DISPLAYING THE FUTURE

Mike Addlesee, Chris Turner and Andy Hopper
Olivetti Research Laboratory,
24a Trumpington Street,
Cambridge CB2 1QA,
United Kingdom.

The computer display is no longer restricted to a single machine. With the advent of high speed communication networks the display provides a window on a wider world. Recent and rapid advances in computers and communications have brought with them major changes in the user interface. Interfacing via speech, sound, pointers and gestures have become possible and many applications now include a multimedia element which combines conventional text and graphics output with full colour, motion pictures. Enhanced communications capacity has yielded increased quantity and diversity of information, all of which must be presented to users in a meaningful way. At the same time there is a discernable trend towards mobile working and communicative applications which is moving devices and their displays off of the desktop and into the pocket, or onto the wall.

Recent work at the Olivetti Research Laboratory, ORL, in Cambridge, England has produced a number of operational systems for multimedia communications and mobile working. Pandora and the Active Badge are two examples. Many projects at ORL are carried out in collaboration with researchers at the nearby University Engineering Department and Computer Laboratory and, more importantly, the overall environment at ORL is one of mobility and communications throughout the enterprise. Through their work researchers at ORL have gained experience and have a vision of future requirements for the user interface and display systems in particular.

Probably the greatest advances taking place in display technology are in liquid crystal techniques. Passive supertwisted nematic (STN) technology has shown a steady improvement through the 1980's in terms of response time, contrast, weight and power consumption. Current displays of this type typically have a response time of 250ms. This speed of response appears to be sufficient for a user of a mouse based system, but is inadequate for the display of image sequences as required for a video conference or for video playback as a component of a multimedia application. Increased response time can be achieved, but this tends to compromise the constrast ratio available from the display. Power consumption and weight improvements are being driven by the notebook computer market. The provision of backlighting is a major component in a machine's power budget. Commonly, cold cathode fluorescent tubes (CCFT) are used to provide backlighting and the size of these, and the light guide which is used to provide even illumination of the display, has been steadily reducing. STN displays are usually monochrome, however colour STNs have recently become available. These still exhibit the rather poor viewing angle range of monochrome STN, and the constrast ratio needs to be good to compensate for the transmission loss imposed by the colour filters on the backlight. Consequently, rather more power is consumed to provide the increased backlighting required by these displays.

Active matrix displays are now the dominant technology in high end notebook computers. These displays utilise a new type of active device, called a thin film transistor (TFT), at each pixel location of the display and are hence structurally more complex. This makes them harder to produce and so more expensive. The presence of an active element at each of the matrix of pixel locations results in reduced coupling between pixels as each is addressed, which in turn improves the contrast ratio. A better response time can also be achieved, typically 80ms, making these units more suited to the display of video in a multimedia environment. Viewing angle is much improved over that obtained with passive STN displays. Some TFT displays now achieve a balanced viewing angle of up to 40 degrees from the normal in all directions.

Maximum resolution of colour LCD modules is typically 640 by 480 pixels, but modules of 1024 by 768 pixels are under development by several manufacturers. The current record for highest resolution in an active matrix LCD is believed to be 3072 by 2048 pixels for a thirteen inch monochrome panel made with an amorphous silicon technology. Another manufacturer has demonstrated active matrix displays measuring sixteen inches along the diagonal. Ferroelectric liquid crystal displays have also been demonstrated having the benefit of a pixel matrix that does not need to be continuously refreshed.

At ORL we have been conducting some experiments in the suitability of liquid crystal displays for use in a networked multimedia environment. This work stems from work carried out earlier under the title Pandora(1) and we shall describe that briefly first. The aim of project Pandora was to provide a number of users with multimedia extensions to their existing Unix workstations. This was done by designing a peripheral called Pandora's box which attaches to the workstation to provide audio and video facilities for the system developers and applications writers through the X graphical user interface (GUI). A number of these boxes were deployed and a variety of applications written including videophone and videomail with which to evaluate the system and obtain performance measurements to aid the design of future systems.

To each Pandora's box is attached a monochrome camera, a microphone, a loudspeaker and a local area network (LAN). The LAN uses Asynchronous Transfer Mode (ATM) in its operation. ATM involves the transmission of data in small packets or cells. This allows fine grain control of the delay and jitter experienced by the cells as they pass through the network. As a consequence a network using ATM can simultaneously carry real-time digital audio and digital video traffic along with the more normal transport of non real-time data files. The LAN used with Pandora was based on hardware designed for the Cambridge Fast Ring (CFR) which is well suited to ATM. Video is displayed on the host workstation's display by the Pandora's box. It does this by introducing an analogue video mixer between the workstations display output and the cathode ray tube (CRT) display device. The Pandora's box can then choose on a pixel by pixel basis whether to display the workstation's graphical output or to display video from any number of alternative sources. No inherent limit is imposed on the number of video or audio streams used and streams established by different applications can coexist. Multiple incoming audio streams are mixed together and sent to the loudspeaker.

The method of video mixing employed in the Pandora system is not the only, nor necessarily the best way of performing this operation, but was constrained by available workstation and graphics controller technology when the research was being carried out. In our more recent research we have taken advantage of the maturing technology to examine a different approach. The speed of processors like the Alpha and the increasing use of high speed bus architectures such as TURBOchannel now opens up the possibility of manipulating raw digital video directly within the workstation. Generalizing this, by providing a high speed ATM Network Adaptor within the workstation we can move the various peripherals that would normally reside as cards within the workstation out onto the network. This we have done in our project Medusa. A number of peripherals have been designed that can source and sink data directly off the ATM network. These include an ATM Camera that can source multiple digital colour image streams in a variety of different resolutions simultaneously and an ATM Audio box that can digitize and reproduce mono, stereo or quadraphonic audio to CD quality and again can handle multiple streams. In addition to this we now see that we need no longer limit ourselves to one display tied to our workstation, but may have any number of displays connected directly to the network each of which may be displaying different information.

