



TECHNOLOGY report




**THE
ARCHITECTURAL
EVOLUTION
OF A HIGH-
PERFORMANCE
GRAPHICS
TERMINAL**

THE 4115B: THE ARCHITECTURAL EVOLUTION OF A HIGH-PERFORMANCE GRAPHICS TERMINAL



Douglas J. Doornink is a senior electronic engineer in Graphic Systems Products (GSP). He was the leader of the hardware development project for the 4115B. After joining Tek in 1974, he participated in the design of the 4112, 4027, 4025, 4024, and 4081. Doug holds an MSEE from Stanford University. His BSEE is from the University of Washington.



John C. Dalrymple is a senior electronic engineer in Graphic Systems Products. He lead the team that developed the picture processor for the 4115B. John joined Tek in 1976 and has done research in color anti-aliasing, display-processor architecture for directed-beam refreshed displays, and display-processor architecture for intelligent oscilloscopes. John's MSEE and BSEE are from Oregon State University.

This article is a case study of the design of the Tektronix 4115B. It details the design constraints, shows how the architecture evolved from the Tektronix 4113, and shows how performance was gained by adding a microcoded picture processor and some special-purpose hardware. This article was adapted from a paper given at Spring COMPCON, 1984.

The task given the designer is usually not to design the "best" or ultimate product but rather to squeeze the most out of a design – given many constraints. These constraints range all the way from how much the customer is willing to pay to how big the product can or must be. In 1981, the GSP engineering group was asked to design a "high-performance" graphics terminal.

High performance in computer graphics requires interactivity, and interactivity hinges on the speed at which an image can be produced or changed on the screen. The system must respond to the user in the appropriate amount of time, or productivity will suffer and frustration will increase. The challenge for the designers of a high-performance graphics terminal is to squeeze out the most speed given a set of real-world constraints.

This article reviews the design of the Tektronix 4115B and, in particular, the picture processor. First, we will discuss the background and motivation behind the design; this discussion will

analyze the Tektronix 4113 and its performance bottlenecks. We will then discuss the system partitioning of the 4115B and the tradeoffs made, show how the right balance of hardware and microcode can substantially increase performance, and look at how the completed design performs.

Background and Motivation

Even though most engineers would like to start each new design from the ground up, this is not the job usually given. Because most products are not designed in a vacuum, product design means specific requirements and hard constraints. Products are designed within an environment that includes other products, customer desires, and technology limitations. This was very much the case with the design of the Tektronix 4115B graphics terminal.

Functional requirements

Since the 4115B was to be a member of the 4110 family, the most important functional requirement was family compatibility. The 4110 family was modeled after the *SIGGRAPH CORE*,¹ which includes the ability to store and transform picture segments. Compatibility with all the 4110 family commands was necessary. There were raster display commands implemented in the 4112 and 4113. These commands were similar to those in the *Raster Extension to the Core System*.² Compatibility with all raster and color features of the Tektronix 4113 was necessary because it was the only color-raster member of the 4110 family.

Because the 12-bit coordinate system in the then current 4110 family was inadequate for many applications being addressed by computer graphics, a 32-bit coordinate system requirement was added to the design goals. This extension had to be compatible with the old 12-bit system, of course.

In addition, the 4115B was to be compatible with ANSI X3.64 alphanumeric control commands. This would allow the terminal to be used with the many alphanumeric-oriented programs available. This compatibility was essential. Even though the 4115B is a high-performance graphics terminal, it is frequently used for alphanumeric data entry and editing. Most users expect their \$20,000 terminal to be able to do what their \$500 terminal can do.

Performance requirements

Interactivity is the primary requirement of a high-performance graphics terminal. Another word for interactivity is *speed*. To be interactive in a man-machine environment, the long-standing rule of thumb was "response within two seconds."³ This rule has been superseded because exposure to personal computers and single-user systems has increased user expectations.⁴ Functions done locally on a terminal are expected to be done quickly, if not instantaneously.

On the 4110 terminals, this meant that the re-draw and manipulation of picture segments, which are stored locally in the terminal, must be done quickly. Dragging and transforming of segments should happen instantly. Picture segments of several thousand vectors are not unusual; pictures with tens of thousands of vectors are common.

Another performance requirement was high display quality. 1280 × 1024 pixels is now mandatory for a high-resolution display. For picture stability, 60-Hz noninterlaced scanning is required; anything less tires users with flicker.

A third requirement was that the new terminal must support a display of 256 colors simultaneously. This meant at least eight bits per pixel, or eight planes of display memory.

