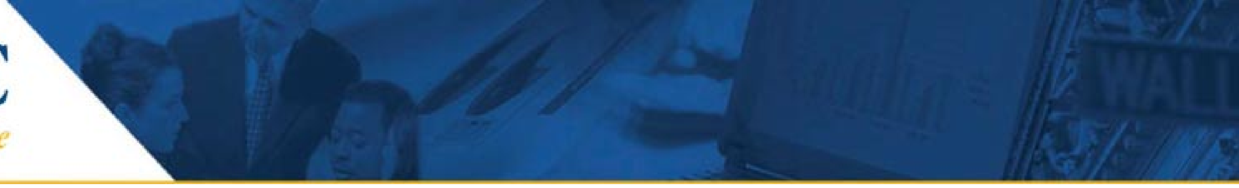

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

The Role of Graphics in the Future of Mainstream Computing

Sponsored by: ATI Technologies

Bob O'Donnell George Bulat
Alan Freedman
October 2005

IDC OPINION

We are on the cusp of a new era in computing, a step up from the period when computing was "good enough" to handle text (e.g., word processing, email) to one where the capability to deal with large digital media files is required. During the next five years, even in commercial environments, digital media — with its large, bit-oriented, less-compressible data types — will become the norm. By the end of the forecast period, robust graphics processing will be standard in the workplace, and high performance graphics, no longer the sole purview of gaming fanatics and entertainment buffs, will be common and expected for handling normal workplace tasks.

Driving this adoption are applications — such as mapping, data visualization, large data set viewing and videoconferencing — that will become widespread. Three to five years from now, companies and public institutions will avail themselves of such applications to conduct ordinary business.

METHODOLOGY

This paper is built around functional specifications, portrayed as "scenarios," for future mainstream commercial applications. The descriptions have been developed based on existing products and demonstrations, knowledge of the directional development of the infrastructure that will support such products, and inductive reasoning as to how such products will likely work.

IN THIS WHITE PAPER

In this white paper, IDC presents the case for the use of 3D graphics in mainstream commercial environments. We discuss graphics — what it is, how it works, how it has evolved, and why it is important — and illuminate a number of future scenarios likely to drive the adoption of performance graphics in everyday business situations.

Although the overall proportion of performance graphics in client PCs (desktops and notebooks) will rise steadily throughout the next five years, corporate systems already run almost entirely on integrated graphics. By integrated graphics, we refer to systems whose main graphics engine is embedded in the system's core logic (sometimes referred to as a combination of the North and South bridges). Core logic directs data traffic between the processor and memory, graphics, and input-output subsystems. In a discrete graphics subsystem, the core logic "talks to" the graphics engine over a system of electrical connections known as a "bus." In an integrated system, this bus and the graphics engine are on the same physical silicon die as the normal core logic functions.

The merits of integrated graphics, first introduced in 1999, have already become apparent. Integration is one of the holy grails of the computer industry. Both hardware developers and buyers value the minimization of three things associated with integrated subsystems: part count (cost savings), the latency associated with longer signal paths (performance), and the heat and power consumption of discrete subsystems (reliability). While discrete graphics maintain a powerful draw for consumers' PCs, integrated graphics have taken over the corporate space simply because they deliver acceptable performance at far greater efficiency and a more attractive price point than discrete graphics.

However, the bar for acceptable performance is rising as new applications that make full use of 3D graphics capabilities become more common among mainstream users.

During the next few years, high performance graphics will come to be used in ordinary commerce in applications such as mapping, data visualization, videoconferencing, and large data set viewing (multiple monitors). By offering high quality graphics experiences to mainstream commercial customers, PC hardware OEMs will be able to bring ordinary PC users in business and public-sector settings up to the next level of performance, matching increased hardware capability with compelling applications.

SITUATION OVERVIEW

How Graphics Processing Arose

Separate graphics engines arose from the need of computers to perform repetitive, yet complex, calculations. These calculations eventually absorbed so much of the main processor's computing power that it had few cycles left over to perform the computing tasks it was actually designed for. At that point, the calculations were offloaded to a specialized processor, the graphics processing unit or GPU.