Another area of research at ORL is in wireless communications, both low rate bidirectional infra-red and higher data rates over radio in the microwave region of the spectrum. We have gained considerable experience through our research on the Active Badge(2), a device designed to assist in the location of people and devices in a computer environment and to provide the wearer with a number of additional facilities such as the ability to 'teleport' ones computer environment to any X display of your choice that happens to be near by. Each badge beacons at regular intervals using short bursts of infra-red light that are received by networked badge sensors distributed at suitable intervals throughout the building to give full coverage. The information gathered by the sensors is made available to users via an X windows interface. Furthermore, a badge may be challenged by a sensor and forced to authenticate itself hence making it suitable as a secure access control mechanism at door entry points. Users can send control or data information to a given badge or set of badges to provide an enhanced paging function. Recent improvements to the Active Badge system have resulted in finer granularity of location through the use of low intensity, low frequency radio fields with a range of around one metre and the introduction of transmit only badges (TOBs) for use on lab equipment.

We foresee a need for user interactive objects that fill the gap between Active Badges and the desk bound Pandora and Medusa Workstations. Our work on LCD modules relates to this. Indeed we envisage the relegation of workstations to a rack somewhere to act as a processor bank for networked peripheral objects with the processor bank or banks providing client-server facilities such as voice recognition and adaptive acoustic echo cancellation (AEC) for the Medusa ATM Audio devices whilst gesture recognition, face identification and JPEG/MPEG compression facilities would be provided for the video streams generated by the ATM Cameras. The ATM Camera has also been designed with a view to including JPEG compression of arbitrary video streams within the device itself.

Our first attempt at producing a potentially mobile LCD based display device for multimedia applications was with a device called the Tablet. This used one of the fastest STN Monochrome displays then available. Support for 64 grey level video was provided by a graphics controller that used a combination of the frame rate duty cycle (FRDC) and spatial stipple techniques. Hardware logic support was provided to ensure that an arbitrary number of video streams could be handled simultaneously. This took the form of a triple ported VRAM framestore. One port being a conventional DRAM port for use by a RISC processor as a memory mapped framestore. The other two ports operated serially, one being used to generate the raster scan pixel stream required to refresh the LCD display. The other was dedicated to the overlay of line segments of digital video into the framestore. These line segments are preassembled as they arrived off the ATM network into a dual ported VRAM double buffered memory and then moved into their correct location within the framestore as specified by instructions passed by the processor. A simple input device in the form of an untethered pen digitizer was provided. The X windows GUI was ported to the R3000 RISC processor used on the Tablet and provision made for control of the X windows system by the pen. A demonstration application was written which allowed the user to move a real-time video stream, originating from a monochrome digital camera also on the Tablet, around the display screen. The results were somewhat disappointing. Although exceptionally smooth movement of the real-time video around the screen with the pen was obtained due to the hardware assist for video, the constrast ratio was very poor. Grey level rendition was average owing to visible flicker and patterning produced by the FRDC and stipple algorithms. We also found viewing angle to be limited. In view of these results we concluded that current STN panels did not provide a quality of display suitable for use with multimedia applications. Work on the Tablet also raised fundamental questions about where the composition of the display should take place. Should the various video streams and computer graphics be combined at the display device itself or elsewhere on the network? In the X paradigm the X server is considered to have direct control over the local display, video streams would need to be combined locally. If the display were composed remotely, by a processor in a processor bank for instance, compression could be applied to the whole display which would allow the mobile display device architecture to be kept simple with the possible provision of a hardware decompression engine. It is important to find a good solution to this problem if best use is to be made of network bandwidth, its impact on mobile wireless communications is particularly serious. A new display device is currently being designed for the Medusa project called the Videotile. This adds a 24-bit plane colour framestore to an ARM processor board with high speed ATM network connectivity. The framestore drives a colour active matrix LCD module capable of displaying three bits per colour (512 colours total). With this we hope to conduct further experiments to resolve the problems mentioned above.

On the software side of project Medusa, active objects called modules are used to represent the Medusa peripherals which may exist either in hardware or software(3). Connections between modules are simple with more complex connections represented by adding special intermediate modules. Security is provided by naming modules with un-forgeable capabilities. We are working to produce an interface to the system that makes it as easy as possible for the user to create custom applications that make use of the Medusa facilities under X. Extensive use is made of the Tool Command Language (Tcl) an embeddable scripting language and the Tk Toolkit X windows widget set extension to Tcl. In general the user can call up an X windows control panel for any Medusa object and manipulate various attributes of that device. For example the control panel for the ATM Audio can be used to pan the spatial position of a mono source when using quadraphonic output. Gestures observed by the ATM Camera have also been used to control the panning function. We are currently in the process of integrating an autostereo display into the Medusa system.

(1) Hopper, A.: *Pandora - An Experimental System for Multimedia Applications*, ACM Operating Systems Review, Vol. 24, No. 2, April 1990.
(2) Want, R., Hopper, A.: *Active Badges and Personal Interactive Computing Objects*, IEEE Trans. on Consumer Elect., February 1992.
(3) Wray, S., Glauert, T., Hopper, A.: *The Medusa Applications Environment*, ORL Internal Report, December 1993.