

# Design considerations for Digital's PowerStorm graphics processor

---

by C. Gianos  
D. Hobson

**The specific development goals for the Digital Equipment Corporation PowerStorm™ graphics processor were improved performance, low product cost, quick time to market, and backward compatibility with existing user software. Achieving these goals required the evaluation and implementation of many new features, enhancements to the existing architecture, and improved development techniques. This paper describes several of the more notable aspects that were considered and includes a discussion of how the underlying technology played a role in meeting the product goals.**

## Introduction

The PowerStorm™ graphics processor is the fourth member of a family of graphics processors based on the Digital Equipment Corporation Smart Frame Buffer (SFB) architecture, as shown in **Figure 1**. All members of the family target the entry level of the workstation market. This market requires that graphics options provide a cost-sensitive solution with high performance. These options


are targeted at applications in scientific visualization systems, electrical CAD, and mechanical CAD markets. The performance of these applications is dependent on the acceleration of the two-dimensional (2D) lines, 2D filled areas, and three-dimensional (3D) wire-frame primitives, as well as the overall performance of a windowing system such as the X Windows® System.

The first member of this processor family is the TURBOchannel™-based 2D HX option, which was introduced at the same time as the first Alpha workstations using the 21064 processor. As with RISC microprocessors, these designs trade off design complexity between the software and hardware design to allow the hardware to implement the simple primitives that make significant performance impacts [1, 2].

The next generation of the SFB family was the ZLX-E series, introduced with the second generation of Alpha workstations based on the 21064A processor. This product was designed to increase 2D performance and to introduce respectable, entry-level, 3D performance, with features such as 12/24-bit visual types, double buffering, and Z-buffering. These 3D features were needed to address the growing needs of the mechanical CAD market segment. The application-specific integrated circuit (ASIC)

©Copyright 1996 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/96/\$5.00 © 1996 IBM



**Figure 1**

Smart Frame Buffer Architecture roadmap.

design for this generation allowed for multiple options from the same chip. Each option, with added memory and more complex video digital-to-analog conversions (videoDACs), provided additional features for applications that could take advantage of them. The entry-level option retained its low cost for cost-sensitive users, and reduced the number of hardware and software resources necessary for product development.

The emergence of the peripheral component interconnection (PCI) standard, and its adoption by Digital workstations, led to the ZLXp-E processor. This generation used the same graphics engine as the ZLX-E series, but replaced the TURBOchannel interface logic with a PCI bus interface. These options provided features such as VGA pass-through, and allowed the economies of the PC cost structure to be capitalized in the workstation market.

The PowerStorm series, which evolved from the ZLXp-E series, was designed to complement the 21164-based Alpha workstations. The primary product requirement was to increase both 3D and 2D performance. VGA support, multimedia, a Windows-NT™ application program interface (API), and higher monitor (1600 × 1280) resolution support were added to this series.

**Table 1** summarizes the features, and **Table 2** summarizes the performance, of the SFB family.<sup>1</sup>

### Design considerations

There were several factors that formed the PowerStorm graphics processor. These included consideration of new graphics engine features and improvements to the PCI bus interface implementation, as well as pixel generation,

memory control, frame buffer memory, and finally, ASIC verification. Details of each of these aspects are now discussed.

#### • Graphics engine feature acceptance criteria

The criteria by which features would be evaluated and accepted for incorporation into the product were one of the first major design considerations. Application performance was a primary goal. A graphical application is typically written to standard interfaces such as X11™ and OpenGL™. Device-specific interface software isolates the application from the hardware and translates X11 or OpenGL protocol requests into operations directly supported by hardware. From an applications perspective, only the throughput of the combined software and hardware system is important. Since the software was typically running on a high-performance Alpha-based workstation, great care was taken to select only those features in which hardware acceleration actually increases application performance. Features that add complexity, risk, or time without appreciable benefit to the application were avoided.

Short time to market was another goal of the product. This, coupled with a relatively small design team, caused features that could not be easily integrated into the existing architecture to be rejected. Because software availability was a key element in hardware verification and debugging, features requiring a manageable change to the hardware, but more significant software changes, were also rejected.

