



## APPLICATION FOCUS

# UMA Redux: The 7MB Case in Detail

In our October 1995 white paper *An Analysis of Unified Memory Architecture*, Alliance Semiconductor presented a price/performance model of a typical UMA system, and discovered that the so-called “savings” advertised by UMA’s apologists result in an unreasonable cost in system performance. We also debunked some of the widely promoted myths about UMA, by showing how faster CPUs will inevitably worsen the problems inherent in the architecture, and how new memories like BEDO and SDRAM will barely be fast enough to keep up with these faster CPUs—with nothing left over to waste on graphics display.

The white paper concluded by proposing an “ideal” UMA system. This system includes 16MB of BEDO or SDRAM memory, 256K pipelined burst SRAM cache, and limited display resolution, color depth, and functionality: in particular, no software MPEG, and no software or hardware 3D rendering.

In presenting our findings to OEMs and analysts, Alliance found that many had already studied the UMA question and had encountered these same problems, as well as several other important ones. System manufacturers typically begin by specifying UMA for an 8MB entry-level configuration. Later, faced with performance and operating system related issues, they end up cancelling or retrenching to a higher-cost 16MB platform. This addendum, which deals specifically with the question of 8MB UMA for 1996 systems, addresses the additional problems not covered in our original white paper.

---

### The 7MB Home PC

#### *Cost-Only vs. Price/Performance*

Many system vendors take a cynical view of the retail market. Average consumers, they argue, do not understand price/performance, so a lower-cost system can be sold to them regardless of its performance. For this market, 8MB UMA appears to offer a competitive advantage.

However, as CD-ROM publishers have learned, the uninformed public can exact a heavy revenge. The heady Christmas rush may be merely a prelude to the dreaded December 26 refrain: “I want my money back.” For every SKU returned, the manufacturer has to sell through five to ten additional units just to cover return costs, not to mention loss of reputation and customer goodwill.

Though an 8MB UMA system is a poor Windows 95 performer, it may be argued that low performance alone will not lead to many system returns. In fact, however, the real specter of high return rates arises not from Windows 95, but from PC manufacturers’ old familiar compatibility nightmare: DOS extender games.

#### *The Descent Phenomenon*

Following on the tremendous success of Doom and Descent, game developers have recognized that it pays to take advantage of the capabilities of new Pentium-class machines.<sup>1</sup>

As a result, some developers have given up on 4MB systems for their 1996-release titles and will specify 8MB minimum configuration. Some of these developers have moved to Windows 95 and DirectX; others remain more comfortable with DOS extenders. The well-known 640K low-memory problem with DOS extenders has a workaround: a separate boot floppy, with minimal CONFIG.SYS, can make most titles run. If a system has less total memory than anticipated, however, there may be no workaround for the end-user.

Certainly not all new titles will break when main memory is reduced from 8MB to 7MB. Importantly,

- 
1. According to one title developer’s research, consumers buy an average of 7 titles in the first 2 years after buying a new computer, and only 5 through the remaining lifetime of the system. Since the number of systems has also been growing year by year, clearly the proportion of demand coming from this “recent installed base” is overwhelming, justifying the abandonment of some of the older installed base.



though, none of these new titles has been developed or tested with 7MB memory, so it is likely that some programmers did allow their games to expand into the full 8MB space.

Game developers are acutely aware of the memory issue. To simplify reporting of system requirements and reduce software return rates, as well as to promote the PC as a game platform, more than 50 software developers and more than 40 hardware developers have become members of the new GamePC Consortium.<sup>2</sup> This fall, the Consortium approved its first “GamePC Compatible” specification. Recognizing the 7MB problem with titles already under development, as well as the other limitations that UMA would place on games, the membership voted to specify a minimum of 1MB dedicated display memory in 8MB systems to qualify for the forthcoming “GamePC Level 1 Compatible” logo. The Consortium intends to offer certification of systems and titles to the GamePC standard.

