# ..... a thousand words
## (The development of computer graphics)

*Inaugural Lecture of the Professor of Applied Computer Science delivered at the College on November 14th, 1988*

*by*

PROFESSOR PETER TOWNSEND
B.Sc., Ph.D., M.B.C.S.

GWEDDW CREFFT HEB EI DAWN

UNIVERSITY COLLEGE OF SWANSEA

First published November 1988 by the University College of Swansea. Obtainable from the Registrar's office, University College of Swansea, Singleton Park, Swansea SA2 8PP.

ISBN 0 86076 056 1

## ..... a thousand words
## (The development of computer graphics)

### 1. Introduction

Inaugural lectures by Professors of Computer Science at Swansea are few and far between. In fact this is only the third in the history of this College.

The last such lecture was given by Professor Niblett in January 1977. The progress made in Computer Science and in the development of computer hardware since that time has been phenomenal. Just one example will illustrate this. Micro-computers - IBM-PCs, BBC-micros, and so on, so commonplace in the College these days were totally absent when Professor Niblett gave his lecture. Indeed the very first micro-computers - the Commodore Pet, the Apple and the Tandy were released by the manufacturers that year.

The main theme of this lecture is one specific branch of computer science, that of computer graphics and here progress has been quite spectacular. I must control my use of superlatives but it is difficult to describe these changes in any other way. One can chart the major steps forward in the field in parallel with innovation in hardware technology, particularly in relation to the computational power of machines and the cost of memory chips.

I should like to go back to the very early beginnings of computer graphics and pick out just a few significant events as set out in figure 1. In doing so I would also like to parallel this progress with the development of 'computing' within the College. It is difficult to pinpoint the exact birth of computer graphics but certainly by 1949 the EDSAC computer at Cambridge had a number of cathode ray tube displays. These were intended to diagnose hardware problems but programmers soon realised that one could use them to display simple graphics. This inadvertent discovery of computer graphics was followed very soon after by a more positive step when the Whirlwind computer at MIT was fitted in 1950 with a cathode ray tube specifically for plotting pictures. Incidentally, it also carried a light-pen which allowed some primitive interaction. Amongst other things, the Whirlwind display could be used to plot the solution of differential equations. The classic example recorded on film is the computation of the path of a ball bouncing on a flat surface.

These early beginnings in the field of computer graphics did not lead to any significant progress for some considerable time. It wasn't until 1962 that Ivan Sutherland (1) published his doctoral dissertation entitled 'Sketchpad: a man-machine graphical communication system'. This work was to establish the basis for interactive computer graphics. Sketchpad was a program which allowed a user, by interacting with a screen using a light-pen, to input basic geometric shapes such as vectors, arcs, circles, rectangles, etc. The software included facilities to automatically modify the graphics so that, for example, a circle could be drawn up roughly with the light-pen and subsequently improved to perfect roundness. Sketchpad itself led on very quickly to more sophisticated ideas, including, even in those early days, the ability to draw and visualize objects in three-dimensions. These ideas were soon adopted by the engineering world and one saw an embryo CAD (Computer Aided Design) environment developing. As early as the mid-sixties General Motors were using this newly