The cost model was a less significant factor in feature consideration for several reasons: First, the workstation market is not as cost-driven as higher-volume PC markets. Cost trade-offs can be more easily justified where value is added. Second, the frame buffer memory chips comprise a significant percentage of the total product cost; few factors actually had a significant impact on total product cost. The basic cost model of "less than or equal to previous generations" was essentially formed by the time-to-market and backward compatibility goals.

#### • PCI bus interface

Targeting a PCI local bus environment was a more difficult problem than was originally anticipated. Optimizing utilization of the bus to achieve maximum throughput, integrating support for the VGA protocol, and simply meeting the timing and electrical requirements of the PCI were the primary areas of focus.

From a graphics perspective, the two most important bus operations are writing commands into the chip and transferring images (data) back to system memory. Commands are written directly by the CPU to a defined address space. The programming sequence necessary to perform different operations varies with the operation and

<sup>1</sup> Information about the SFB family of processors is available on the Internet World Wide Web at <http://www.alphastation.digital.com/announce/graphics.html>.

**Table 1** Graphics features of product options in the SFB family.

| <i>Option</i>   | <i>Frame buffer<br/>memory<br/>(MB)</i> | <i>Overlay<br/>planes</i> | <i>Frame buffer<br/>depth<br/>(bits)</i> | <i>Z buffer<br/>depth<br/>(bits)</i> | <i>Double buffer<br/>depth<br/>(bits)</i> |
|-----------------|-----------------------------------------|---------------------------|------------------------------------------|--------------------------------------|-------------------------------------------|
| HX              | 2                                       | 0                         | 8                                        | Software                             | Software                                  |
| ZLX(p)-E1       | 2                                       | 0                         | 8                                        | Software                             | Software                                  |
| ZLX(p)-E3       | 16                                      | 4                         | 8, 12, or 24                             | 24                                   | 8, 12, or 24                              |
| PowerStorm 3D30 | 2                                       | 0                         | 8                                        | Software                             | Software                                  |
| PowerStorm 4D20 | 16                                      | 4                         | 8, 12, or 24                             | 24                                   | 8, 12, or 24                              |

**Table 2** Unix graphics performance summary of the SFB family.

| <i>Option</i>   | <i>Machine</i>         | <i>Xmark</i> | <i>PLBwire</i> | <i>PLBsurface</i> | <i>3D vectors<br/>(M/s)</i> | <i>3D triangle<br/>(K/s)</i> |
|-----------------|------------------------|--------------|----------------|-------------------|-----------------------------|------------------------------|
| HX              | DEC3000 Model 500      | 7.43         | —              | —                 | —                           | —                            |
| ZLX-E1          | DEC3000 Model 900      | 18.06        | 89.5           | 44.5              | 2.60                        | 66                           |
| ZLX-E3          | DEC3000 Model 900      | 17.50        | 102.1          | 45.7              | 2.60                        | 66                           |
| ZLXp-E1         | AlphaStation 250 4/266 | 17.03        | 86.0           | 45.2              | 2.56                        | 80                           |
| ZLXp-E3         | AlphaStation 250 4/266 | 14.80        | 101.3          | 52.6              | 2.57                        | 80                           |
| PowerStorm 3D30 | AlphaStation 600 5/333 | 33.07        | 168.2          | 101.7             | 3.45                        | 256                          |
| PowerStorm 4D20 | AlphaStation 600 5/333 | 28.04        | 185.6          | 134.3             | 3.43                        | 257                          |

can cause many small transfers to discontinuous addresses on the PCI bus. Because compatibility with existing software was a design requirement, it was necessary to optimize bus utilization in the presence of this behavior. The bus interface was designed to perform what the PCI calls a "fast DEVSEL," or fast decode operation. In other words, a transaction would be accepted on the first possible cycle rather than stalling the bus to complete the transaction decode process in a subsequent cycle. For a one-word transaction, fast decode could make the transaction up to 33% faster. Unfortunately, the PCI timing specification allows only 7 ns to receive, decode, and set up the acknowledgment of a fast-decode transaction. Careful design and fast gates were required to meet this goal.