The basic graphics calculation is about lighting up pixels on a screen. Back in the 1970s, when commercial computing got its start, once the processor arrived at a result, say, a decision to print the words "Name" and "Address" next to each other starting at a screen location defined by an x,y coordinate, another calculation was undertaken to figure out which pixels had to be lit up in order to display this result. The calculation wasn't too difficult in a world where 40 simplified font elements composed of lit and unlit pixels in a 6 by 8 matrix sat next to each other in 25 rows per screen. The graphics calculation would determine that the upper-left-hand-most pixel of the "N" in Name was lit but, after calculating proper offsets, the upper-left-hand-most pixel of the "A" in Address was not. Thus, the calculation was binary: a pixel was either on or off. To display a screenful of data, the graphics processing function that calculated 48,000 times whether a pixel was supposed to be lit or not to refresh a screen 60 times per second involved fewer than 3 million calculations. The math was simple, and, for the machines of the day, not too taxing.

However, modern graphics computation is much more complex. The work still involves decisions about lighting up pixels, but in a 3D color environment, such visual rendering has become far more intricate, subtle, and multilayered. Not only must math be done 60 times per second for a screen (now denominated in pixels rather than characters), say 1,400 pixels wide by 1,050 pixels long, but for each pixel a determination of brightness must be made among 32 choices for each of three colors (red, green, blue or RGB) for each pixel. Thus, the number of raw calculations has gone up to more than 8 billion per second. To make matters worse, exactly what goes into these calculations has grown in complexity. No longer just a question of "to light or not to light," the graphics function must decide how much to light each color of each pixel. In a 3D environment, these results are derived from multiple passes, each pass determining the value of a particular visual effect.

For example, if a room depicted in 3D has a particular pixel on the back wall partially obscured by a foreground element, say, a wisp of smoke, then after the graphics engine has determined the exact way to attain the proper shade of "back-wall brown" by giving the proper RGB values, it has to make another calculation to determine by how much to lighten each value to achieve the effect of foreground smoke at that exact point. The output of this second set of calculations is added to that of the first to obtain modified pixel values. And so on for each effect (backlighting, smoke, water, and so on) that cues the viewer to see a picture as three dimensional. And for each pixel, this set of equations must potentially be run 60 times every second.

Now, obviously, if these calculations had to be made by brute force, they would overwhelm even today's (and tomorrow's) graphics engines. But as the math has become more complex, so has the bag of tricks at the disposal of graphics subsystem designers. For example, if, in a display of video, the movement of an object across the screen does not change the value of a particular pixel, that calculation can be skipped (the instruction from the rendering engine to the pixel painter being "Paint again what you just painted."). Another way graphics engines save on processing is to determine that even though a 3D model of the object to be displayed is maintained in memory, it is not necessary to calculate values for surfaces that would be invisible to the viewer's eye from the view angle to be displayed. Also, by caching some results nearby, a graphics processor can save on some types of repetitive work. For example, math for a complex texture can be maintained and replicated for a defined patch of screen, saving on individual calculations for each pixel and repeated calculations for the same pixel. Thus, through a combination of computational muscle and clever, work-reducing art, the graphics engine is able to produce a realistic full-motion, 3D video effect.

When graphics calculation was split off from normal processing, the only way to supply such a function was as a discrete subsystem, a separate board that communicated with the main processor through an I/O bus. At the highest level, the relationship went like this: the processor would say, "Here's what I want," and the graphics subsystem would reply "Here's what you have to light up (and how much) to achieve it." Moore's Law (which says that the transistor count of silicon parts will double every 18 months) applies to graphics just as it applies to traditional processing, and it is no surprise that by the late 1990s, chip designers were able to squeeze a full 3D engine into a core logic module suitable for commercial computing.