## A History of Computing

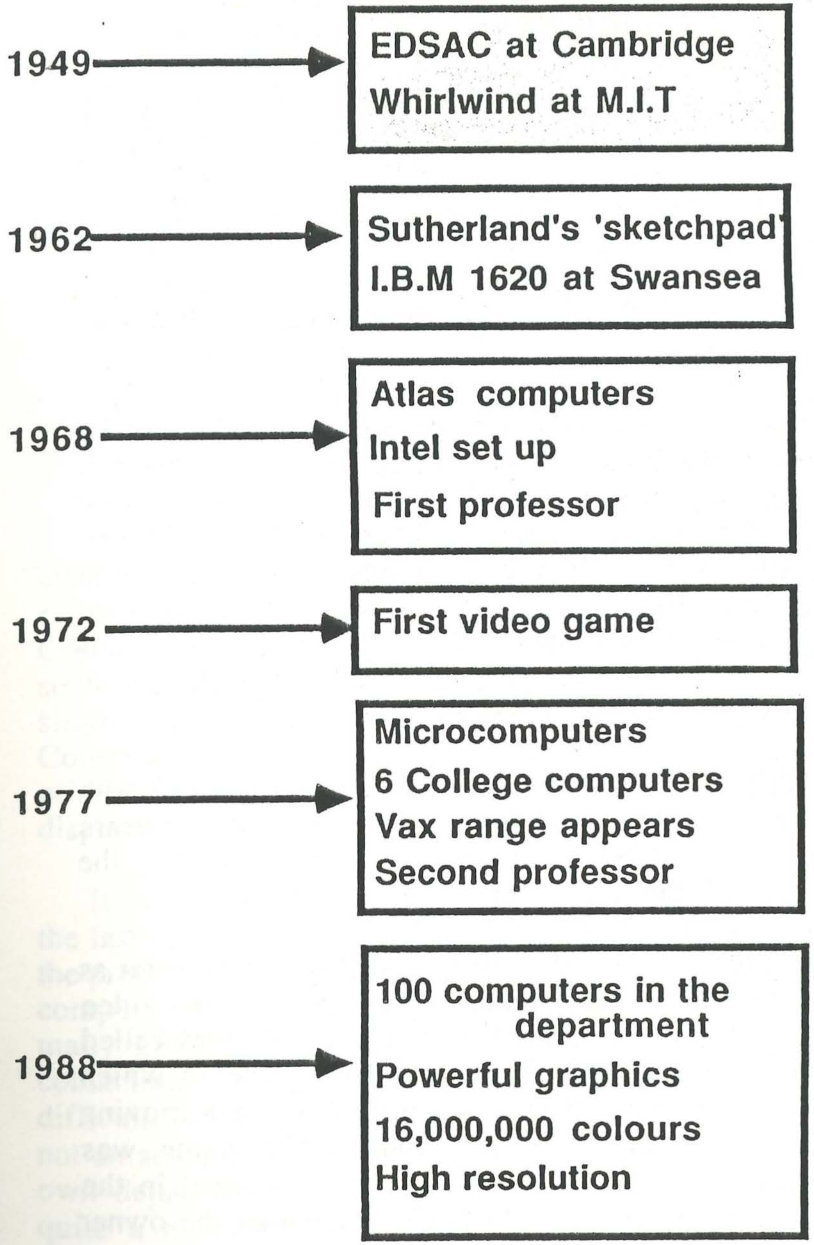

Fig.1

discovered tool to help design car bodies.

This exciting time for computer graphics was exciting also for the College, for it was in 1962 that the first computer was installed in what was then the Department of Applied Mathematics. The computer was an IBM 1620 and in the College's annual report for that year it states that a five day course on programming for computers was held ... attended by over 100 members of the College - clearly an enthusiastic start.

The year 1968 is not a major turning point in the history of computer graphics but the College that year made a significant step forward for they appointed the first professor of computer science at Swansea, David Cooper. At this time I remember, how, as a postgraduate student at Aberystwyth I would send off my 'large' computing jobs to the mighty Atlas computer at Rutherford. The Atlas computers when they were first introduced were thought to be the most powerful in the world.

One little pointer to the future also arrived that year with the setting up of the Intel company, who within two years was producing the first microprocessor chip, the 4004 - the forerunner of all our microcomputers.

In 1972 there occurred quite a significant event as far as computer graphics is concerned. In that year the first video game was introduced in the United States. It was called 'Pong' and was a simplified table-tennis game in which players, using simulated bats, attempted to hit a moving square of light representing a ball. The game was subsequently to be renamed 'Ping' when it arrived in the U.K. Its inventor, Noland Bushnell, persuaded the owner

of a local bar to install the equipment on a trial basis for use by the customers. Bushnell was disappointed to receive a call very soon after to say that the equipment was broken. However it transpired that it was not the equipment that was at fault. The game was proving so popular that the coin collection box inside had been jammed full. This fault was easily rectified.

This popularity of video games led much later to the boom in ownership of 'home micros' which in turn provided one of the motivations for the development of computer graphics techniques particularly in relation to interactive software.

Home micros did not appear, as I've already mentioned, until 1977. At the same time the very successful VAX range was introduced by the Digital Equipment Corporation. In the College there were just half a dozen or so computers in total. Virtually all of these were single-user machines and all were very expensive. Computer graphics in the College was still at a quite primitive stage with only vector facilities and monochrome displays.

It is only a relatively short jump from that year, when the last inaugural lecture in computer science was given, to the present time. There are certainly more than half a dozen computers on campus today. I tried to estimate just how many computers there are in the College but without contacting almost every member of staff this task is very difficult, for now every single department is a potential if not an actual computer user. Suffice to say that within my own department there are over 100 separate computers, quite a number of which are more powerful than that 'ever-so-powerful' Atlas computer I referred to earlier.

In terms of graphics capability we have very powerful high resolution workstations capable of displaying photographic-like images in over 16,000,000 different colours. But I must not say too much for how foolish this will look in just a few short years.

## 2. ... a thousand words

The origin of the saying ' a picture is worth a thousand words is a little vague but most attribute it to the ancient Chinese. But there are pictures and there are pictures. Some will presumably be worth more words than others. What qualities are we interested in?