Physical constraints

The 4115B was to be an addition to a family; this predetermined physical constraints. The 4110 family was pedestal based; therefore the 4115B had to fit into the pedestal. (See figure 1(A).) The established card cage had room for only 18 circuit cards; all of the display system plus all options had to fit. We knew six circuit cards would be needed for the display memory and controller and that nine would be needed for the options common to the 4110 family. We could use only three 8.5 × 11-inch boards for the picture processor (figure 1(B)). This definitely restricted the amount of circuitry that could be used for the picture processor.

The power for the system as a whole was fixed at 500 watts. This limit was important when it came to making tradeoffs between hardware, software, and microcode in the system architecture.

Technological constraints

Because the development schedule for the 4115B was to be short, we had to use off-the-shelf parts – no custom VLSI. We also had to use as many existing 4110 circuit cards as we could. This meant staying with the processor (8086), option cards, and bus structure used in the previous 4110 terminals. This limited system architecture to 1-M byte address space and 16-bit data paths.

Bottlenecks in the 4113

The 4113 was the functional model for the 4115B, but the 4113 had some shortcomings for high-performance graphics; the most important was its picture-segment re-draw speed. The 4113 outputs constant-time vectors to the screen at about 1000 fully transformed vectors per second.




Figure 1. (a) Front view of the 4115B pedestal. The card-cage dimensions, and (b) circuit-card dimensions were the primary physical constraint for the new terminal.

Also, the 4113 couldn't keep up with sustained communications at 19.2 kilobaud. The 4113 could sustain only 9600 baud without overflowing the communications buffer.

8086 overloaded

In the 4113, there is a definite imbalance between hardware and firmware.

Although there was a hardware vector generator in the 4113, the firmware in a 8086 microprocessor did most of the work. Running in the 8086 was a message-based multitasking operating system, which also handled all host communications and peripheral management. In the graphic pipeline, 8086 firmware did all command parsing and decoding. It also did the display-list creation and traversal, 2D transformations, panel-scan conversion, and dot-matrix character setup.

The hardware vector generator was a slave to the 8086, doing only actual vector generation and controlling writing to the frame buffer. Also, the hardware performed block moves in the frame buffer for dialog-area scrolling since the dialog area in the 4113 was in the frame buffer. The hardware vector generator could do 2600 full-screen vectors per second and 100,000 short vectors per second; so it was not the bottle neck.

Data flow in the 4113

Figure 2 shows the data flow in the basic 4113. The data originates from one of three sources: the RS-232 interface, the keyboard, or the graphic input devices. After going through the command interpreter, the communication system, or the graphic

input system, all data – except report data – goes to the display driver. The display driver is shown in detail in figure 3. The heart of this driver is the transform system: Ultimately, all of the data paths lead to the transform system, and the hardware is not invoked until the very end of the data flow.