Investment Protection

The theory of investment protection posits that buying "up" at the time of purchase gives purchasers "headroom" inasmuch as they don't have to buy again to be able to handle future standards as they become broadly accepted. Future standards, in this case, refers to graphics capabilities that will enable more visually-oriented interfaces to allow users to comprehend simultaneously many sets of data that may be in motion. For managing large enterprises, this capability is increasingly essential. Applications that make use of today's high performance graphics processors will become common in the enterprise during the next three-to-five years, during which time a major upgrade in graphics software and tools from Microsoft will be made available through the upcoming Windows Vista operating system.

Even for companies who don't immediately migrate to Windows Vista, the need to plan for purchasing both notebook and desktop PCs that are capable of delivering the full benefits of the graphical enhancements to be found in Vista is critical. Enterprise PC lifecycles are 3-4 years in most cases and even longer for some companies, so given that the Windows Vista launch is currently targeted for the second half of 2006, any future PC purchases need to consider the recommended requirements for the new operating system.

In the investment protection equation, hardware is but one part. In the complex world of performance graphics, software has emerged as an imperative element that it is integral to a good user experience. Industry proven driver software ensures consumer and commercial users have broad compatibility and top performance for their commonly used applications. This competency also provides hardware vendors with an important tool to be highly reactive to changing market dynamics and user requirements.

Core logic embedded or used in tandem with the best of today's graphics solutions will be able to stand up to the challenge of tomorrow's application demands. The high end of the current generation of hardware designed for commercial organizations will be able to run future demanding operating environments and applications.

The Effect of Emerging Technologies

Several technologies on the horizon, both near and far, will have an accelerating effect on the adoption of performance graphics in mainstream commercial computing, including the next generation operating system from Microsoft, named Windows Vista, and the introduction of a faster system bus in the form of PCI Express (PCIe).

Windows Vista: A Platform for Advanced Graphics Applications

Windows Vista, the next-generation operating systems from Microsoft, will be key to the adoption of 3D graphics processing in mainstream applications. By the end of 2006, when the first version of Windows Vista debuts, advanced graphics applications will be ready for adoption by mainstream customers.

An example, illustrated at one of Microsoft's Windows Hardware Engineering Conferences (WinHEC), demonstrated what could be done with computational fluid dynamics, a simulation technique that calculates the movement of particles interacting with each other as they flow through an environment. Computational fluid dynamics are involved when creating a realistic simulation of water in a pipe or air circulating through a building. Such, simulations require a lot of data, much interaction among the parts of the simulation, and many feedback loops.

For now, such simulations must run on 64-bit multiprocessor memory-heavy systems to produce near-real-time graphics output, but with 64-bit Windows desktops now available, it is only a matter of time before standard business desktops can run such programs. The door to "personal supercomputing" is opening.

This type of application can be used in the life sciences, medicine, molecular dynamics, architecture, and aerospace engineering, to name a few fields. In the next several years, organizations will being investing in advanced computation of this sort for business intelligence and other areas.

Aiding this type of application will be advances in display resolution and size as well as the ability to make use of multiple displays, all of which will change the user interface radically. Most desktops will have displays that are 20 inches or larger. People coping with huge sets of information will have either large or multiple displays.

Windows Vista software will be "camera aware," dovetailing with hardware developments. Many commercial PCs will have a camera for video conferencing, presence awareness, and meeting recording. Between the various cameras in a conference, software will be able to provide different points of view and record a white board and electronic minutes. Average users will have the kind of rich display that makes the PC the best reading device. Navigation and annotation experiences will be immediate and rich.

Applications Today

One of the most exciting software developments to come along in some time is Google's impressive Google Earth application, a free program that brings high-end mapping and GIS (Geographical Information System) capabilities to the mainstream. Google Earth combines the searching capabilities of the company's namesake search engine with satellite imagery and 3D terrain maps to create a complete visual map of the Earth with a level of detail that was previously only available to a privileged few. Users of the program, which is also available in more feature-rich paid versions, can type in virtually any location on Earth and then see a satellite image enhanced with 3D terrain information of that spot. You can also easily zoom in or out to any location and the program will automatically adjust the screen to display the further in or further out view. As simple as this may sound, the visual impact is stunning.