When data are transferred across the PCI bus, utilization is maximized if new data are transferred every clock cycle. Historically, driving a new data word onto the bus using consecutive cycles has been difficult because of the simultaneous switching characteristics of inexpensive packages and the drive characteristics of PCI-compliant drivers. Because of this, the previous-generation design was forced to use two cycles to drive each data word, effectively halving the transfer rate. Selecting a packaging technology capable of transferring data at full bandwidth was an important consideration.


Graphics options intended to operate on a variety of PCI-based platforms are generally required to support VGA, a low-level interface common to most, if not all, PCI platforms. Since device-specific software is often unavailable during initial system boot sequences, graphics

adapters must support VGA in order to facilitate the display of messages during these sequences. The PowerStorm design team lacked access to a usable VGA design and lacked the experience to develop a custom VGA functional unit from scratch. Because ISA bus VGA chips are low-cost commodity items, using an external VGA chip with integrated DAC and a small DRAM frame buffer was an attractive option. Costwise, it would have been better to have an integrated solution, but time to market and product risk mandated the use of an external solution.

The ability to operate in a variety of Alpha- and Intel-based PCI platforms imposed the additional requirement of compatibility with both 3.3-V and 5.0-V signaling environments. Simple 5-V tolerance is not enough; meeting the required *I-V* characteristics often necessitates using a split power rail for the PCI drivers. A technology compliant with both signaling environments was an attractive alternative.

#### • Pixel generation

The previous-generation ZLX processors introduced some basic features necessary for hardware acceleration of 3D graphics. True-color pixel processing, linear interpolation of color, and depth buffering were among the features added to the architecture, but the fundamental set of primitives remained 2D lines, stipples, and area copies. In the development time frame of the PowerStorm processor, it was clear that the market would require 3D performance beyond what was achievable using the available 2D primitives. Traditional mid-range 3D graphics



**Figure 2**

The unique frame buffer organization arranges four independent memory controllers to minimize the cost of crossing pages in memory when drawing lines.

processors provide direct hardware support for primitives such as triangles and polygons, but require some specific and complex logic to decompose and render these objects. For consistency with the Smart Frame Buffer philosophy of carefully balancing what is best done by software and what is best done by hardware, a triangle-span-mode operation was added. Rather than add the full hardware overhead of complete triangle rendering, span mode takes full advantage of the tremendous line-drawing performance of the existing architecture by rendering triangles or almost any other complex polygon with spans of lines. Software no longer must compute or pass the complete set of parameters for each span, and the hardware must add only storage and control for 2D interpolation, but can reuse the existing interpolation logic to step from span to span. Although this is not the highest-performance solution, this feature significantly increased the 3D rendering capability without significant impact on the time-to-market requirements.

Although the ability to process true-color pixels introduced with the ZLX generation enabled the products to compete in imaging markets, it was the AccuVideo™ dithering technology that attracted the most attention [3]. To further penetrate the imaging markets, the pixel generation path in the PowerStorm processor was enhanced to include color space conversion and image-

scaling capabilities. The color space conversion logic was required because many imaging software applications manipulate data only in YUV formats. Image-scaling and filtering capabilities are also common, especially in video editing and videoconferencing, and can significantly reduce the system bus overhead if accomplished with hardware. For example, consider a videoconference where a QSIF<sup>2</sup> (160 × 120) YUV image is transmitted between participants, but each participant displays the image in a 640 × 480 window. The transmitted image requires only about one megabyte per second to achieve full motion, but uses more than 35 megabytes per second of display bandwidth. Providing hardware assistance for these features was necessary to support and compete in the multimedia and imaging markets.

As with today's RISC CPU designs, increasing clock speed was an obvious way to achieve additional performance. The ZLX family operated the pixel generation logic at 38 MHz. Performance estimates were modeled for faster clock rates and the initial target was set at 50 MHz, but with careful attention to timing during logic design and a much faster ASIC technology, the final product shipped at 75 MHz, approximately twice the clock speed of the previous generation.