For many the drawing of just a few lines on a screen may convey all the information that is necessary. For others the addition of colour shading of the picture is a requirement. For the true connoisseur one's ultimate goal must be a photographic quality image with full realism. In many applications these days there is also a requirement for real time interaction or animation. This conflicts with the image quality attainable within available resources. One must compromise and settle for either static photographic quality or real time interaction with much reduced quality. Both are illustrated later in this lecture.

I referred above to a realistic image. What do we mean by realism? How do we inject realism into our displays? Consider a child's picture of say, two people (fig.2). We would recognise immediately that there are two people in the picture but the level of realism is low. A professional

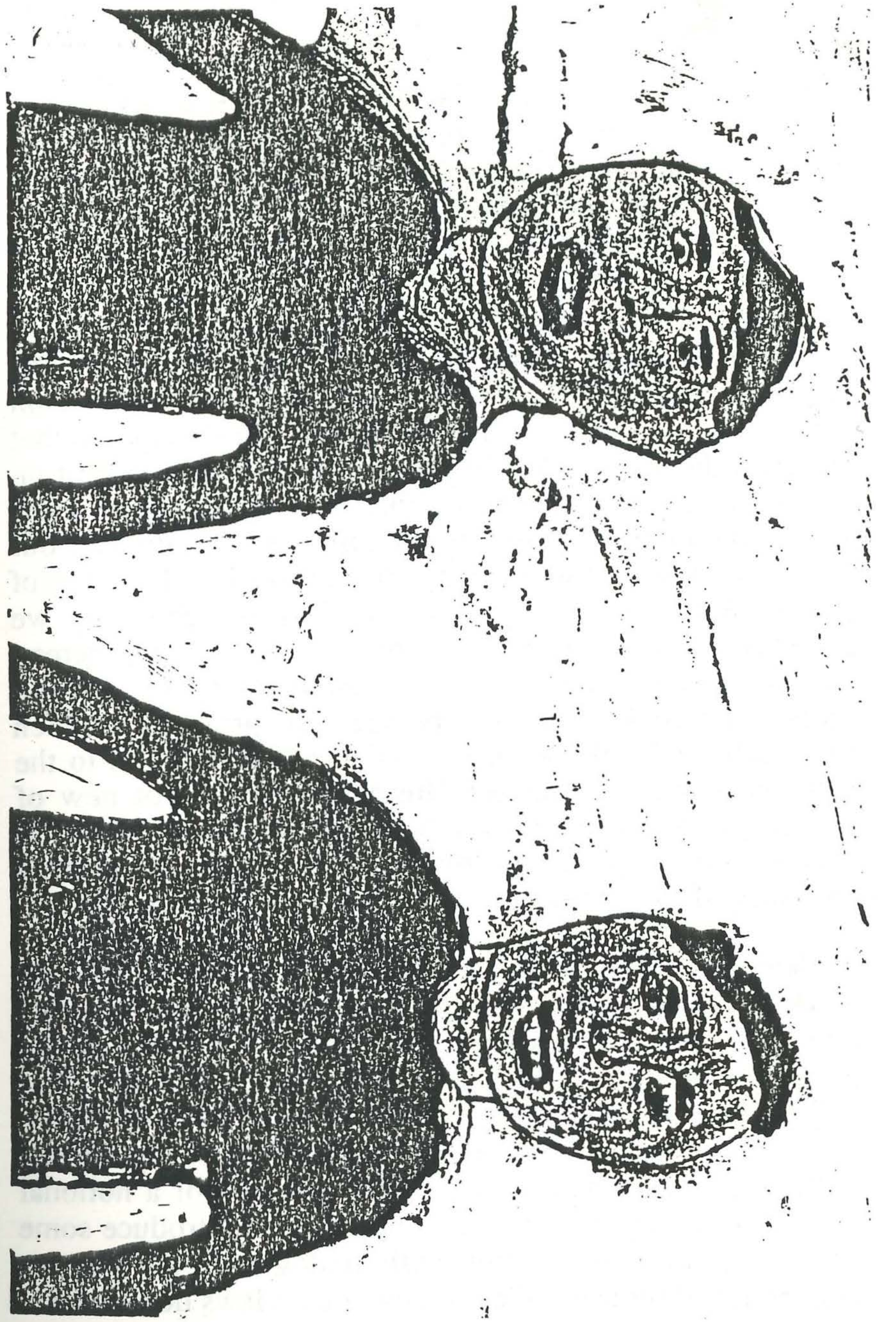

Fig.2

artist painting a similar picture would, apart from making use of a more developed drawing ability, introduce effects which would help our brains see the image with three dimensional depth. But still the image would lack photographic realism, often of course quite deliberately.