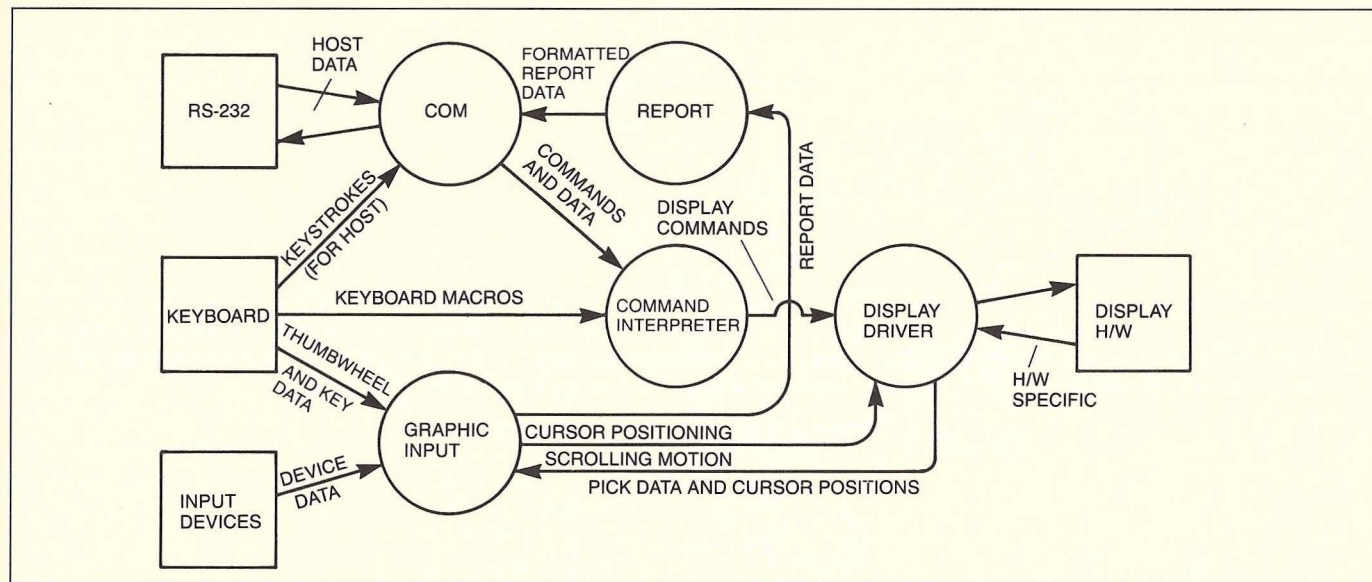


Figure 2. 4113 basic data flow. Data originates from three sources: the RS-232 interface, the keyboard, graphic input devices. The display driver is shown in detail in figure 3.

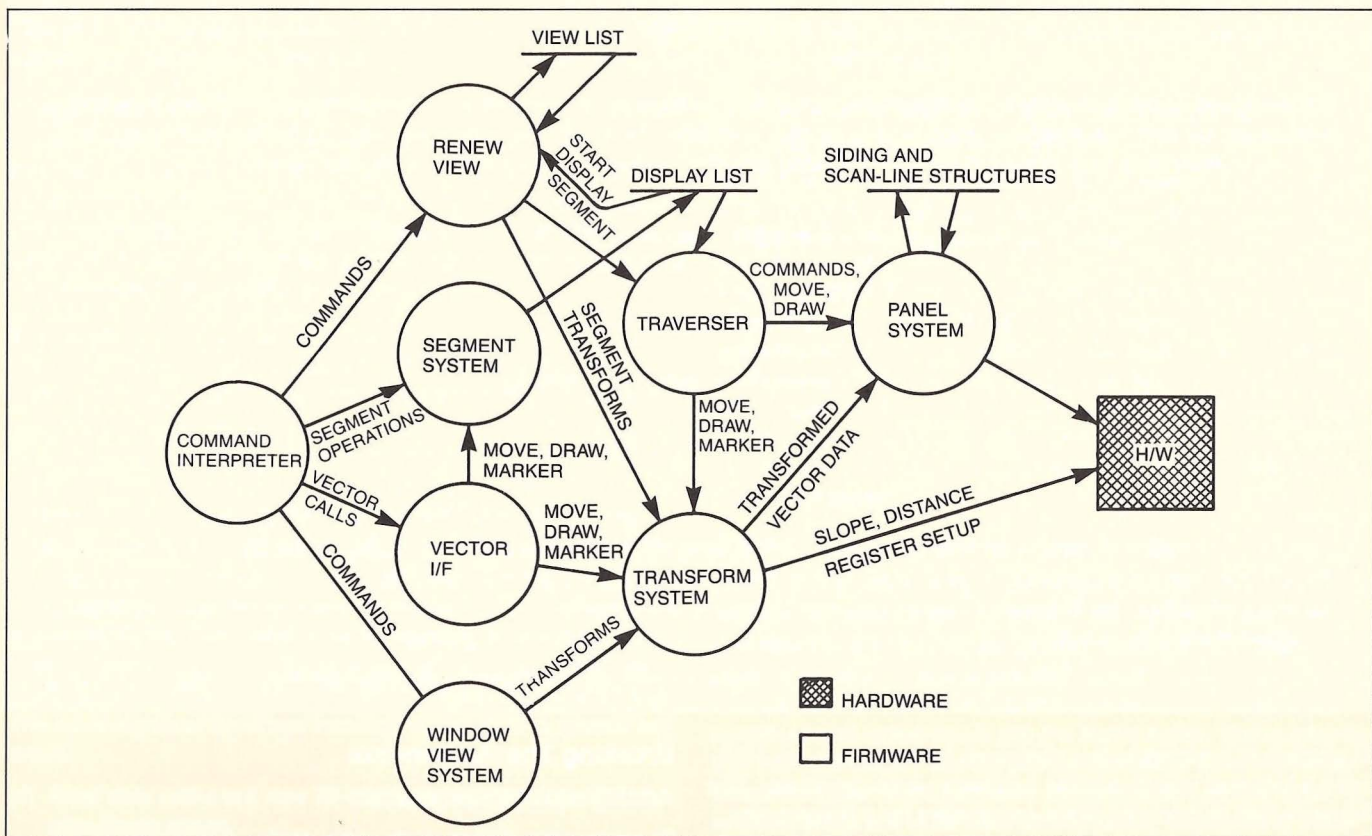



Figure 3. 4113 display-driver data flow. Ultimately, all paths lead to the heart of the system, the transform system. The hardware is a slave to the 8086, invoked only at end of data flow. There was a software/hardware unbalance.