#### • Memory control

Prior generations of the Smart Frame Buffer architecture have all used traditional fast-page-mode video RAMs, and have emphasized maximizing the available bandwidth for rendering. The previous-generation ZLX processor was developed to use a (patent-pending) method of organizing and accessing frame buffer memory that significantly improves the available bandwidth for objects such as lines.<sup>1</sup> The actual implementation, however, was forced to limit the memory access because of a lack of pins. The critical element of this feature was the ability to operate the 64-bit interface to memory as four totally independent 16-bit slices. By arranging frame buffer memory so that neither vertically nor horizontally adjacent pixels are serviced by the same 16-bit slice, the cost of crossing pages in frame buffer memory can be overlapped with the painting of previous pixels by other slices (Figure 2).

Removing the physical pin limitations of the 240-pin plastic quad flat package (PQFP) used for ZLX not only allowed a truly independent slice design, but removed all of the external components that were required to latch and decode values on the multiplexed address/control buses, resulting in the module block diagram shown in Figure 3. Latency was reduced by eliminating stalls when the multiplexed bus was in use, the design was simplified

<sup>1</sup>Information about the SFB family of processors is available on the Internet World Wide Web at <http://www.alphastation.digital.com/announce/graphics.html>.  
<sup>2</sup> QSIF—quarter size of the standard image format (SIF), 160 × 120 for NTSC, 192 × 144 for PAL.





**Figure 4**

Block diagram of the PowerStorm ASIC.

**Table 3** Bandwidth comparison of potential memory solutions.

|                | Peak transfer rate (Mb/s) | Streaming transfer rate (Mb/s) | After screen refresh (Mb/s) |
|----------------|---------------------------|--------------------------------|-----------------------------|
| Rambus         | 500                       | 350                            | 260                         |
| Fast page VRAM | 152                       | 142                            | 132                         |
| EDO VRAM       | 320                       | 310                            | 300                         |

alone. Single-port memory structures force the screen refresh circuitry to share the data bus with the rendering operations. This meant that DRAMs were too slow to support workstation displays, SGRAMs were not expected to be available in the proper time frame, and although Rambus might appear attractive with its low-pin-count 500-megabyte-per-second interface, when the protocol overhead costs and screen refresh were factored in, it would have necessitated using more than two channels. Of the dual-ported structures, only video RAMs were viewed as stable and price-competitive. In addition, extended data out (EDO) features which improve random port bandwidth were becoming available on VRAMs. Assuming an ideal memory subsystem, the results of comparing potential bandwidth of the viable solutions is shown in **Table 3** for the most common workstation class resolution of  $1280 \times 1024$  with 8-bit pixels.

With project goals including cost and time to market, we decided to continue the use of video RAMs, and to exploit EDO modes and next-generation ASIC processes to shorten the cycle times as much as possible. The result of all of the above considerations is shown in the block diagram of the PowerStorm ASIC in **Figure 4**.

#### • Verification

Verification of both hardware and software is becoming one of the most difficult and time-consuming portions of the product development cycle. Many design groups have as many (or more) resources dedicated to verification as to traditional design tasks. The PowerStorm development team was a small group of experienced hardware designers, several software engineers responsible for enhancements to the interface software and device drivers, and no dedicated verification resources. The team considered many verification alternatives before settling on a verification strategy.

Building upon the experience and work from the previous generations, C was adopted as the modeling and simulation language, and a C model of the ASIC was developed. The PowerStorm products are PCI option cards with a simple read/write interface. The C model provided a similar interface by defining access routines called `BusRead` and `BusWrite`. The device-specific software was designed to send all read or write operations through the `BusRead` and `BusWrite` routines, allowing the ASIC model to be used in place of functional hardware. When hardware became available, only the `BusRead` and `BusWrite` routines were recompiled, and the same interface software was operational. The C model allowed production quality software to be developed in parallel with the hardware, increasing the product quality and reducing the time to market.

The usefulness of the software development environment described here extends beyond software

verification. Notice in **Figure 5** that the interfaces between the software model, the X server, and the frame buffer map directly to the pin interface of the ASIC. Simple routines that log the activity at the interfaces were incorporated, allowing the generation of traces equating roughly to stimulus and response. Each time the software developers ran an application, the traces were generated. If the display was correct, the traces were saved and formed the basis for the hardware verification test suite. When the final, fully structural hardware model matched the response of the software model, the hardware design was considered verified. With this method, the development environment virtually eliminated the need for a verification team, an important factor given the limited resources of the group.