If we are to use computer graphics to display a three-dimensional scene we must mimic many of the techniques used by the artist. But first we must decide on some method of representing the surfaces of the objects that make up the scene. A popular way to do this is to split up the surfaces into planar polygonal patches which are more easily manipulated than curved surfaces. To display our scene we then start with a wire frame model of the edges of our polygons and using a little elementary geometry we project our wire-frame image onto a two-dimensional screen (fig.3). In doing so, in order to retain some depth in our picture we make use of a perspective projection which essentially reduces the size of objects in proportion to the distance from the viewer. This technique is not new of course. It has been used by the artist since the thirteenth century. If we produce a number of such images using a sequence of viewpoints the effect can be quite dramatic.

But these images are far from realistic. There is much work to be done of we are to approach photographic quality. Our first task is to remove those parts of the picture which are hidden from the viewer (fig.4). This helps greatly to enhance the feeling of depth and this feeling is further enhanced if we shade the surfaces of our objects with an intensity determined by the position of a notional light source (fig.5). There is scope here to introduce some physics into our model but, with apologies to my physics colleagues, it turns out that inverse square laws don't 'look

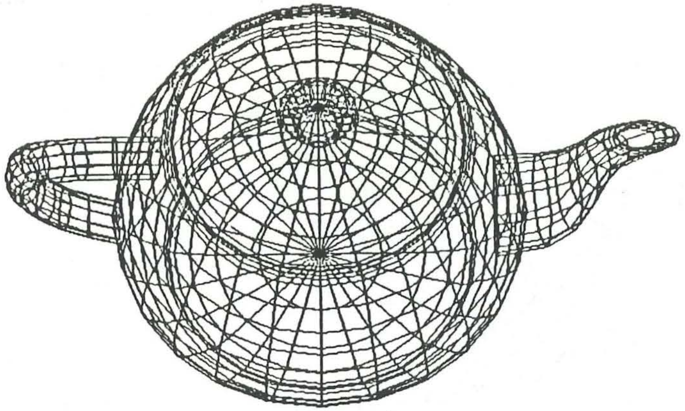

Fig.3

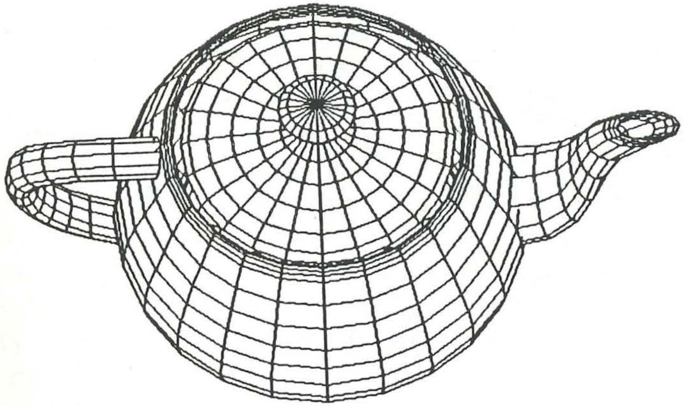

Fig.4

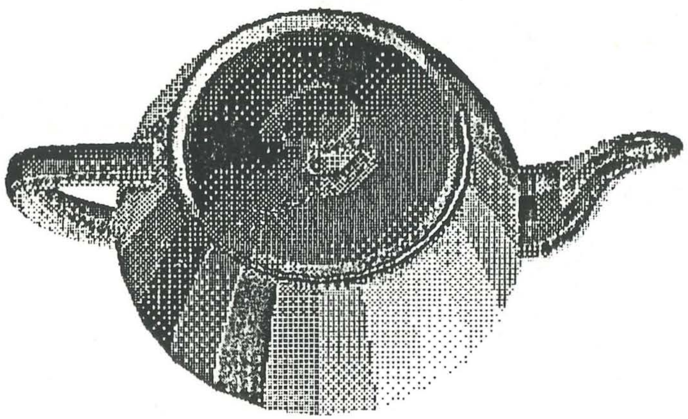

Fig.5

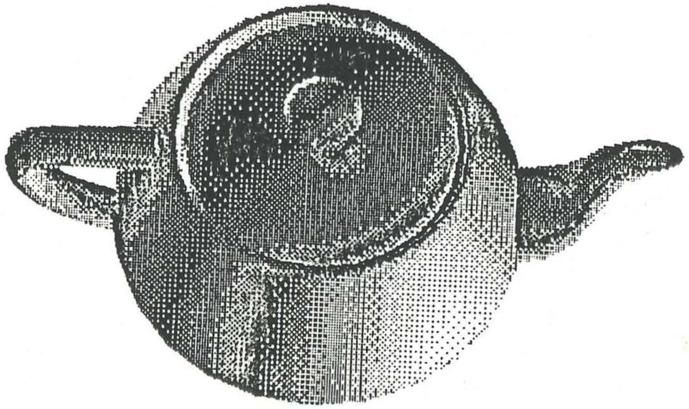

Fig.6

right' and we settle for a visually more pleasing ad hoc shading law which has no physical basis whatsoever. If we now use a little interpolation of surface normals to smooth rough edges (fig.6), add a few highlights to create an illusion of a shiny surface (fig.7), and finally map onto our object some texture pattern (fig.8) we