In order to provide these realistic map images, Google Earth makes extensive use of 3D graphics hardware, for both texture mapping and shading. The program requires an OpenGL and/or DirectX 8 or higher compliant driver and prefers at least 32 MB of dedicated video RAM (and can use as much as 512 MB) to perform its extensive texture mapping operations. Although it can run on some PCs with only integrated graphics, performance is compromised—the program really comes to life with dedicated graphics controllers. As consumers and business users start to realize the capabilities that detailed mapping software such as Google Earth offers, we expect applications like these to become an important tool even for mainstream business computer users.

Another area where high-performance graphics will make a difference can be seen today in the products of Instant Effects, which makes a package called OfficeFX that adds 3D extensions to PowerPoint presentations. To run these extensions, a system needs 3D acceleration. Although awareness is still low for these products, they will become more mainstream over the next few years.

OfficeFX places advanced graphics within reach of ordinary business people, opening the way for a broader use of 3D tools. Any computer user can integrate studio-quality, real-time 3D graphics into a presentation. With full PowerPoint compatibility, OfficeFX lets the user add dynamic 3D backgrounds and foregrounds to a presentation, making use of 3D shading and lighting. For venues in which the medium is the message, presenters can make use of rich media PowerPoints to deliver the sort of compelling presentations previously achievable only with complex and costly rendering — that produced a presentation that could be used just once. OfficeFX delivers the high production values of a telecast with known, customizable tools that produce a consistent look and feel by using a thematic template.

The Window Vista User Interface

It's not only the applications running under Windows Vista that will be more graphically rich—so will the operating system itself. In fact, one of the major highlights of the forthcoming operating systems is its new graphically intensive look and feel. Windows Vista promises to bring a rich, 3D experience to everyday business computing with graphical enhancements, such as "see-through" windows, that not only make the screen more attractive, but easier to use.

Windows Vista's graphical interface will be made available in two forms and they will be determined by the graphic hardware capabilities of the computer in use. For systems with only basic graphics hardware (as well as for any user who prefers), Windows Vista will offer a classic interface that looks and feels exactly the same as Windows XP. The next tier of graphical experience—which will require a graphics subsystem capable of supporting DirectX 9.0 or later as well as a Windows Vista Display Driver Model (WVDDM)-compliant driver—offers the enhanced, 3D view, showing translucent windows, live previews of applications and documents, as well as other graphical enhancements.

Windows Vista's 3D-centric architecture allows images to run smoothly in multitasking environments. In addition, the new operating system will rely on 3D graphics for both performance and stability. 3D, formerly the domain of gamers, game developers, and workstation gurus, will become key to Windows Vista. Crashes, which will be rare, will not cause the graphics subsystem to hang because it will be able to do a snap-reset. The user will not even notice the failure. Under the Windows Vista hood, Windows Graphics Foundation (WGF) will have a "common shader core" that will contain unified vertex and pixel shader operations, making it easier for developers to create 3D elements. The Windows Vista Display Driver Model will:

- allow graphics processor sharing by multiple threads,
- reduce driver complexity,
- enable seamless hang recovery,
- improve monitor support,
- enhance the boot experience,
- permit high-resolution graphics earlier in the boot process,
- support programmable DirectX shaders,
- improve fast user switching,
- allow driver upgrades without reboot (Plug and Play Stop support),
- improve power management, and
- allow virtualized graphics memory, and
- support PCI Express (PCIe).

Hardware advances and software advances are coming together in parallel. On the hardware side, in addition to advanced graphics acceleration from companies like ATI, infrastructure is being put in place that will enable mainstream systems of the future to handle these increased loads. Prime among them is PCIe.

