"TRANSPUTERS ON THE COMPUTER IMAGE GENERATION"

João Pereira *
Francisco Reis *
Carlos Vinagre *
Mário R. Gomes *

This paper discusses the design of a graphics board with parallel architecture based on Transputers and a resolution of $1024 \times 1024 \times 8$ [1], namely: the processing unit (it plays the role of a display processor), the organization of the frame buffer and the video output hardware which includes the video controller and a RAMDAC (lookuptable + DACs).

DISPLAY PROCESSOR

Parallel processing can be successfully applied as a solution to the Computer Image Generation (CIG) Problem [2]. In fact, the parallelism may be obtained by decomposing the CIG problem into a number of smaller, simpler tasks which can be executed in parallel [3].

One of these tasks is a process called "transformer" that transforms the objects defined in Real World coordinates into viewing coordinates (Image Space) according to the viewing position and direction. Objects are represented by planar and convex polygons. In this way the viewing transformation of objects is reduced to the transformation of points, namely the vertices of the polygons. The transformer process will perform (4×4) -vector-matrix multiplication using homogeneous coordinates, clipping and, then, perspective transformation to obtain the illusion of depth.

Each object transformed is, then, passed to a drawing process (the pixel processing part of the CIG pipeline). Additional transformers may be used to increase the number of drawn objects. As a transformer becomes free, the object store can send another object to transform (fig. 1).

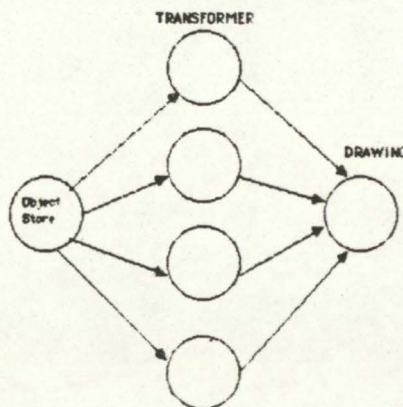



Fig. 1 - Parallel processes in Computer Image Generation



N-transformers can process data at up to N-times the rate one transformer can. This speed-up depends on the rate at which the object store is able to supply objects and on the drawing process being able to draw objects fast enough. The later factor can, again, be achieved by introducing parallelism in the drawing process. This parallelism is reached by considering Image Space partition or Object Space partition [2].

Facing these considerations, IMS-T800 Transputers (20 MHz version) from INMOS [4] were chosen to implement a concurrent display processor. The concurrency of the display processor has two levels: the software level (existence of concurrent processes in each transputer) and the hardware level (several transputers processing concurrently).

The IMS-T800 transputer is a CMOS microcomputer which integrates a 32-bit 10-MIPS integer CPU (20 MHz version), a full 32/64 bit IEEE 754 floating point processor capable of a sustained 1.5 MFLOPS, 4 Kbytes on-chip fast static RAM which can be accessed at 80 Mbytes/s, a memory interface with integral DRAM controller to address up to 4 Gbytes of external memory at a data rate up to 26.6 Mbytes/s and four full-duplex interprocessor communications links on a single chip.

High performance graphics support is provided by microcoded block move instructions which operate at the speed of memory. These kind of instructions can be used in graphics operations such as windowing, zooming, text manipulation and screen updating.

The transputer is designed to implement the OCCAM language, but also efficiently supports other languages such as C, Pascal and Fortran.

Its four serial links provide a simple way to connect multiple transputers together, thus allowing concurrent operation. Interprocessor communications is implemented by connecting a link interface on one transputer to a link interface on the other transputer by two uni-directional signal lines. Each of these interprocessor links has two on-chip DMA engines associated with it, one for input and the other for output. Each link runs at to 20 Mbits/s, providing sustained data rates of 2.35 Mbytes/s, independent of the processor and the other links. So, the four links can run simultaneously, providing data transfers at up to 9.4 Mbytes/s.

The existence of these links together with an extremely fast hardware scheduler on-chip (switching processes takes on the order of 1 μ s) allows the power of parallel processing to be exploited.

Taking advantage of the transputer capability to form networks easily, a ring network with four transputers was implemented [1] (fig. 2).




Fig. 2 - Transputers in a ring network

All transputers are connected to the system bus, so an arbitration on the access to the bus must be done. Each one of them has its own local memory. Using its local memories and its links the problem of saturation is avoided because the transputers accesses to the system bus decrease [1] (fig.3).