end up with an image which definitely has some feeling of depth. But we are not finished! There is more to be done. The human brain is used to even more depth cues such as for example shadows. Without them we feel there is something missing. The picture seems odd. Then there are other effects. Shiny surfaces give rise to reflections; transparent or translucent materials transmit light. To account for these and other effects is difficult and one must make peace with the physicists again and start to look more closely at the optics. The current trend is to attempt to follow rays of light as they leave a source and reflect off or refract through, objects and finally reach the viewer. The computational effect required to generate such images is considerable but even on a relatively cheap home micro-computer the results can be quite startling, particularly if you save a set of such images in memory and then display them in an animation sequence.

We have concentrated in the above on the efforts made to render surfaces in a realistic manner according to the light falling on them. The approximation of surfaces by polygonal patches is a relatively straightforward task if those surfaces are mathematically well-defined. The vast majority of objects in our world are unfortunately not that easily defined. Most of us would find great difficulty drawing realistic images of objects such as plants or trees or mountains. Such objects have a multitude of component parts arranged in a structure which has just that appropriate degree of randomness that we find difficult to copy. For the

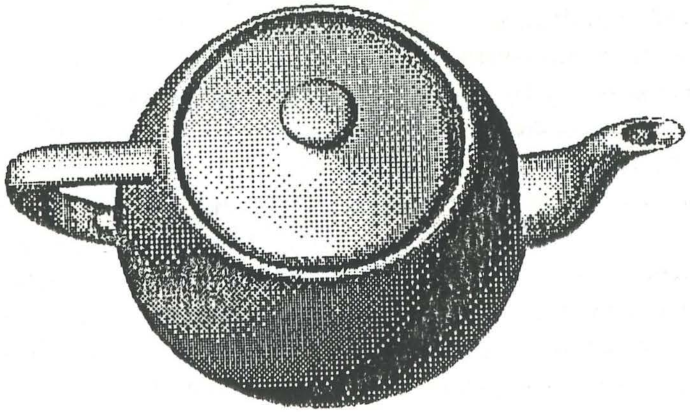

Fig.7

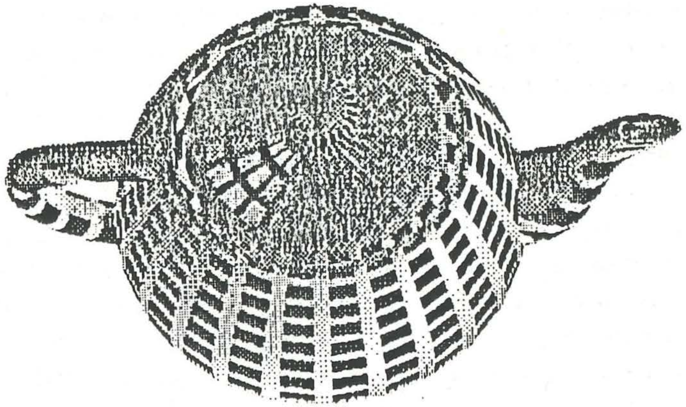

Fig.8

computer, of course, the task is immense, for if we are to use the techniques described previously, each leaf on a tree, each blade of grass and each rock on a mountain side must be represented faithfully as a set of polygons. Fortunately we are not very observant. We might recognise Snowdon or Cader Idris from their general shape but we are unlikely to remember any great detail.