PCI Express: Wider Pipes for Delivering Real-time Multimedia

Aside from graphics hardware itself, the most important hardware development that will enable performance graphics in future platforms is PCIe.

PCIe represents the reunification of the graphics bus with the rest of the system transports. With the existing system, called AGP, only limited uplink bandwidth is available. PCIe, with runs full duplex at 2.0Gbps, is far faster than AGP, which runs at only 170Mbps upstream. PCIe runs downstream twice as fast and upstream 15 times as fast as AGP 8x.

Other Important Infrastructure Elements

Of the key infrastructure elements, PCIe will have the greatest effect on graphics performance since visual traffic will travel over the PCIe bus (or channel). However, the reader should be aware that a number of other elements that are already contributing and will continue to contribute to the higher performance that will be necessary for a transformed visual world. These elements include:

- **Serial ATA (SATA) and Serial ATA II —** high speed communications links that connect the computer to its storage
- **USB 2.0 —** a high speed communications link that connects the computer to external peripherals such as cameras, storage, and printers
- **Gigabit ethernet —** a high speed communications link that connects the computer to the network
- **Faster memory and buses —** Memory types and speeds will continue to evolve, and the internal transports, or buses, that carry data around inside the system will advance to accommodate other elements such as SATA as they increase in power, thus removing potential inhibitors to overall increased system performance.

FUTURE OUTLOOK

Scenarios

Based on the direction of technology development, we have developed several compelling scenarios for why and how users of commercial systems will increasingly turn to integrated graphics subsystems capable of producing realistic 3D effects. These scenarios were developed with knowledge already available about coming technologies, computer usage models, and trends in marketplace developments, to which was applied a process of inductive reasoning to arrive at a likely picture of the future of commercial computing.

There are many scenarios that highlight the need for 3D graphics processing in the commercial systems of tomorrow, but three are truly compelling: data visualization, the ability to view large data sets and videoconferencing.

Data Visualization

The scope and quantity of data that managers must cope with rises all the time. These vast stores overwhelm human capability to absorb them. An ordinary mortal simply can't keep all relevant facts in his or her mind at the same time. For this reason, data visualization — the ability to see huge amounts of information as a picture, even a moving picture — will rise in importance. The following section illustrates how one such visualization application would work.

Fly Through of My Empire

The chief operating officer (COO) of a large U.S.-based retail chain was looking at his company. What he was actually looking at was a huge flat-panel monitor with a glowing map of the United States on it. Each state was depicted with a sense of its topography: Colorado was green and uniform on its eastern half and brown and textured on its western half. California showed strips of gnarly brown and uniform green where the Sierra Nevada and other mountain ranges ran parallel to the Central Valley. Florida was solid green. Each state had a ball apparently hovering over it in 3D space. The size of the ball represented the current month's dollar volume of the chain's operations in that state. California and New York had large ones. South Dakota and Maine had small ones. Hitting a key caused a second ball, representing the previous month's volume, to superimpose on the first one. A different key called up the prior year's volume in the same month. Hitting the spacebar in each view caused the national average to show up as a red circle on top of each ball so that the COO could determine whether a given state was under or over the mean at a glance.

The COO toggled the views, watching to see if patterns were normal. The month-on-month figures showed a general decline, but the COO called up yet another view that showed, with purple circles, average monthly patterns for the past decade. He decided that the pattern related to normal seasonality: the averages indicated that sales declined from January to February as the effect of post-holiday sales tapered off and midwinter set in. But the year-on-year comparison showed a larger-than-average decline in a group of states in the Southeast. Most states were up, but a cluster of them was definitely down. He grabbed the mouse and started to "fly down" into the Southeast region. The scale of the map changed continuously and smoothly as the region came into focus. Of interest was that, at this level, he was able to determine that the states with the unusually poor results were all under the aegis of the same regional manager.

