the Win32 API (Application Program Interface).

**Rings** – Access and instruction execution privilege (protection) levels controlled by Pentium processors.

- Ring 0 – Most privileged protection level on a Pentium. Also known as System Level or Supervisor Level. Code running at this level (such as comemL.vxd) can execute any instruction and access any memory or I/O location, including the page tables.
- Ring 3 – Least privileged protection level on a Pentium. Also known as User Mode or Application Level.
- Ring 1 and 2 – though built into Pentium processors, Ring 1 and 2 are not used by Windows operating systems, or likely by any other operating system.

**Thread** – a single uninterrupted flow of execution through a program.

A Process owns one or more Threads. When created, a Process owns only one Thread. This parent thread may then explicitly create (“spawn”) more Threads as it runs by calling OS services. The Process would then *appear* to be doing more than one thing at a time.

Another key component in this slight of hand is the Windows Scheduler. The Scheduler time slices between all the active Threads in the system. It gets invoked periodically when timer or other interrupts occur. For example, you could write a device driver that would invoke the scheduler whenever the AN3042 DMA completion interrupt fires.

In general, the Scheduler selects the highest priority Thread from the current list of ready-to-run Threads, from any Process.



## References

1. *CO-MEM Lite, AN3042 Integrated Circuit, Technical Reference Manual*. Anchor Chips Incorporated.
2. Intel Architecture Software Developer's Manual *Volume 3: System Programming Guide*. Intel Corporation, 1997.  
The Intel Architecture Software Developer's Manual consists of three books and three addenda, all available at <http://developer.intel.com/design/pentium/manuals> :

Volume 1: Basic Architecture	Order Number 243190
Volume 2: Instruction Set Reference Manual	Order Number 243191
Volume 3: System Programming Guide	Order Number 243192
Addendum: Volume 1: Basic Architecture	Order Number 243691
Addendum: Volume 2: Instruction Set Reference Manual	Order Number 243689
Addendum: Volume 3: System Programming Guide	Order Number 243690.
3. *PCI Local Bus Specification, Revision 2.1*. PCI Special Interest Group, Portland, OR.
4. Shanley, Tom, Anderson, Don. *PCI System Architecture, Third Edition*. MindShare, Inc., 1995. Addison-Wesley.
5. Hazzah, Karen. *Writing Windows VxDs and Device Drivers*. R&D Books, Lawrence, KS, 1997.
6. Shanley, Tom. *Protected Mode Software Architecture*. MindShare, Inc., 1996. Addison-Wesley.