We take advantage of this human failing and use a technique which introduces a 'natural' randomness and yet, at the same time, preserves the basic prescribed overall structure such as the general shape of the mountain. This technique, which is derived from the study of fractals in mathematics, uses a recursive algorithm to generate the required polygonal representation automatically from an initial specification of just a few polygons. These are chosen to be just sufficient in number and detail to roughly represent the overall shape of the surface to be modelled. Each of the initial set of polygons is subdivided recursively and the co-ordinates of the points of subdivision, perturbed at random in three-dimensional space. The effect is that from one initial polygon one obtains a set of smaller polygons each of which has a slightly different orientation, and consequently a different light reflection property. One terminates the recursion at whatever level of detail that is required. Superb examples of scenes generated by fractal-based algorithms are given by Benoit Mandelbrot and Richard Voss of IBM research (2), and by Loren Carpenter of Boeing Aerospace(3). Another technique used to model 'natural' effects is that of particle systems. Very small dots of light are generated in different ways to compose various effects. The particles are 'made to move' and their tracks recorded like streaks on a photographic film. Such a technique has been used by

Reeves (4) to generate, for example, images of flamés or explosions. One can also adopt this method to even grow blades of grass.

The technology that is applied to the generation of computer graphics images becomes ever more sophisticated, limited only by the computational power needed. As that power continues to increase we will see the generation of ever more realistic images until the day arrives when it becomes difficult to distinguish between computer generated images and photographs.

## 3. 'A solution looking for a problem'

Despite the comparatively early beginnings of computer graphics and despite even, the enthusiastic adoption of the technology by companies like General Motors, the take up by other application areas in the early years was small indeed. Much effort was put into developing the mathematical techniques and computing algorithms needed, particularly for three dimensional work. Graphics still attracted derogatory remarks from some 'non-graphics' computer scientists who referred to the field as the 'solution looking for a problem'. From the mid 70's cheap memory chips and the readily available raster scan display technology began to bring about quite dramatic changes. Today the adoption of computer graphics is so widespread that to cover all possible applications would require a course of many lectures. We choose instead to pick out a few representative examples, many of which may be so familiar to us that we already take them very much for granted.

We start with computer games, an area which, as indicated earlier, has had quite a significant influence on the development of computer graphics. Anyone who has children must surely at some stage have been dragged into that area of the motorway service station where racing cars roar and space battles rage - the games arcade. The level of sophistication of such games today is quite remarkable. The player interacts in real time with a three-dimensional model, often using stereoscopic pairs of images and special viewers to enhance the three dimensional nature of the display. No great level of realism is needed of course. The real time response is far more important and requires very clever programming techniques. Arcade computers have special purpose hardware for a specific game, but even general purpose home computers have sufficient power to implement quite advanced games. For example, one can interact with a three-dimensional model of the air space over various cities in the United States and, using a simulated instrument panel, fly an aircraft from one airport to another. Of necessity the graphics is very simple but certainly adequate to provide the necessary feedback for controlling the interaction.

We, in education have certainly found a place for computer graphics. At the primary school level, for example, the use of microcomputers is now commonplace. Ideas can be introduced to young children pictorially, without the need for a great reading ability. As the children develop, they progress through a wide range of available software which allows them to experiment for themselves. In some of the early work in this field, notably in the Pluto project in the United States it was found that the introduction into the curriculum of sessions at graphics terminals greatly enhanced the motivation of the children. Indeed they found that the only inducement required to persuade the children to complete set projects was to allow them extra computer time to work on projects of their own.

At a more advanced level in education, computer graphics can play another role. Laboratory experiments which involve the use of expensive equipment or materials are often now replaced by simulations with graphical feedback as the experiment proceeds. The advantages are obvious. The student can repeat a simulated experiment as many times as he or she wishes without additional expense, without harming possibly delicate equipment, and without harm to himself or herself if the experiment concerns some dangerous process.

Computer graphics seems almost an unwelcome intruder into the traditional disciplines of science. After all the body of scientific knowledge was firmly established long before computer graphics appeared on the scene. However, it has a role to play and scientists have been quick to realise the potential of this new tool. For example, in Chemistry, molecular structure may now be studied using simulated three-dimensional models. The viewing point can be changed to allow one to fully appreciate the underlying structure, and animation can be used to visualize processes which occur within time spans of a few millionths of a millionth of a second. Interaction with a model allows one to understand the means of attachment of various drug molecules and in this way new drugs may be 'designed'.

In medicine, computer graphics has been used as a guide during plastic surgery. An image of a patient who has suffered from severe burns or facial fractures is compared with a normal one and the differences quantified to aid the work of the surgeon. In another application internal organs of the body have been digitized to build up display models which may subsequently be used to teach medical students. At San Diego, for example, a human

brain was sliced into very thin sections and the data entered into a computer. Computer graphics was then used to conduct the students on a guided tour around and through the brain.

In engineering, too, there are so many applications which make use of graphics. In electronics it is an essential tool in the design and manufacture of VLSI chips, in some of our own work here at Swansea, graphics plays the major role in research into computer-aided control system design. The control engineer is given the opportunity to interactively draw up pictorial repsentations of control systems and simulate their performance. The system can be edited pictorially on the screen and all of this without any need for programming expertise.