Simulation performance and debugging were enhanced dramatically by organizing the model with compile-time switches that allowed mixed levels of abstract, behavioral, or fully structural chip implementations. For a small design team, this allowed rapid turnaround of a block of the design, even if fully implemented in gate-level detail, by choosing to use higher levels of modeling for the rest of the chip.

The C-model-based simulation environment had positive project impact that saved enormous costs and simulation time. The model was not tied to a proprietary "fee per licensed CAD simulator," such as popular industry Verilog® or VHDL simulators. For this reason, simulations could be run on as many systems as were available on the network, without affecting the CAD tool budget.

#### • Technology selection


Technology selection was influenced by a number of major factors, including packaging, off-chip drivers, on-chip RAM, gate speed, compatibility of design methodology, and on-chip interconnections, as discussed in the following sections.

After examining the merchant ASIC supplier market for a company that could meet our requirements, we selected IBM and its CMOS 5L process with a 360-pin enhanced ceramic ball grid array (CBGA) chip. A quick snapshot of some details of the alternatives is shown in **Table 4**.

A die photo is shown in **Figure 6**.


#### Packaging

The 340-pin PQFP was the only other option capable of actually meeting our cost goals; however, its relatively poor electrical characteristics greatly reduced the actual number of available I/Os if supplier power and ground distribution rules for signals were met. Because we had chosen a direct drive scheme for large frame buffer loading, ASIC power dissipation would increase compared to ZLX. This made the metal quad flat package (MQFP),



**Figure 5**

Verification environment. The software model and the hardware are virtually interchangeable, allowing the software developers to visually verify the functionality by utilizing real applications in parallel with hardware development.



**Figure 6**

Die photograph of the PowerStorm ASIC showing large areas of on-chip RAM.

with its poorer thermal performance, a somewhat distant runner-up. The IBM 360-pin CBGA solution could have been considered slightly more expensive than our alternatives, but we were able to offset all of this through module-level component cost reductions.

**Table 4** Technology characteristics of the alternatives considered for the PowerStorm design.

|                                             | <i>IBM CMOS 5L<br/>with 360 CBGA</i> | <i>Supplier X with<br/>304 MQFP</i> | <i>Supplier Y with<br/>340 TBGA</i> |
|---------------------------------------------|--------------------------------------|-------------------------------------|-------------------------------------|
| Signal I/Os                                 | 287                                  | 292 (243–197<br>effective)          | 328 (272–224<br>effective)          |
| Minimum die size for I/Os                   | 6.3                                  | 11.7                                | 9                                   |
| CMOS process drawn/ $L_{eff}$               | 0.5/0.45                             | 0.7/0.5                             | 0.7/0.5                             |
| Layers of metal                             | 3, 4, or 5                           | 2 or 3                              | 2 or 3                              |
| Usable gates                                | 150K (4LM)                           | 150K                                | 150K                                |
| 5-V-tolerant I/Os                           | Yes                                  | Yes                                 | Yes                                 |
| PCI I/Os                                    | Yes                                  | Yes                                 | Yes                                 |
| Self-correlated I/Os                        | Yes                                  | No                                  | No                                  |
| Max. power ( $t_j = 85^\circ$ max., 50 lfm) | 5 W                                  | 2.5 W                               | 2.5 W                               |

**Table 5** Comparison of Smart Frame Buffer chip implementations.

|                               | <i>HX</i> | <i>ZLX-E</i> | <i>ZLXp-E</i> | <i>PowerStorm</i> |
|-------------------------------|-----------|--------------|---------------|-------------------|
| Package                       | 184 PQFP  | 240 PQFP     | 240 PQFP      | 360 CBGA          |
| CMOS process drawn/ $L_{eff}$ | 1.2/1.0   | 1.0/0.8      | 1.0/0.8       | 0.5/0.45          |
| Die size (mm/side)            | 8         | 11.7         | 11.7          | 6.3               |
| Layers of metal               | 2         | 3            | 3             | 4                 |
| Gates used                    | 21K       | 65K          | 70K           | 150K              |
| Typical power (W)             | 0.6       | 1.7          | 1.9           | 2.5               |
| Clock rates (MHz)             |           |              |               |                   |
| Graphics core                 | 25        | 38           | 38            | 75                |
| Memory                        | 25        | 38           | 38            | 80                |
| I/O bus                       | 25        | 25           | 33            | 33                |

#### Off-chip drivers

The two most timing-critical interfaces to the ASIC were the PCI bus interface and frame buffer memory interface.