BRESENHAM'S LINE-DRAWING ALGORITHM

Bresenham's algorithm is designed so that each iteration changes one of the coordinate values by ± 1 . The other coordinate may or may not change, depending on the value of an error term maintained by the algorithm. This error term records the distance, measured perpendicular to the axis of greatest movement, between the exact path of the line and the actual dots generated. In the example below, where the x axis is the axis of greatest movement, the error term e is shown measured parallel to the y axis. The following description of the algorithm assumes this particular orientation of the line.



At each iteration of the algorithm the slope of the line, $\Delta y/\Delta x$, is added to the error term e . Before this is done, the sign of e is used to determine whether to increment the y coordinate of the current point. A positive e value indicates that the path of the line lies above the current point; therefore the y coordinate is incremented, and 1 is subtracted from e . If e is negative the y coordinate value is left unchanged. Thus the basic algorithm is expressed by the following PASCAL program:

```
{Note: e is real; x, y, deltay are integers}
e := (deltay/deltax) - 0.5;
for i := 1 to deltax do begin
  Plot(x,y);
  if e > 0 then begin
    y := y + 1;
    e := e - 1
  end;
  x := x + 1;
  e := e + (deltay/deltax)
end;
```

The weakness of this sequence of operations lies in the division required to compute the initial value and increment of e . This division can be avoided, however, since the algorithm is unaffected by multiplying e by a constant: only the sign of e is tested. Thus by multiplying e by $2\Delta x$ we produce the following program, requiring neither divisions nor multiplications:

```
{Note: all variables are integers}
e := 2 * deltay - deltax;
for i := 1 to deltax do begin
  Plot(x,y);
  if e > 0 then begin
    y := y + 1;
    e := e + (2 * deltay - 2 * deltax)
  end
  else e := e + 2 * deltay;
  x := x + 1
end;
```

A full implementation of Bresenham's algorithm involves allowing for other cases besides $0 \leq \Delta y \leq \Delta x$, the case discussed above. At the same time the algorithm can be somewhat simplified by using only integer arithmetic. Bresenham's algorithm avoids generating duplicate points. Because it also avoids multiplications and divisions, it is well suited to implementation in hardware or on simple microprocessors.

— From *Principles of Interactive Graphics*, Newman and Sproull

In addition to performing few processes by hardware, the 4113 used a low-level interface to the hardware vector generator. The firmware had to calculate all the parameters for a Bresenham vector algorithm⁵ and load the parameters into registers on the vector generator. The Bresenham algorithm was then executed in the hardware.

Breaking The Bottlenecks

Adding a picture processor

To achieve the design goals of the 4115, we had to open the bottlenecks of the severely overloaded 8086 and mostly idle vector-generator hardware. We took a conventional approach:

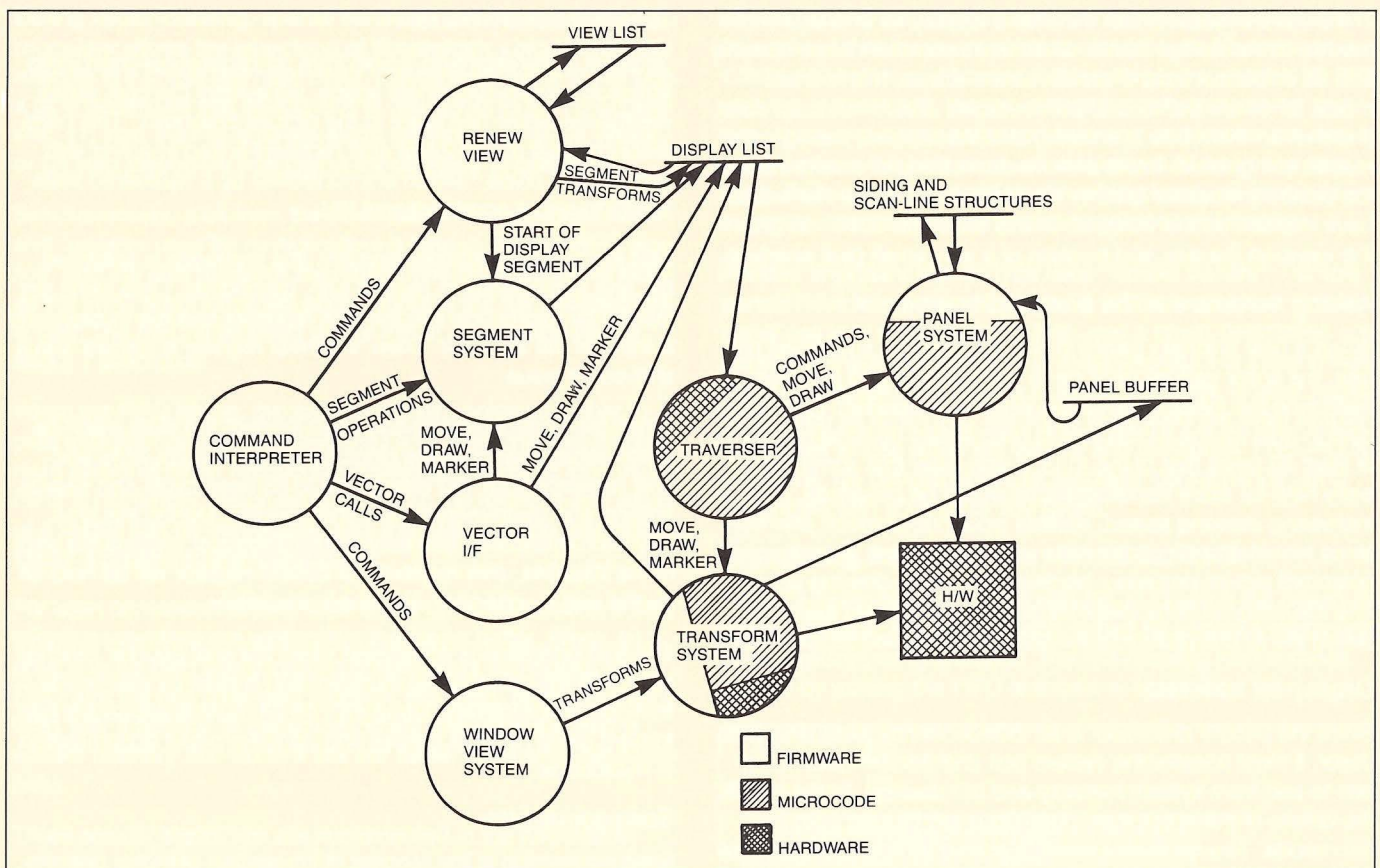


Figure 4. 4115B display-driver data flow. We offloaded lower-level graphics processing from the 8086 by replacing the vector-generator hardware with a microprogrammable bit-slice "picture processor."

replacing the vector-generator hardware with a microprogrammable bit-slice "picture processor," so that the lower-level graphics-processing functions could be offloaded from the 8086. This new division of labor can be seen in figure 4. Along with the bit slice, this processor contains several hardware accelerators for speed-critical tasks. It was constrained to fit on two standard-sized (11.25 × 8.5 in) cards.

The 4115B picture processor is an instruction-set processor that executes programs (display lists) built by code running on the 8086. The initial specification of the display-list format was done by software engineers who were designing the 8086 code and by microcoders who would be implementing the instruction set. The specification evolved as the implementation proceeded; however, task partitioning between the 8086 and the picture processor did not change drastically after the first specification.

8086 tasks

In the 4115B, the 8086 retains the multitasking operating system, host communication, peripheral management, and display-list management functions discussed in the description of the 4113. The code significantly differs from that of the 4113 in several key areas:

- (1) Most data paths are now 32-bits wide to support the 32-bit coordinate space.
- (2) New algorithms and data structures allow faster, more space-efficient creation of many small graphic segments.

- (3) New code drives the hardware dialog overlay and cursor overlay (not present in the 4113).
- (4) The numeric co-processor (8087) is used for the precise arithmetic operations needed to generate graphic-image transforms for the picture processor.

Additionally, the 8086 assists the picture processor in the scan conversion of panels (filled areas). For this task, the picture processor transforms coordinates from the display list and sends them to the 8086. The 8086 then builds temporary data structures, which it passes back to the picture processor. The picture processor uses these data structures to compute and fill the set of pixels interior to the specified area. Finally, the 8086 deletes the temporary structures.

Picture processor tasks

The picture processor executes commands from a display list resident in system memory. It transforms graphic primitives, described in a 32-bit integer coordinate space, into 1280 × 1024 pixel screen-coordinate space and clips the results to rectangular viewports on the screen. It scan converts the transformed primitives and writes pixels into the frame buffer. Using information from the display list, the picture processor controls the appearance parameters (primitive attributes such as line style, whether areas are to be filled or left hollow, background transparency of dot-matrix characters, etc.).