**Table 1: GamePC Level 1 Minimums**

General	<ul style="list-style-type: none"> <li>• Meets MPC Level 2</li> </ul>
Memory	<ul style="list-style-type: none"> <li>• 9MB total</li> <li>• 8MB system</li> <li>• 1MB graphics/video</li> </ul>
Hard disk	<ul style="list-style-type: none"> <li>• 600 KB/sec sustained</li> </ul>
CD-ROM	<ul style="list-style-type: none"> <li>• 300 KB/sec sustained</li> <li>• AutoPlay or equivalent</li> </ul>
2D graphics	<ul style="list-style-type: none"> <li>• 30 fps Fox and Bear [640x480x8-bit]</li> </ul>
Video	<ul style="list-style-type: none"> <li>• MPEG from CD at 15fps [352x240x16 in 640x480x8]</li> <li>• OM-1 compliant</li> </ul>
Sound	<ul style="list-style-type: none"> <li>• Mono/stereo 8-44KHz</li> <li>• 4 simultaneous buffers [22KHz 16-bit]</li> <li>• General MIDI</li> </ul>

2. For details of the Consortium’s membership, mission, standards, etc. see: <http://www.mmwire.com/gamepc/gpchome.html>

## The 7MB Business PC

### *Windows NT’s Oncoming Freight Train*

Business PC buyers, whether professional MIS managers or small business owners, generally are better informed than retail consumers about price/performance and system specifications. However, because many business applications do not entail heavy computing demands, some system vendors have considered introducing 8MB UMA as an entry-level office system. They argue that businesses will be swayed by UMA’s cost savings, believing that an added 1MB graphics memory, though superior in terms of price/performance, is not necessary for day-to-day work.

For a business, however, initial equipment price is only one part of the purchase decision. Usable lifetime, or obsolescence horizon, is an even more important factor in cost of ownership than initial cost. And 8MB UMA systems will be obsolete—or will require expensive upgrades—significantly earlier than 8MB non-UMA systems, because of Microsoft’s intention to drive Windows NT into the workplace over the next several years.

Microsoft is requiring that all Win32 applications be tested under Windows NT as well as Windows 95 in order to receive the coveted Win95 compatibility logo. Bill Gates has stated publicly that Windows 95 and NT code bases will be merged in a future version. A key to the NT strategy, though, is building a large installed base of NT-capable machines. To accomplish this goal, Microsoft must optimize its NT operating system to run in 8MB and must strongly encourage users to buy a minimum of 8MB memory. Making NT run in 7MB main memory, if possible at all, would require a tremendous additional investment and schedule slip, so it is likely that Microsoft will not support the 7MB configuration.

Even in a rich-memory environment, UMA presents problems for a secure operating system. Crash-proofing, one of the major goals of NT and other advanced operating systems, requires memory protection to ensure that one process or thread does not corrupt another. Even device drivers are not permitted unfettered access to memory. A back-door hardware path into system memory, such as that proposed by Unified Memory Architecture, exposes the system to software bugs and potentially to new



types of viruses which the operating system cannot control. So a simple memory upgrade does not necessarily transform the 7MB system into an acceptable NT platform.

---

## The 12MB Case

---

### *Three Bad Options*

Faced with the severe shortcomings of UMA in an 8MB system, several chipset vendors have proposed moving to a nonstandard 12MB configuration. This intermediate configuration could potentially solve the DOS extender, GamePC, and Windows NT problems. However, while conceptually appealing, the 12MB configuration ends up being difficult to implement and difficult to position.

Three different 12MB implementations are being discussed. The first uses three standard 72-pin, 4MB SIMM modules, or equivalent devices soldered on the motherboard. Two 1Mx32 SIMMS form a standard 8MB bank, and one 1Mx32 SIMM forms a second bank. The problem here is that the second bank is only 32 bits wide. If the 32-bit bank is mapped to the low-order address range (0-4MB), the most frequently used locked operating system code and data pages will be accessed at half the usual speed. If, on the other hand, the 32-bit bank is mapped to the high address range (8-12MB), then it includes the shared frame buffer (11-12MB), which means that display refresh and display update accesses take twice as long as in the standard 64-bit case; CPU activity stalls twice as long as the standard UMA model predicts.<sup>3</sup>