Another valuable indicator on the COO's screen was the color of the balls. Color indicated inventory issues. Green meant levels were normal, yellow was a warning of overstock, and red meant a stock-out. In the Southeast region he noted a number of yellow and red balls mixed in with the green. He "flew up" to the national level again and "flew down" into the other regions, one at a time. They were mostly green, with an occasional red or yellow, but never both in the same region. Because inventories in the Southeast were all over the map, so to speak, and the problem was isolated to one region, the COO decided that he had a management problem. He lifted the phone and called the national sales manager.

The map he was viewing was a composite rendered in real time with inputs from satellite cameras; local cams on location in stores, offices, warehouses, and parking lots; and 3D interpolation at each level of scale. Each level showed as a realistic view of the world from that point of view. With lots of detail, he could fly down through the layers to look at actual stores and parking lots. The system allowed the viewing of overlapping maps and visualized data. By setting certain views to 30% or 50% transparency, he could see data overlapped on photos or other representations, including publicly available data such as customer demographics, and company info such as store payroll.

Turning back to his map, he wanted to determine what exactly the inventory issues were, whether there were certain items or categories that were causing problems, whether stock-outs of certain items could be attributed to one-time conditions (e.g., a run on batteries after a big storm, a stock out on tires after a big promotion) or were endemic (e.g., a long-term trend away from preparing home meals was causing permanent elevated demand for complete frozen dinners).

At the 50-mile level, he viewed his whole empire with the inventory overlay, noting a region where a red X indicated stock-outs of a single item. As he flew down through the levels, he noticed that the problem seemed to emanate from Texas. All the stores in the Houston area seemed to have the same issue. So, he chose the store on Travis and zoomed in on it. Right away, he could see a full parking lot, and as he "entered" the store, he could see that the red X was hovering over the outdoor equipment shelf. One notch further in, and he saw that the stocked out item was insulated containers, portable ice chests. Clicking on the item itself opened up a view of the database behind the visual system, revealing a history in calendar form of ice chest shipment dates, quantities, prices, and inventory levels. He saw that the problem had been happening for two weeks, pretty steadily. He had a suspicion of what the issue was. Zooming back out, he called up a map of temperature and summaries of weather-related news for the Houston area and, sure enough, learned that the city had been experiencing higher-than-normal temperature for two weeks and that the city's electricity had been on and off due to several violent thunderstorms that had caused trees to drop onto power lines. Thus, the run on ice chests.

He turned to a weather forecast for the Houston area and discovered that the heat wave was due to break in two more days. Based on this information, he decided not to change the overall stocking policy of ice chests in Texas, but just to have one more replenishment sent. He picked up the phone and called his operations manager.

Viewing Large Data Sets

The utility of a wide array of applications — financial trading, patient monitoring in hospitals, strategic and tactical military communications and decision support, modeling for drilling, manufacturing, architecture and graphic design — rises as the amount of space available to display results increases. In each case, two or more monitors sharing the real estate best accommodate the sheer amount of information that needs to be accessible visually. These large-surface, multi-monitor environments are enabled by high performance graphics. As the scenario below illustrates, one of many that might be considered, the convenience and productivity afforded when two or more monitors share the viewing area has real implications for the multi-tasking graphic designer.

A Jump to the Left...and a Toggle to the Right

When Tilman first joined the studio as a proofreader he felt rather special, being the only employee besides the head designer to have a multi-monitor set up. Of course, having two monitors (with twice the space available to him) came in very handy as he jumped between drafts, now with this layout and text wrapped around it, now yet another with more paragraphs requiring proofreading. The designers, he felt, were always wont to be "creative"; they were artists, after all. But this proofreader was ready for the rapid succession of drafts, especially around Annual Report Season, able even to accommodate the corporate liaisons who also came up with their own seemingly never-ending "ideas", resulting in more changes midstream. Companies have an image to preserve, a message to convey (or so he was always reminded).

Tilman had an advantage, though. As soon as the latest draft was emailed to him, up it went on the left hand monitor, the eventual successor of the work already rendered, now sitting forlornly on the twin monitor to the right, not quite rejected, merely relegated to the status of "reference" document. There was a surprisingly large amount of text to wade through – often in both English and French, which itself posed a problem since French words typically take up more space than their English equivalents and changes to either can affect the designer's layout.