The picture processor can traverse a display list in one of three modes. In the *default* mode, all attribute commands are obeyed and all visible portions of primitives are drawn. In the *erase* mode, all visible portions of primitives are drawn in a solid color (fixed when the mode is entered) and the attribute-setting commands are ignored. In the *pick* mode, nothing is drawn; instead, the picture processor informs the 8086 about items that would have intersected the viewport. In this case, the 8086 has given the viewport the size and position of a very small pick aperture.

Adding hardware to the picture processor

When a user is locally panning and zooming on a retained picture, the entire picture must be re-transformed and re-drawn each time the "view" key on the terminal is pressed. A principal design goal for the 4115B was to both speed up local redraw to at least 20 times that of the 4113. It was clear that some hardware help for point transformations would be beneficial. Also, in increasing drawing speed, vectors would need much attention since terminal applications were vector intensive. Therefore, we optimized picture-processor hardware for vector-drawing speed. The following sections describe how we partitioned tasks between the 4115B picture-processor hardware and its microcode.

Display-list traversal

Since the picture processor is an independently executing processor, it must acquire the system bus and perform data transfers to and from system memory and I/O devices. In the 4115B, the details of these low-level operations are hidden from the microcode. Two hardware state machines (both resident in one registered PAL) implement the bus acquisition and data-transfer protocols.

The machines are activated by a single microinstruction. The microcode can then continue executing until it needs the results of a bus-read or until it tries to start another bus operation. At that time, a hardware "wait" mechanism temporarily halts the picture processor until the original cycle has been completed. Thus, microcode does not have to test any status flags to see if a transfer has been completed prior to starting another transfer.

The terminal bus has 20 address bits (referencing 1 Mbyte) and 16 data bits. Maximum bus bandwidth is obtained by performing 16-bit transfers on even-address boundaries. The picture-processor instruction set (display list) consists of one- and two-byte opcodes, with operand lengths ranging from zero to tens of bytes.

Most operands are immediate data following an opcode. Therefore, a display-list fetch-ahead system was built in microcode. Routines needing operands from the display list make subroutine calls to this system, which manages a three-byte fetch-ahead queue and always does 16-bit bus transfers at even addresses.

If the execution of an opcode references data from outside the instruction stream, another set of subroutines is used, which destroys information in the queue. These may be eight- or 16-bit transfers at arbitrary addresses. The queue state must be restored (by calling a subroutine) before the next opcode is fetched. Transfer-of-control opcodes flush and refill the queue.

2D point transformation

To assist in transforming points from 32-bit terminal-coordinate space to screen-coordinate space, we added low-cost serial/parallel multiplier hardware to the picture processor. This hardware consists of two 24-bit shift registers linked in a ring with two 25LS14 chips,⁶ which are controlled by a hardwired state machine and a combinational PAL.

Like the bus-data transfer circuitry, this hardware is activated by, and can operate in parallel with, microcode. Data items are loaded by microcode into the shift registers and the multiplicand is input to the 25LS14s. Then, a control register is loaded with a shift-count value.

When the control register has been loaded, the state machine automatically switches the clock period for the 25LS14s and shift registers from 163 nsec/cycle to 65 nsec/cycle and begins shifting and decrementing the shift counter. During this time, any microinstruction that attempts to access or change the data in the 25LS14s, shift registers, or control register will be held off by the wait mechanism described previously.

The shift operation is completed when the shift counter reaches zero, and the result of the multiplication is read by microcode from the shift registers. Up to 48 × 48-bit multiplications are performed using multiple passes and partial-product accumulation in microcode.

Scan conversion

Scan conversion is the process of converting primitive descriptions (such as vector endpoints or polygon vertices) together with attribute information (such as line color or area-fill pattern) into the set of pixel addresses in the frame buffer to be modified and into the pixel data to be written at those addresses.

The 4115B includes special frame-buffer-interface (FBI) hardware, resident on one standard-sized card. It is designed to hide the details of frame-buffer memory organization from the picture processor. The frame buffer appears as a 2D array of 8-bit pixels. Also included in the FBI hardware are X- and Y-address registers (both registers can be incremented, decremented, held, or loaded on each cycle) and two sets of pixel-data registers. Either set of registers can be selected on each microcycle. A four-pixel cache with automatic swapping hardware contains a copy of the current frame-buffer region being accessed by the picture processor.

The FBI is augmented by additional hardware on the picture processor for boosting vector performance. The control signals that drive the FBI can come either directly from the current microinstruction (during vector setup) or from a 32-deep FIFO memory that queues up address-stepping and data-register-selection commands for the FBI. These commands are used during the actual drawing of vectors. The address-stepping commands trace out the trajectory of the vector, and the data-register-selection commands select between foreground and background colors for dashed lines.