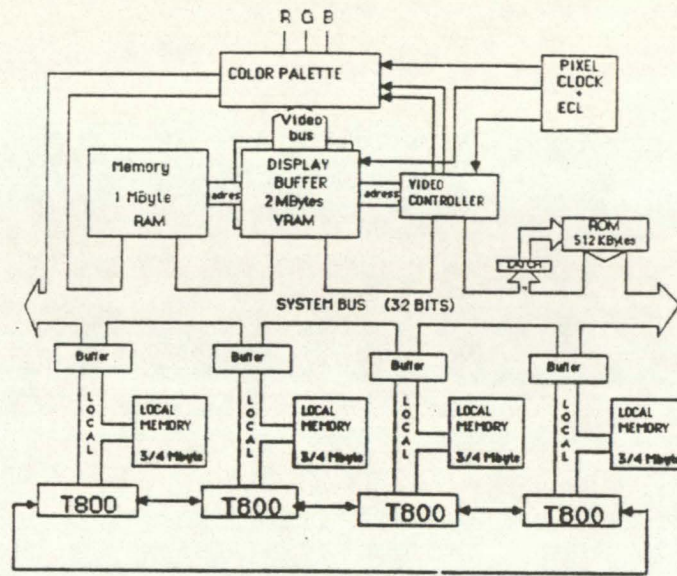


Fig.3 - Blocks diagram of the graphics system.

Lets examine the splitting of the CIG problem into several tasks and how these are distributed among the four transputers.

In the drawing part it was decided to implement the image space partition which leads to several processes, each one responsible for a region of the screen. This partition is suitable to implement the Z-buffer. If the object space partition was used, there would be a lot of dependency between the processors that must be avoided.

To take full advantage of the CIG pipeline, there will be four "transformer" processes whose outputs (transformed polygons) will be local delivered, in a FIFO scheme, to a pipeline of image processes (fig.4).

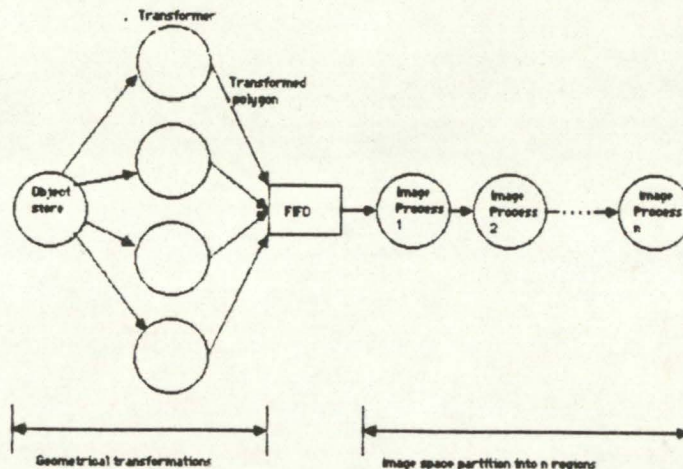


Fig. 4 - CIG tasks

It will be necessary a scheduler process that, according to the throughputs of the geometrical and drawing parts, will control the feed rate of transformed polygons to the image pipeline.

The picture is defined by a display file, containing 3D polygons definitions. To simplify the implementation, the polygons will be restricted to triangles. The display list is located in local DRAM memory and is loaded by the host .

At each vertex of the polygon a depth and intensity is stored. The scheduler process traverses the display file and distributes the polygons among the

"transformer" processes. Each one of these processes can, in turn, be distributed over a pipeline:

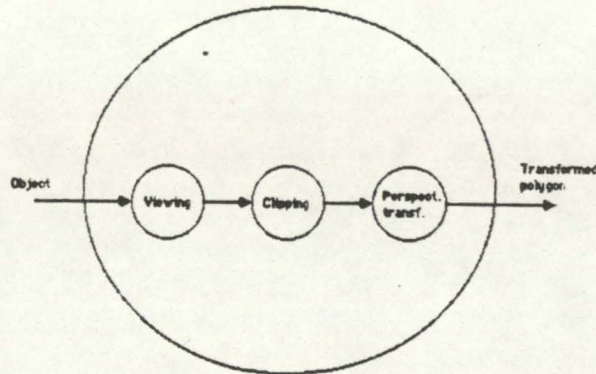


Fig. 5 - The transformer process

Then, the scheduler process determines if the FIFO is ready to accept the transformed polygon. In the positive case, the polygon is passed and the "transformer" process receives another polygon to transform. If not, the "transformer" process must wait.

In the beginning, the display will be broken up into 16 regions. So, in each transputer there will be four image processes (IP), each one with its own local Z-buffer and local frame-buffer ($64K * 2$), occupying a total of 1/2 Mbyte of local memory. The remaining 256K (see fig. 3) will be dedicated, mainly, to geometrical transformations and to the scan-conversion algorithm.