Thanks to the IT manager of this graphic design shop, who realized that the proofreader could benefit from having two monitors at his disposal, this potentially tedious and time consuming job – making sure that every word has the proper spelling and every sentence is grammatically correct – was made easier. The dual monitors enabled Tilman to fly through one draft after another as minute corrections to the text were readily identified and effected. The studio owners soon noticed Tilman's facility as a proofreader and (unusual in a proofreader) his inclination to suggest possible changes to the layout itself. It did not take the owners long to realize that not only did they have a proofreader with a design sensibility but that multi-monitor environments had a wider application and that perhaps others within the studio besides the proofreader could benefit. Tilman was soon promoted to Junior Designer, keeping his multi-monitor system to render a whole new set of tasks using relatively sophisticated applications. He was not alone, though, as now every graphic designer, production artist, and Web designer was equipped with the same. Oh yes, and so was the new proofreader.

Words, words forming sentences, sentences making up paragraphs; now there were images to contend with, photos and drawings to select, a particular layout, or color, or typography, one of many that might "work" aesthetically – and, of course, always, a message to convey. Although most of this was new to him, Tilman was up to the challenge. Where before the multi-monitor application was confined largely to text, enabling Tilman to easily shift from one draft to another, this new application was more complex. Instead of simply words, there was a comparatively rich visual content that needed to be manipulated, refined, and finally rendered within the context of a suitable layout. The multi-monitor system enabled Tilman to have two independent (and sometimes very different) views of the same layout. The high-end 3D graphics also allowed him to clearly see detail, without having to zoom in, and to overlay images, juxtapose them and move them around. As Tilman quickly learned, having two monitors affords larger production real estate than you would have with only one. Indeed, Tilman was able to maximize one as his primary monitor for production (by force of habit, the left hand monitor). This was the monitor dedicated to the finished layout, effectively his "working file", where he would view changes to the layout as it evolved from early to final stages. The right hand monitor, or secondary viewing area, is where the various palettes resided and where he could open multiple applications, such as Quark Express and Adobe Photoshop. This monitor would serve as the space where Tilman would vary the layout, choose styles for type or build color profiles, transferring those elements that he would finally incorporate in the primary layout to the left hand monitor.

Just as Tilman was becoming accustomed to the world according to Quark Xpress, Adobe Photoshop, and Adobe Illustrator, the owners of the graphic design studio decided to expand their roster of clients and to move into video production as well. They also determined that they would need someone special for this role, someone who could combine often-disparate graphical elements and make appropriate aesthetic choices. They needed someone who could also move from pre- to post-production without losing stride, who was a born editor, someone too that was well-versed in multi-monitor use – you see, there was this new application …

Collaboration and Conferencing

The benefits of video conferencing, a credible substitute for expensive and time-consuming travel, are generating significant demand in the market, and therefore, significant investment amongst the vendor community. Users envision the day when a video conference will be a viable method of conducting "face-to-face" meetings. The ability to clearly recognize intonations, body language, audience comprehension and visual aids are critical to the success of today's personal meetings, and need to be on par, or superior for the success of video conferencing. Insufficient refresh rates, and unsatisfactory 3D rendering capabilities are technical issues that require attention before mass adoption of products will occur. The capabilities of 3D graphics processors begin to address the limitations of today's offerings with more realistic visual acuity and "real-time" performance. 3D-capable graphics processors can improve the quality of 2D video by offloading work from the CPU, as well as accelerating the playback of the video frames. The smooth functioning and high quality of the audio and video, as illustrated in the scenario below, will make videoconferencing a viable option.

Saving Face — The New Videoconferencing Archetype

A smile came across the CTO's face as she went to turn on the recently installed video conferencing unit. The application had never failed to generate "oohs and ahhs", as participants raved about the clarity of both the picture and the sound. More importantly, video conferencing had already saved the organization a tremendous amount of money on travel expenses, and was making the branch operations more efficient and accessible.