During vector drawing, special hardware around the bit-slice processor allows the inner loop of Bresenham's algorithm^[5] to execute in a single 163-nsec cycle. During each cycle of this loop, a bit is generated that selects one of two inputs to the FIFO pipeline: (1) step the FBI address along the long axis of the vector, or (2) step diagonally one unit in the direction of both axes.

Simultaneously, a bit from the dash pattern (stored in the same shift registers used for multiplication) is loaded into the FIFO pipeline, and the shift registers are rotated one bit position. As long as the FIFO is not full, it will accept input at the full 163-nsec/pixel rate. The FIFO is emptied by the FBI at an average rate of about 1 μ sec/pixel.

When the FIFO fills up, the wait mechanism holds off further inputs until the FBI has unloaded the FIFO. Also, the wait mechanism prevents direct microcode access to the FBI as long as the FIFO is not empty. Because of the FIFO, the fetching and setup of a vector can be overlapped with the writing of the previous vector's pixels into the frame buffer.

For scan conversion of solid-filled areas, up to 80 pixels at a time (along a scan line) are written into the frame buffer. When an area is to be filled with a pattern consisting of two colors, the FIFO and the two pixel-data registers are used. When an area is filled with a general fill pattern (arbitrary size up to the limit of system memory and up to 8 bits per pixel), there is no hardware help.

Three area-fill algorithms are implemented in the microcode:

- (1) A speed-optimized rectangle-fill algorithm with trivial rejection and clipping for rectangles whose sides are vertical and horizontal.
- (2) A microcode-only algorithm for unclipped polygons with up to 16 sides.
- (3) A general panel algorithm for areas that may be clipped, may have holes in them, and are unrestricted as to the number of edges.

In the general panel algorithm, the picture processor passes transformed vertices to the 8086 (through a shared buffer in system memory) and the 8086 builds data structures for the microcode to use when it fills the panel. Because of the handshaking overhead the general panel algorithm has the slowest panel-filling performance.

System Performance

It is not enough to say "the system must be fast" and use that statement as a design goal. To know if the design is good enough, one must have numbers against which design-goal achievement can be measured. Early in the design of the 4115B, we established performance metrics that we felt were appropriate in light of the application targets for the 4115B. Our metrics dealt with drawing speed: vectors per second, segments per second, and simple panels per second. All of these metrics assumed the application of a 2D transform as part of the drawing process.

After analyzing sample pictures from typical applications, we determined that the average vector length was 10 pixels on a 1280 x 1024 display. Pictures in this category averaged 30,000 vectors; pictures with longer vectors typically have fewer vectors. The worst-case application had only one vector per segment, while in the best case, all vectors would be in one segment. Both cases needed benchmarking because each picture segment incurs significant picture-processor overhead.

Another important picture type uses "simple panels" – simple panels have fewer than 16 edges. A prime example of simple-panel use is in a solids model that uses a mesh description and generates the image with a lot of quadrilaterals. For these images, the quadrilaterals have an area of about 100 pixels. In other applications, such as VLSI CAD, the areas to be filled are even simpler – rectangles whose sides are vertical and horizontal. For CAD applications, we benchmarked rectangle-fill performance using rectangles with an area of 100 pixels.

Comparing Performance: 4115B vs. 4113

Figure 5 compares the 4115B with the 4113 using the previously defined metrics. The performance gained by adding a microcoded picture processor is dramatic. The line-drawing performance gained by adding the FIFO is shown in figure 6.




Figure 5. 4115B versus 4113 performance. Scale is logarithmic. Performance metrics assume that graphic primitives are repainted from retained segments.




Figure 6. 4115B vector drawing performance. The effect of having a FIFO between the picture processor and the frame buffer can be clearly seen. The initial slope is 163 nsec per pixel (microcode instruction time). The final slope is 1.22 msec, the worst-case rate for writing pixels to the frame buffer.

From figure 6, one can see that the basic setup time for a vector, including the transform time, is about 18 μ sec. The initial slope is 163 nsec per pixel, the microcode instruction time. As the FIFO becomes more than half full, the output port has priority and we see the knee in the curve. The final slope of the graph is 1.22 μ sec per pixel, which is the worst-case rate at which the pixels can be written into the frame buffer.

It is also interesting to compare the performance of the 4115B with that of other architectures and implementations. Four architectures are compared in figure 7. The 4113 is at one end of the spectrum, an example in which a single microprocessor does most of the work. We included the Apollo DN420 as an example

of a graphic system having a microprocessor and hardware optimized for BITBLT-type functions. The third example is the 4115B, in which the work is divided between a microprocessor and a microcoded picture processor. The final example is the Seillac-7 in which the work is done by a pipeline of a 32-bit bit-slice processor, two 16-bit bit-slice processors, two 16-bit microprocessors, a 4×4 matrix multiplier, a clip circuit, and a perspective circuit.

