17

Applying Solid Modelling and Animated Three-Dimensional Graphics

Paul Reilly*

Stephen Shennan†

17.1 Introduction

Spatial distributions of archaeological material have long formed one of the main sources of information for the archaeologist at both the site and the regional level. In the last twenty years they have been the object of intense interest and investigation as the aims of archaeologists have grown more ambitious and they have tried to use the evidence of such distributions to make inferences about the nature of past forms of society and economy. The main method of carrying out such work has been the use of statistical methods of spatial analysis derived from geography. These remain important and useful for certain purposes but it has become increasingly clear that they are too simplistic for many archaeological needs. The assumptions behind them and the information they can provide do not do justice to the subtlety and complexity of the patterning in spatial distributions of material, or the questions which archaeologists ask of them. This is especially the case at the intra-site level, where answers to such simple questions as the degree of randomness or clustering in a distribution are rarely of interest. For reasons such as this, among others, there has recently been a renewed interest in data visualisation: visual examination of the data distributions as a basis for exploring patterning in the data and attempting to understand the processes which brought it into being. Such an approach undoubtedly has dangers in the well-known capacity of the human eye and brain to detect patterns where there are none. Nevertheless, as a basis for data exploration and hypothesis construction, it is of vital importance.

Until the advent of computer methods, data visualisation on any scale and with any degree of sophistication was impossible owing to the sheer volume of work required to map distributions by hand. Some time ago, archaeologists started using two-dimensional contouring programs and three-dimensional surface generating programs, such as SYMAP and SYMVU, in order to overcome some of these difficulties, but such methods were still relatively unsophisticated and their batch mode of operation too slow and unwieldy to enable the exploration of data in a systematic fashion, and graphics screens were inadequate for anything more than a basic presentation of the data.

Recent developments in hardware and software have started to change this picture rapidly. Real-time three-dimensional vector graphics models of data have begun to be used, which enable genuine interaction between the investigator and the data in a stimulating fashion. However, current state-of-the-art graphics systems rely on solid modelling and, more recently, animation techniques. The ability to move around solid models of data or reconstructions enables further insights to be obtained. Such

* IBM UK Scientific Centre
St. Clement Street
Winchester S023 9DR

† Department of Archaeology,
University of Southampton,
Southampton S09 5NH

techniques have important advantages over real-time vector graphics models of the same data.

In this paper some of these advantages will be identified and discussed in relation to the excavation of an early bronze age settlement at St.Veit-Klinglberg in Austria, where the questions being asked and the complexity of the distributions of material require the use of such advanced techniques. The early Bronze Age settlement at St.Veit-Klinglberg is noteworthy for being the first such settlement to be examined on any scale using modern methods. Area excavation techniques are being used in the main part of the site. In the upper levels, where stratigraphic layers are not detectable, the basic recording unit has been arbitrary spits 2.0 × 2.0 × 0.15m. Where layers and features have been found they have been excavated stratigraphically in the usual way. In the case of layers of large extent, however, spatial control is maintained by sub-dividing the layer into artificial box contexts. So far no coherent structure plans have emerged, and the relationships between finds and features cannot be considered straightforward, especially in the light of the site's undoubtedly complex depositional history (Reilly *et al.* 1988, p. 289).

17.2 Preliminary experimentation

Data defining the location of all eight vertices of each arbitrary context and information relating to the objects found within that volume of space have been entered into the Winchester Graphics System, which combines all the facilities of a relational database with the power of interactive three-dimensional vector graphics in a single environment. Previous research on a bronze age midden at Potterne, where arbitrary box contexts had also been employed, had shown that a lot could be learned about the distribution of objects and the formation processes present at the site by studying appropriately colour-coded wire-frame representations of the contexts dynamically, on a suitable workstation (*ibid.*) The great advantage of such interactive systems is the ability to move around the models to obtain a large number of observations in a very short period of time. But while these tools represent a considerable advance on previous methods for examining the recorded data, one disadvantage is that as the number of contexts in the model increases so it becomes more difficult to isolate individual wire-frame objects and their possible relationship to neighbouring contexts. In particular, as the number of vectors increases, the picture produced tends to become more confusing. This became particularly problematic when we attempted to explore the relationship between the objects found in the real archaeological features such as post-holes and those recovered in the arbitrarily defined contexts in the deep stratification overlaying them. Attempts to combined the digitised plans of archaeological contexts with the spit-context data were unsatisfactory. Although a degree of comparison could be achieved when using plan views, it proved impossible to relate the two sets of data from any oblique view due to problems caused by parallax.

17.3 Solid modelling

Fortunately, much more powerful solid modelling systems offering superior data visualisation capability have been and continue to be developed.

Solid modelling systems should not be confused with other graphics systems which produce pictures of solid looking objects. For example, a common practice is to construct something called a face-model from a set of rendered polygonal panels. Although the object has the appearance of being solid, it does not actually conform to a truly enclosed solid object. Solid modelling systems were developed mainly by engineers to perform analysis functions on the resulting models. This criterion excluded the use of face-models since the engineers required a process intrinsically geared to the production of representations of solids rather than relying on the user to create the required set of faces.