One cannot mention engineering in a Swansea context of course without some reference to the finite element method. Finite elements are used for a wide range of engineering analyses. The method requires that one generates a set of points or nodes in space which form the vertices of the so-called elements. In three-dimensions these elements might be tetrahedra or rectangular blocks and it is a major task to subdivide a complex three-dimensional shape into a very large number of elements. Having generated the data, computer graphics becomes an essential tool for checking its validity. For example, do the elements fill the space required? One also uses graphics interaction to modify the placement of nodes and to input the bounding surfaces of the solution region and the boundary conditions. In some of our own work we use finite elements to simulate the flow of polymer melts through very complex three-dimensional geometries. Here, as well as the problems of generating the description of the flow region and the finite element mesh, an equally difficult problem is to animate the transient solution derived from the finite element analysis.

An almost integral part of engineering these days is the field of computer-aided design which itself relies heavily on computer graphics as a tool for inputting, editing and displaying designs of all sorts of articles, including machine tools, textiles, car bodies, buildings, etc., etc. For some such applications the skills of an artist are used in conjunction with so-called 'painting' software to input designs or images. Such software allows the artist to select from a wide range of brushes and colours and to paint on the screen using a light-pen or mouse as input device. Tedious repetitive operations that one might need in say textile design, are avoided by providing 'cut and paste' operations. For architectural applications, the advantages are clear. One can 'graphically' construct a new building and view it alongside the existing buildings to see how well it fits in with existing architecture. One can 'walk through' the three-dimensional model and hopefully detect any design shortcomings before committing to the expense of the actual construction phase.

A somewhat less likely candidate for graphical displays is that of the office situation. It is true that statistical information is often presented in the form of pie-charts or histograms or some other graphical form, and certainly modern spreadsheet packages incorporate such graphics facilities. Also project planning software provides the manager with graphical interfaces to construct networks for critical path analyses. But these tools are essentially management tools. Where does computer graphics come into the work of the average secretary? The answer is in the operating system interface provided by the manufacturer. A secretary does not want or need to know about the operating system's command language. To make this unnecessary some manufacturers now provide their systems with a graphical interface which is in the form of a picture of a desk top. On the desk top are little images (icons) of familiar items such as folders, in and out trays, a waste-paper bin and so on. The secretary interacts with the interface in a way that should mimic tasks that are quite familiar. To locate a document one opens a folder. To mail (electronically) a document one places it in the out tray. To discard a document it is put in the waste-paper bin and so on. Each of these operations is carried out by pointing and placing with a mouse. There is no need for the secretary to be aware of a command language. The use of icon-driven interfaces was first introduced by the Xerox corporation (5) but is now readily available on most computers used in the office environment. One other less obvious use of computer graphics in the office is in the area of desk-top publishing. When the 'golf-ball' typewriter was introduced it brought with it the capability of using different golf-ball heads i.e. different fonts for typing documents. More recently, with the introduction of laser printers, the typographical opportunities have expanded enormously. We now expect our secretaries to use a multitude of different fonts and point sizes. Documents take on the appearance of having been professionally printed. All this is achieved through the use of computer graphics both to display the multitude of different character sets on the screen and to drive the laser to 'draw' the characters on the page.

Perhaps the most obvious exposure of the general public to computer graphics is through television and advertising.

On television it is now not sufficient to display titles or station logos in a static form. Instead they must come alive, move around the screen, deform or explode and re-establish themselves in some other form. One of the earliest software packages for this type of work was Antics developed originally at Rutherford in the early 70's. Amongst its many varied applications, one of the more noteworthy is the version sold to Swedish television in 1978 and used for animating the smiling or frowning faces of politicians in the Swedish general election. Television news services are big users of computer graphics tools for without them it would not be possible to react sufficiently quickly so as to produce for example illustrative animated maps for a news item which may only have become known a very short time before the broadcast.

Although used in other areas of television for many years, computer graphics has had comparatively little impact on television commercials until quite recently. The reason for this is not entirely clear but it may be that the quality of image was not acceptable to the advertising executive. Certainly today those advertisements which rely on computer generated graphics - and there are now very many - make use of some of the most powerful computers and the highest quality displays available. Live action is often mixed with graphics sequences so that one might, for example, show a real can of oil being poured into a car engine and then, using graphics, follow a simulation of the oil's journey as it passes through the various parts of the engine.