The operations required to render a scene includes :

- scan-conversion
- hidden surface removal
- shading

Simple Gouraud shading will be incorporated in the scan-conversion algorithm, since it involves only linear interpolation between intensities defined at polygon vertices. The hidden line removal is performed by the Z-buffer algorithm.

Now, it can be seen what operations an image process (IP) will perform. At first, details of the edges of a polygon will be calculated to each pixel so that it can decide if it is inside or outside the polygon. Only pixels inside the polygon are eligible for subsequent depth comparison.

Each IP has its own local Z-buffer, so it will compute the depth of the polygon at the pixel and will store it if it is less than the previous minimum. In this case, the IP will calculate the intensity at the pixel by interpolation between values stored at the polygon vertices and will store it in its local frame-buffer.

Again, parallelism can be exploited. In fact, each IP can be splitted into three tasks and distributed over a pipeline (fig. 6).




Fig. 6 - An image process

After a transformed polygon have been processed by an IP, it is fed into another image process that can run on the same transputer or in the following transputer of the ring.

When the last polygon had been processed by the last IP of the image pipeline, the scheduler will spawn a display process that moves the sixteen local frame-buffer to the final frame-buffer. This move involves system bus access by the four transputers, so it will be necessary to do an arbitration of it. This is solved, simply, by making a "token" to travel around the ring. The display process only begins its action with one transputer when this processor has the token. The others cannot access the system bus. After the display process have finished the move, it passes the token to the following transputer and starts again another move.

FRAME BUFFER

The systems design for real-time graphics performance is expensive and difficult because it implies to generate moving images on a computer screen in real time, i.e. the system must be able to recalculate and redraw the image 30 to 60 times per second to give the illusion of smooth motion. An analysis of the architecture of a graphics system shows that the major performance problem is the organization methods used to implement the frame buffer [5]. This memory is accessed by the display processor, which writes data into the memory, and by the video refresh controller, which reads from the memory and routes the pixel data to the video output circuitry and CRT monitor. For the viewing image to be stable, the video controller must supply the data to the monitor according to strict timing requirements. To meet this required data rate, the video controller must have adequate access to the memories. As a result, the display processor and the video controller contend for a finit number of available memory cycles.

So, three problems have to be considered in the design of the frame buffer: the first is the contention for memory cycles and the others two are the adequate bandwidths required for the image memory by the display processor and by the video controller.

The contention of the frame buffer is totally eliminated with the use of double-buffering and almost eliminated with video Rams (VRAMs) [5].

Double buffering with VRAMs was used in the implementation of the frame buffer. One of the goals of this design was the possibility of obtaining real time animation. Only double buffering enables that feature. If it wasn't used, images with parts of the previous one and parts of the actual could be seen on the screen.

The option for VRAMS was done, simply, because it was necessary to guarantee the bandwidth that video controller needs to do the display refresh. In fact, the VRAM internal shift-registers provide a mean to get the necessary pixel rate. Of course, the other solution was to the faster memories but it would be more expensive.

The frame buffer uses TMS4461 64K x 4 memories from Texas Instruments [6]. With a 1024 x 1024 display, 32 chips are required to implement one buffer (so, in total, 64 chips are required for the two buffers) which are organized in 4 rows of 8 memories each, as illustrated in figure 7. The address lines are the same to every line except the upper two bits which are used to select what row is active.