**Bus interface** For our system platforms, simple 5-V tolerance is not enough; meeting the required  $I$ - $V$  characteristics often necessitates using a split power rail for the PCI drivers. Using a technology compliant to both signaling environments was an attractive alternative. CMOS 5L was able to deliver 5-V-tolerant I/O drivers including PCI 5-V and 3.3-V-compliant drivers and no split power rail.

**Memory interface** With the chosen direct-drive scheme for the frame buffer, memory controller physical interface issues would typically have limited the actual frame buffer bandwidth performance to something less than the theoretical cycle time limits imposed by a supplier's VRAM specifications. Typical problems that reduce achievable cycle time are I/O-to-I/O delay mismatches between signals on the bus caused by on-chip ASIC process variation; drivers incapable of delivering clean, fast-edge-rate signals to the wide variety of loading situations found in the different module-level implementations (e.g., 2MB frame buffer versus 16MB);

and noise from the large number of switching I/Os being coupled back into the core logic. CMOS 5L was able to deliver

- Well-matched pin-to-pin delays by using on-chip I/O delay compensation circuitry.
- A variety of impedance-matched drivers capable of driving heavy loads without requiring the use of additional power and ground pins.
- A significant reduction in noise by virtue of the high-quality electrical path provided by the package as well as the very generous number of chip-level power and ground connections.

General I/O signal quality was found to be better than anything we have experienced. The low package electrical resistance, capacitance, and inductance of the C4-bonded, enhanced CBGA with its high-frequency decoupling capacitors resulted in outstanding signal quality (see **Figure 7**).

#### On-chip RAM

Our design decisions resulted in the need to more than double the amount of on-chip RAM. In addition to this, most of the RAM had to be fast multiport on-chip RAM for deep buffering through pipelines and across

asynchronous boundaries. CMOS 5L was able to deliver on these requirements. This was accomplished by using compiled memory arrays with rapid availability. This would normally require custom diffused memory with longer lead times.

#### Gate speeds

The published gate speeds for CMOS 5L were among the best available in this time frame for this class of technologies. It was important to have a technology capable of running memory-controller state machines at up to 100 MHz. This would allow the use of the most aggressive memory speed bins projected to be available in our product time frame. With the choice of CMOS 5L, we felt we would not be making any compromises.


#### Compatible design methodology

While it was possible to achieve many of our goals with other supplier technologies, IBM was the only supplier that could achieve them all with minimum risk. Because IBM was fundamentally a new supplier, a new ASIC sign-off process had to be developed. Early on, many test cases were performed to evaluate the CAD process and ensure that our design met the IBM design-for-test level-sensitive scan design (LSSD) manufacturing test requirements; the LSSD implementation paid long-term dividends in that we did not have to generate our own test patterns during the ASIC sign-off time frame. In addition to this, IBM EinsTimer™ fully static timing methodology mirrored practice that was already common at Digital. It also provided unique time-saving features such as chip-level power optimization to improve critical path timing. These features were particularly important given our time-to-market requirements and an estimated savings of at least one person-month.

#### On-chip interconnections

Cost-sensitive designs are always gate-limited, but to facilitate further hardware and software debugging, the first pass of PowerStorm was done in a 7.2-mm image. For the final pass, logic design optimizations and more efficient clock planning, accomplished with the involvement of the IBM ASIC design center, allowed the design to fit into a 6.3-mm die image. The four-layer-metal process and the IBM layout tools were easily able to route some of the historically difficult physical structures in this design, such as large crossbar switches.