The compromise between the level of realism and the time taken to update an image is never more apparent than in flight simulators where pilots sit in simulated aircraft cockpits and, using the normal controls, interact with display screens placed where the windows would normally be. It is not necessary for images to be of photographic quality but sufficient cues must be provided for a pilot on a low flying exercise to be able to 'feel' the three-dimensional nature of the landscape and to appreciate the speed at which the aircraft is travelling. Early simulations used flat shading for mountains and pilots just flew through them! In the present day simulators the landscape is carefully textured with regular features such as trees or rocks. The simulators are now so successful that, for example, Royal Navy helicopter pilots spend as much as 30% of their initial 'flying' training on a simulator before moving to the real thing.

The level of realism used in flight simulators is not high, being limited by the need for real time interaction. In graphics sequences used in feature films, however, realistic animation is of paramount importance. In the past, where it was not possible to film actual action, when that action was dangerous, very costly or even impossible, scale models were substituted. A classic example is the dividing waters of the Red Sea, a model which was used for the film 'The Ten Commandments' and which can still be seen today at Universal Studios in Los Angeles. For the spate of recent science fiction films in many cases models have been replaced by computer generated imagery. The computing power required in order to satisfy the quality constraints is however, immense. If we take just one example - that of one second of animation from the film 'The last starfighter', in the most complex cases as many as 1,000,000 polygons were used to represent scenes. A tremendously powerful Cray X-MP/22 computer was used to generate the required images but even this needed as much as ten hours computing time to produce one second of animation. To put this into perspective and relate the task back to where we started ,- the IBM 1620 computer, a rough estimate shows that if in 1962 we had set that machine to do this work, and assuming the machine continued to run without fault, we would expect to complete the calculations needed for our one second of animation in the year 2402! What is surprising is that the use of such a powerful and expensive Cray computer is cost effective – in fact half the cost of producing the scene any other way. This has led to some quite unbelievable predictions that one day complete feature films will be entirely computer generated, including appearances by 'simulated' actors. This may well be quite some time in the future but with the multitude of TV channels with ever more hours to fill the motivation for automated films is clear. The quality is the big unknown but then can they be any worse than today's soap operas!

In summary, it is clear that computer graphics has now come of age and 'solves many problems'. From scientific experimentation to commercial exploitation, from education to leisure time relaxation, one could go on and on. But what of the future? Have we reached the limits of the field?

## 4. The future

In writing this lecture I did of course read very closely the two previous computer science inaugural lectures. Both the previous professors were very careful not to speculate about the future and I should follow their example. When the ACE computer was introduced in the late 1940's it was capable of multiplying 500 nine digit numbers every second. It was claimed that we in the United Kingdom would never ever need more than four such computers to do all the computing required in the country because there would never be sufficient mathematicians to feed the computers with work at a faster rate than these four could deal with. Today's supercomputers are capable of 10,000,000,000 such multiplications every second! It's clear that predicting the future in a computing context is a difficult task. There are a few pointers however. Computer manufacturers are already promising us a tenfold increase in the power of conventional graphics worksations by the very early 1990's. At the same time we are seeing the development of specialised multi-processor graphics hardware which gives even greater display capabilities. There are machines now, for example, which contain over 65,000 processors capable of simultaneously adjusting each pixel of a 256 x 256 block. Great power isn't everything of course. One must make use of it efficiently. There is the story of how the Siamese government were persuaded by UNESCO to perform a census in 1947. A whole array of IBM equipment was supplied to do the job but UNESCO officials were surprised that after eighteen months no results had appeared. They called in experts to sort out the problem. It transpired that the Siamese were so distrustful of the IBM machines that all output was fed to one single Chinese girl who checked the results with an abacus. Assuming we do find an efficient way to use our new hardware, it is probably not too rash a prediction, therefore, to expect in the 1990's to be able, using a desktop machine, to interact in real time with a fully rendered realistic three-dimensional scene.

We shall have to wait and see .... but not for too long.

## References

1. 'SKETCHPAD - A Man-Machine Graphical Communication System' by I.E.Sutherland, M.I.T. Lincoln Laboratory Technical Report # 256, 1965.

2. 'The Fractal Geometry of Nature' by B.B.Mandelbrot, Freeman Press, New York, 1982.

3. 'Computer Rendering of Stochastic Models' by A.Fournier, D Fussell and L.Carpenter, C.A.C.M. Vol. 25 No.6, 1982

4. 'Particle Systems - a Technique for Modeling a Class of Fuzzy Objects' by W.T.Reeves, A.C.M. Transactions on Graphics Vol.2 No.2, 1983.

5. 'Designing the STAR User Interface' by D.C.Smith, C.Irby, R.Kimball and B.Verplank, BYTE Vol.7 No.4, 1982.

*Throughout this lecture a number of slides and video sequences were used to illustrate various points. It is clearly not possible to reproduce this material here other than the few monochrome figures quoted.*