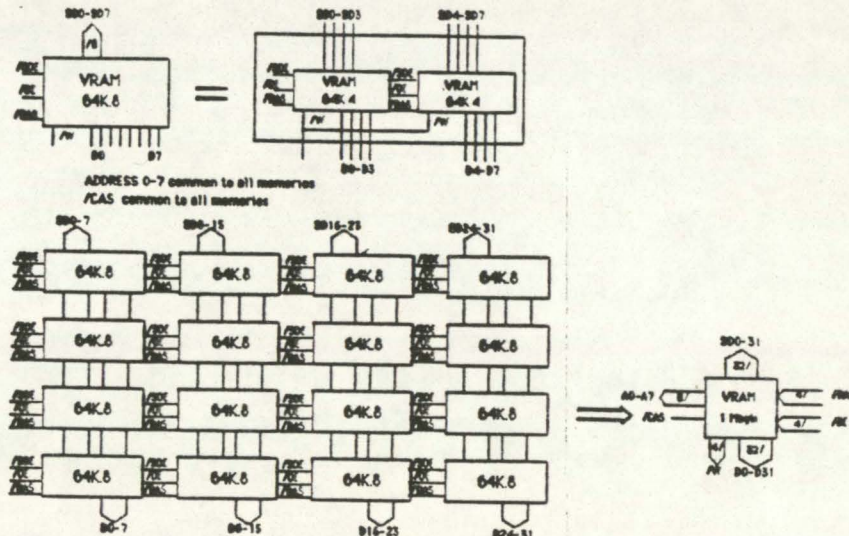


Fig. 7 - Frame buffer organization

It is essential to know how often a new pixel must be supplied to the video output hardware in order to determine how fast the frame buffer must be accessed to support display refresh. In this case, a 1024 x 1024 display with a refresh rate of 60 Hz implies a pixel time of 11.42 ns [5]; so it's necessary a bandwidth of 87.57 Mbytes/s (8 bits per pixel).

Looking again at figure 7, it can be seen that a read operation of the frame buffer by the video controller, taking advantage of its internal shift-registers, occurs only once every four scan lines. So only, a memory cycle is required by the video controller to access 4096 pixels (this read operation generates the data transfer from memory to internal shift-registers). It's important to note that between two read operations, generated by the video controller, data are shifted out a rate controlled by the shift on TMS4461, which handles a maximum speed of 25 MHz, or 40 ns. With this configuration 16 pixels can be accessed per shift operation, thus, the maximum of bandwidth that the frame buffer can guarantee, in this mode, are 400 Mbytes/s which is more than enough.

The overall architecture of the system is represented in figure 3.

VIDEO OUTPUT HARDWARE

TMS-34061 from Texas Instruments [7] plays the role of the video refresh controller. This device controls the video display and the dynamic memory of a bit-mapped graphics system. Primarily designed to provide control of VRAMs, it also controls the conventional 64K DRAMs. Beside generating the video timing signals necessary to interface to a raster scan CRT display, it relieves the processing unit of the burden of controlling the system memory, refreshing video memory and reloading VRAM internal shift-registers. It has 18 programmable registers and an arbiter

which determines whether the host, the video shift-register reload logic or the DRAM-refresh logic can access the memory.

When the video controller accesses the frame buffer it fetches and stores multiple pixels in a video buffer. In this case the VRAMs internal shift-registers function as a large video buffer with a organization of 1024 x 4 pixels (fig. 7) . Then, it is necessary to convert the data stream from parallel pixels to serial pixels. This function is performed by the BT458/125 MHz CMOS RAMDAC from Brooktree [8].

The architecture of this device enables display resolutions up to 1280 x 1024 pixels (up to 8 bits per pixel plus up to 2 bits of overlay information), minimizing the use of costly ECL interfacing, as most of the high speed (pixel clock) is contained on chip. The multiple pixel ports (it may operate either with four pixel ports or five pixel ports) and internal multiplexing (serialization) enables TTL compatible interfacing (up to 32 MHz) to the frame buffer, while maintaining the high video data required for the graphics board. It has, also, a 256 x 24 colour lookuptable with triple 8-bit video D/A converters, supporting up to 256 simultaneous colours from a 16.8 million colour palette.

As illustrated in figure 8, 4 :1 multiplexing is required because the RAMDAC has in its inputs four pixels.

Due to the high clock rates at which BT458 can operate (90 MHz in this case), it is designed to accept differential clock signals (CLOCK and CLOCK*). These clock inputs are designed to be generated by ECL logic operating at +5 Volts. They must be differential signals due to noise margins of the CMOS process. On the rising edge of the signal LOAD, four pixels are latched into the device. This signal is derived by externally dividing the signal CLOCK by four (fig. 8).