Remaining Issues

The 4115B project did not address several functional extensions and performance issues: we did not implement 3D transforms or picture-segment hierarchy, although these are natural extensions to the 4115B feature set. The main reason these features are not in the 4115B now is that our resources were limited. When we prioritized features, these two fell below the cut line.

Another functional issue is standardization, GKS^[7] in particular. Tektronix supports GKS as a standard, but since the 4110 family was not an implementation of GKS, the 4115B could not implement GKS and still be a part of the 4110 family.

Some performance areas were not addressed by the 4115B project. Although one can do picture dynamics on the 4115B, it was never designed to do animation; the performance necessary for animation was beyond the scope of the 4115B.

We did not attempt more than eight planes, which would allow the display of more than 256 colors at a time. Restricting the 4115B to eight planes was necessary because of limits in the frame-buffer system – not because of picture processor capacity.

Finally, we did not optimize the system for BITBLT-type operations as has been done in some work-station architectures. The Apollo DN420 engineering work station is a good example of this architecture. Our application and system architecture required optimization toward line drawing.




Figure 7. Vector-drawing performance of different architectures.

Conclusion

The architecture development of the 4115B was an evolutionary step rather than a revolutionary jump. The main reason for this was the set of constraints placed on the project and architecture by the product-family environment in which it was developed. In spite of this, dramatic performance improvements were made, and the requirements of interactivity were met by adding a micro-coded picture processor and partitioning the work among firmware, microcode, and hardware. Because of this balancing, brute force was not necessary for high performance.

References

- [1] "Status Report of the Graphics Standards Planning Committee of ACM/SIGGRAPH," published as *Computer Graphics*, vol. 11, no. 3 (Fall, 1977).
- [2] "Status Report of the Graphics Standards Planning Committee of ACM/SIGGRAPH," published as *Computer Graphics*, vol. 13, no. 3 (August 1979).
- [3] J. Martin, *Design of Man-Computer Dialogues*, Prentice-Hall: Englewood Cliffs, N.J. (1973).
- [4] J.D. Foley and A. Van Dam, *Fundamentals of Computer Graphics*, Addison-Wesley: Reading, Mass. (1982).
- [5] J.E. Bresenham, "Algorithm for Computer Control of Digital Plotter," *IBM System Journal*, vol. 4, no. 1 (1965).
- [6] "Bipolar Microprocessor Logic and Interface," Advanced Micro Devices, Sunnyvale, CA (1983).
- [7] Graphical Kernel Standard (GKS) International Standards Organization (ISO) Draft International Standard (DIS) 7942.

TECHNICAL STANDARDS

Technical Books Available

The Technical Standards library of technical volumes is growing. Here are some recent additions:

Mark's Standard Handbook for Mechanical Engineers, a dictionary.

IEEE Standard Dictionary of Electrical and Electronic Terms.

IEEE STD 141-Recommended Practice for Electric Power Distribution for Industrial Plants.

If your interest is in a field affected by standards, we may have the book you need. You don't have to buy; we'll lend it.

Books and Reports

IOOC '81 — *Third International Conference on Integrated Optics & Optical Fiber Communication* — April 27-29, 1981. This is a technical digest. A number of test procedure documents from this document are also available. May be purchased or borrowed.

NTIS — *Directory of Computer Software Applications on Energy* — Available for loan.

NASA — *A Catalog of Selected Computer Programs* — Title, number and a brief description of the program included. Available for loan.

New Standards

EIA-RS-505 — *Packaging for Return CRT Glass Component*

EIA-RS-455-47-1983, FOTP-47, *Output Far-Field Radiation Pattern Measurement*, \$6.00

EIA-RS-455-51-1983, FOTP-51, *Pulse Distortion Measurement of Multimode Glass Optical Fiber Information Transmission Capacity*, \$6.00

EIA-RS-455-87-1983, FOTP-87, *Fiber Optic Cable Knot*, \$5.00

ANSI X3.99-1983, *Optical Character Recognition (OCR) — Guidelines for OCR Print Quality*, \$6.00

ANSI X3.103-1983, *Unrecorded Magnetic Tape Minicassettes for Information Interchange*, Coplanar 3.81 mm (0.150 in) \$6.00

Information regarding standards, publications, and workshops can be obtained by contacting Technical Standards, 627-1800, Leah D'Grey.