Smaller average on-chip wire delays and the CBGA improved thermal characteristics compared to PQFP performance led to the ability to exceed many of our performance/clock rate goals. The results are shown in **Table 5**, providing comparison with the previous-generation Smart Frame Buffer product.



**Figure 7**

Oscilloscope photograph showing the superb frame buffer signal quality attained with the PowerStorm ASIC.

#### Summary

The Smart Frame Buffer architecture in PowerStorm is not limited by bandwidth or the graphics engine. Even with a 300-MHz Alpha CPU, PowerStorm is often limited by the CPU that delivers the drawing primitives. For this reason, PowerStorm will probably be the last generation of the SFB architecture; future designs will require a new architecture with new, more complex primitives and a cost model to support the silicon necessary to implement them.

The design and verification strategy used for PowerStorm enabled us to achieve our product cost and time-to-market goals. Using a technology that clearly represents a paradigm shift in how cost-sensitive designers may realistically solve problems with more I/O, a superior CMOS process, and ASIC sign-off verification tools that were consistent with our internal process, we were able to obtain performance levels which by previous conventional wisdom would require more esoteric memory architectures. These technology features enabled the design team to

achieve its time-to-market goals with limited manpower, while maintaining strict containment on cost despite added functional features.

The greatest risks were related to IBM being new to externalizing the technology offering. Problems did occur in this area throughout the course of the project, but with the resolve of both parties, all barriers were ultimately removed.

### Acknowledgments

The PowerStorm graphics processor is the culmination of the efforts of a small team of hardware developers and a large team of software developers, without whose constant creation of numerous verification scripts this product would not be possible. The authors wish to acknowledge Julie Druckerman for her helpful suggestions in preparing this manuscript, and our ASIC support team at IBM Burlington, especially Shawna Moquin and Bob Savaglio.

PowerStorm, TURBOchannel, and AccuVideo are trademarks of Digital Equipment Corporation.

X Windows is a registered trademark, and X11 is a trademark, of the Massachusetts Institute of Technology.

Windows-NT is a trademark of Microsoft Corporation.

OpenGL is a trademark of Silicon Graphics, Inc.

Rambus is a trademark of Rambus, Inc.

Verilog is a registered trademark of Cadence Design Systems, Inc.

EinsTimer is a trademark of International Business Machines Corporation.

### References

1. J. McCormack and R. McNamara, "A Smart Frame Buffer," *Research Report 93/1*, Digital Equipment Corporation Western Research Laboratory, Palo Alto, CA, 1993.
2. J. McCormack, R. McNamara, and C. Gianos, "The Smart Frame Buffer Goes Hollywood: 3D and TV," *Proceedings of the Hotchips VI Symposium*, 1994, pp. 143-152.
3. Kenneth W. Correll and Robert A. Ulichney, "The J300 Family of Video and Audio Adapters: Architecture and Hardware Design," *Digital Tech. J.* 7, No. 4, 20-23 (1996).

*Received December 11, 1995; accepted for publication April 8, 1996*

**Chris Gianos** *Digital Equipment Corporation, 129 Parker Street, Maynard, Massachusetts 01754 (gianos@eng.pko.dec.com).* Mr. Gianos is a consulting engineer at the Digital Workstation Graphics and Multimedia Development group in Maynard, Massachusetts. He joined Digital in 1985, and has since worked on the development of numerous CAD tools, graphics modules, and ASICs. Mr. Gianos received his B.S.E.E. from Cornell University in 1985.

**David Hobson** *Digital Equipment Corporation, 129 Parker Street, Maynard, Massachusetts 01754 (hobson@eng.pko.dec.com).* Mr. Hobson is a principal engineer at the Digital Workstation Graphics and Multimedia Development group in Maynard, Massachusetts. He joined Digital in 1987, and has since worked on the development of CPU, graphics modules, and ASICs. Prior to that, he worked for IBM in East Fishkill, NY, from 1982 to 1987, writing automatic test software and performing physical failure analysis on DRAMs and SRAMs. Mr. Hobson received his B.S.E.E. from Worcester Polytechnic Institute in 1982. He is a member of the Institute of Electrical and Electronics Engineers.