In evaluating competing products, the CTO came to realize that most conferencing technologies delivered a static, 2-dimensional presentation that enabled the conference leaders to control the flow and content of the meeting. Real value-add requires a much more robust set of features on the technology side, and a commitment to process change on the users side.

As the meeting progressed, the CTO marveled at how the technology automatically sensed, by tone of voice, who was the primary speaker and smoothly transferred that person's image into the main window. All the while, other participants were shown around the edge of the screen, allowing all the attendees to gauge reactions to the discussion, introduce visual aids or exhibits, and even see who was dozing off. The ability to incorporate data, decision support tools and wireless communications was just icing on the cake. It was then that a new feature became apparent; the capacity to have two main speakers enlarged in the centre as two members of the executive committee engaged in a heated debate. The sensation was so realistic she could see all the other people cringe, it was as if all the participants were in the same boardroom.

Other applications she had seen had trouble keeping the graphics up to speed with the audio, forcing speakers to pause, causing the presentation or discussion to lose flow and the usefulness of the meeting to plummet. Performance 3D graphics capabilities in the conferencing unit and at the users' sites resulted in seamless transitions, synchronized video and voice, and most importantly, user confidence in that effective meetings can take place in a virtual environment.

3D conferencing is applicable in many more scenarios than just for the replacement of traditional meetings. Learning applications for one, thought the CTO, will benefit tremendously by the usage of 3D simulations providing a more realistic, plausible and effective visually-based training environment. In fact, collaborative applications such as graphics design, engineering and even everyday marketing communications can be enhanced with 3D capabilities. 3D images provide a richness and attractiveness that enhances spatial understanding and promotes differentiation and recall, key attributes in advertising and marketing. The CTO recalled some applications, such as Siemens OpenScape, that merge real-time communications with messaging, business information and conferencing, resulting in improved personal productivity by removing inefficient processes and providing timely information to the right users.

The meeting ended the same way that all of them had for the past two weeks, with all participants commending the CTO on her selection of conferencing technology. Morale and operational efficiency was up, travel and related expenses were way down. The utilization of 3D capabilities was just beginning, and the CTO was starting to target her next home run application.

CHALLENGES/OPPORTUNITIES

Some of this picture of the future is Windows Vista dependent, and delays in Windows Vista could set our scenario rollouts back. The main challenge for the industry will be to educate customers as to the benefits of enhanced computing platforms. While not wanting to sacrifice anything in the way of cost, a user will, over time, come to expect the kind of performance graphics that make possible video-collaboration, visualizing masses of data, and being able to see huge amounts of detail simultaneously.

Companies such as ATI supplying advanced graphics capabilities as add in cards and embedded in core logic must convince risk-averse corporate buyers that their logic is as good as any other available platform on the market. For example, ATI's DirectX 9 core, which has reached maturity in discrete subsystems, is now also available in an integrated format with robust, user-friendly software.

CONCLUSION

It is likely that all suppliers of infrastructure silicon — core logic and graphics — will strive to supply the highest performance graphics that can be fit within the exacting envelope of a commercial environment. Companies with a head start on the graphics side have an advantage because of the amount of intellectual property represented by advanced graphics technology. As long as end users are able to qualify the core logic functions of such suppliers, they are likely to adopt, subject to positive results of price-performance analysis, the best graphics solution they can get.

Users buying systems now will want to make sure they are 3D-graphics capable, even in advance of widespread availability of applications that make full use of such capabilities. IT managers want to get the maximum life out of the systems that they buy, and one way to do that is to ensure that there is sufficient "headroom" in the systems they specify to accommodate the computing loads of tomorrow.

Although Windows Vista won't hit the market until the end of 2006, systems bought today will be only a year old when it arrives. Buying advanced graphics capabilities now will ensure that today's systems will still be viable tomorrow.

Copyright Notice