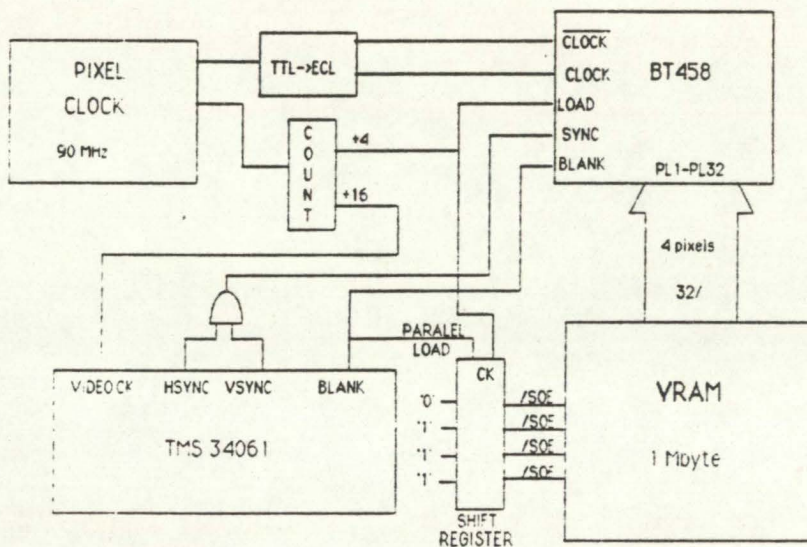



Fig. 8 - Functional block diagram of video output hardware

These were two ways of generation the CLOCK and CLOCK* signals : designing a Pixel Clock with levels TTL and, then, translating them to ECL input levels or implementing, at once, an ECL Pixel Clock. The first method was choosed.

Although, the use of ECL is minimized, some rules must be followed in the design with this high speed logic.

The major extra design constraint of ECL systems compared to any other lower speed logic is that the gate propagation delays and output transition times (~2 ns) are comparable with the signal "round-trip" time along typical on-card interconnects. In these conditions the interconnects exhibit transmission line properties. This has two results. First, propagation delays along interconnects now become a significant part of the total propagation time of signals between circuits nodes. Circuit layouts must



be designed to minimize interconnect distances on critical paths, and most important, clock line lengths must be matched in synchronous systems to ensure that hazards do not occur.

The second effect is that reflections which always occur at the end of a mismatched line are no longer concealed in the transition of the driving rate. A gate with a rise time of 2 ns driving a one foot long interconnect has almost completed its transition before the voltage at the end of the remote end of the line starts to change. Reflections manifest themselves as overshoot and undershoot at the end of the line. Excessive overshoot must be avoided in ECL, because it can cause the input transistor to saturate, and hence slow the circuit down. Equally important, the undershoot can put the voltage into the threshold region and cause false transitions at the output of a driven gate.

The solution to this problem is to use constant impedance and properly terminated transmission lines so as to minimize reflections.

One of the greatest advantages of ECL is the availability of both true and complementary output functions with, essentially, matched delays; these outputs are required to drive the signals CLOCK and CLOCK* of the RAMDAC.

ECL circuits normally operate with ground on VCC pins and a negative -5.2 Vdc power supply on VEE pin. While ECL may be used with ground on VEE pin and +5 Vdc on VCC pins, the negative supply operation has noise immunity advantages and is recommended for larger systems. However, in accordance with the BT458 technical specifications ECL must operate on a single +5 Volts supply. ECL works well in this mode if care is taken to isolate the TTL generated noise from the ECL +5 volts supply line. Translators for interfacing TTL and ECL in this mode are built with discrete components since integrated circuit translators do not operate on a single +5 volts supply (except the 10H350 from Motorola).

The TTL to ECL translator shown in figure 9 consists of three resistors in series to attenuate TTL output to ECL inputs requirements [9].





Fig.9 - Interfacing the BT458 to the TTL clock

As it was before, to minimize the reflections, both the CLOCK and the CLOCK* lines require standard ECL termination and equal length.

A typical ECL parallel termination to a 120 Ω transmission would be the Thevenin equivalent resistor pair (220/330) to +3 volts, as shown in the figure 9.


CONCLUSIONS

The parallel architecture was designed to reach real time performance. However, this requirement depends on the software used in the implementation of the visualization pipeline. In fact, the performance of the display processor is highly dependent of the degree of parallelism that is possible to get from algorithms



(concurrent programming) and how efficiently the tasks are distributed over the four transputers.

Inserting more transputers in the network will be improve the performance of the graphics system but implies a new distribution of the tasks by a major number of processors.



References

- [1] C. Vinagre, F. Reis, J. Pereira, "Placa gráfica baseada em transputers", 3º Simpósio da Electrónica das Telecomunicações, May 1988
- [2] W. Strasser, "Tutorial B1: VLSI-Oriented Graphics System Design", Eurographics 1986
- [3] David May, Roger Shepherd, "Communicating Process Computers", Technical Note 22, INMOS
- [4] INMOS Limited, "Preliminary Data: IMS T800 Transputer", April 1987
- [5] M. c. Whitton, "Memory Design for Raster Graphics displays" IEEE CG&A, March 1984
- [6] Texas Instruments, "MOS Memory DataBook European Edition"
- [7] Texas Instruments, "TMS34061: User's guide", 1985
- [8] Brooktree Corporation, "Preliminary Information BT 458/451", 1986
- [9] Motorola Inc., "MECL System Design", 1983