The second possible implementation requires two new nonstandard 6MB SIMM modules, each organized as 1.5Mx32, or an equivalent 1.5Mx64 bank on the motherboard. The SIMMs presumably would use 1Mx4 or 1Mx16 devices along with 256Kx16 or 512Kx8 devices. The system would address them like 8MB SIMMs but would not access the last 512K of device address (or 4MB of system address). In this case the SIMMs themselves need decode logic on board, since the external interface is as one bank but internally the module is organized as

---

3. In Alliance's *An Analysis of Unified Memory Architecture*, mapping the 32-bit bank to low memory doubles the *MISS-CLK<sub>STD</sub>* parameter in Equation (1); mapping it to high memory doubles the *GBW* and *GSTALL* parameters.

two banks. The cost of this glue logic, as well as the speed penalty it imposes, make this option impractical.

The third, and least problematic, implementation involves two 4MB SIMMs and two new 72-pin 2MB SIMMs, the latter organized as 512Kx32. The drawbacks here are: first, that two of the four SIMM banks are taken, limiting memory expandability, and second, that 512Kx8 devices are unpopular, while using 256Kx16 memories would still require decode logic to make two 256Kx16s appear as 512Kx16.

Finally, there remains the problem of positioning the system. A 12MB UMA system costs \$75 more to build than an 8+1MB non-UMA system, so it cannot compete in the high-volume entry-level slot. For the midrange market, 12MB price/performance is not as good as 16MB, leaving the proposed 12MB system "stuck in the middle."

---

## Conclusion

---

### *Who Stands to Gain?*

If Unified Memory Architecture doesn't perform as advertised, why are so many companies hyping it? The simplest explanation is that UMA's inventors, after developing an elegant and appealing concept, failed to understand its drawbacks. An analysis based on 486DX4 or P90 CPU, running DOS or Windows 3.1 only, with game titles still targeting 4MB platforms, could reach a different conclusion than we do in today's market.

In evaluating claims, it is also instructive to consider whose interests would be served by a shift to UMA. Chipset companies believe that they can gain competitive advantage by being first with UMA; those who introduce early hope to capture some of UMA's "savings" for themselves, by pricing UMA-capable chips higher than equivalent non-UMA offerings. CPU makers also have an interest in UMA, since a UMA system needs a faster (and more expensive) CPU to achieve the same performance as a similar non-UMA system.

SRAM manufacturers like Alliance Semiconductor would also gain from a shift to UMA. Unified memory systems need more cache memory than non-UMA systems to get acceptable performance, which would increase demand for Alliance's core 3.3V fast SRAM products. Nonetheless, Alliance is not



promoting UMA for 1996 systems, because we believe the offerings currently available are a bad deal for consumers and a bad deal for the PC industry. We believe that enough problems remain with the architecture that attempts to introduce it will hold back growth of the PC industry as a whole, and will create significant opportunities for those “traditionalists” able to ship compatible, robust non-UMA systems in 1996.

### ***Traditionalists Fight Back***

As system vendors join the 1996 battle, we expect the non-UMA camp to make the most of its advantages. In the retail market, system vendors may begin advertising systems with “9MB memory” to distinguish from competitors’ UMA systems. They may also promote GamePC Level 1 compatibility, as well as MPEG and 3D capabilities, for further differentiation. In the business market, resellers may begin touting systems as “ready for Windows NT,” to emphasize that today’s investments in hardware will be preserved.

Alliance Semiconductor is continuously at work on new silicon technologies to reduce cost and improve performance of personal computer systems. Contact Alliance to see how our industry-leading fast SRAMs, high-performance ProMotion<sup>tm</sup> graphics and video accelerators, and new DRAM and Flash memory offerings can reduce system BOM cost without sacrificing compatibility, functionality or performance.

---

### **For more information**

#### **Alliance Semiconductor Corporation**

3099 North First Street

San Jose, California 95134

tel. +1-408-383-4900

fax +1-408-383-4999

Attention: Systems Products Marketing