Most solid models are based on one of two data-structures, and often both are present (e.g., Woodwark 1986). The *boundary* model is similar to the face-model, except that the faces, edges, and vertices of the model are linked together into a structure which is assured in its topological consistency. That is to say, there are no extra or omitted faces, edges, or vertices of the object. The *set-theoretic*, or constructive solid geometry model, on the other hand, is defined as the combination of simple, or primitive, solids using operators derived from set theory.

Once the computer model has been constructed it is possible to generate a multitude of views with correct perspective and accurately modelled lighting, including highlights, shadows, and reflections. Such sophisticated image synthesis is extremely CPU intensive. Despite this, although boundary models have not yet been used much in archaeology set-theoretic models have.

So far these solid modelling systems have only made an impact in reconstructions of former building complexes, where, for example, they have stimulated the production of a number of interesting ideas concerning the use of space (Reilly 1988, Eiteljorg II 1988). For example, analysis of views of the solid model reconstruction of the Temple Precinct from Roman Bath convinced Professor Barry Cunliffe that the temple precinct builders had made use of the space to emphasise the symbolic relationships between the various elements of the precinct. About 140 views of a solid model of the Roman legionary bathhouse at Caerleon were created to help people understand the scale and organisation of this structure. This idea of creating multiple views to illustrate the monument was taken to its logical conclusion when a fully animated tour of a computer model of the Old Saxon Minster of Winchester was produced at the IBM UK Scientific Centre. The two-minute version of the movie that was exhibited at the British Museum's 1987 *Archaeology in Britain* exhibition required many hundreds of views to be computed and involved a considerable investment in central processor time.

17.4 Animating solid models of St.Veit-Klinglberg data

These techniques are now being developed further to help in the analysis of the Klinglberg excavations, where the archaeological contexts are being modelled using the WINchester SOlid Modeller, WINSOM (e.g., Fig 17.3). Each of the Klinglberg box contexts is actually a tetrahedroid, which is defined in WINSOM using a series of planar half-spaces. A database system is used to retrieve the coordinates of vertices of the required contexts together with the value of some property of interest (e.g., absence/presence of pottery type x). These eight triads of coordinates are read into a WINSOM program, in a fixed order, where they are used to define the locations of the vertices of the twelve triangular faces. The intersection of the planar half-spaces passing through these twelve triangular faces defines the solid geometry of each box context. By this method different solid models were constructed to show the distribution of box contexts containing both copper objects and bronze age pottery, bronze age pottery only, copper objects only and all contexts.

17.5 First video sequence

A series of tours around these models were then created by computing a sequence of views of the solid models taken at incremental steps along a chosen view path and transferring them to video tape as individual frames. The view path used in this animation was that of a simple circle placed above the models, so that the viewer would look down towards the centre of models from an angle of about 30°. In fact, the illusion of this circular perambulation around the model was achieved by keeping the viewing parameters constant and rotating the model. Parameters providing for suitable ambient lighting and natural perspective were supplied in the model definition, but no other special attributes, such as texture, were requested.

Since these animations were intended to help in spatial analyses, they had to be slow enough for the viewer to *take in* what was being shown, rather than appearing as a meaningless blur of action. On the other hand, the tour shouldn't be too slow either as this will incur the cost of producing unnecessary extra frames. Experience has shown us that a comfortable circular tour should last about 25 seconds for relatively uncomplicated models such as the Klinglberg box contexts. Intricate sculptured forms, however, may need much longer to fully appreciate the finer details. Since the frame rate of a video is fixed at 24 frames per second, 600 frames are required for a twenty-five second circular tour, with the increment between each step on the tour calculated as the total number of degrees of revolution divided by the number of frames (i.e., 360/600 = 0.6°). Considerable savings in production costs and time could have been made by halving the number of computer pictures but placing them down as two repeated frames of video. This was unacceptable because the interval between increments would then have been 1.2°, which would have had the effect of producing lurching movements around the model. In contrast, with the smaller interval the viewer perceives a smooth gliding motion. Finally, the several tours around the different models mentioned above were concatenated to produce the final animation.

So what did this animation reveal? While this is not the right place to offer a detailed interpretation, which would in any case be premature at this stage, because other related analyses which which have a direct bearing on this matter are planned, some provisional observations can be made. Very briefly, the animation indicates that the copper objects and bronze age pottery may be separated into distinct spreads of material down the slopes, which suggests a degree of layering, but that the two sets of material are much more mixed-up in the deeper deposits which lie on the level plateau at the bottom of the slopes.

17.6 Second video sequence

Another animation was designed to help investigate whether or not the distribution of material in the large layers reflected what was found in the archaeological features underlying them. Once again, the variables under examination were extracted directly, or were calculated, from the data base on a feature basis (e.g., the cuts of post-holes and pits). Like most modern excavations the features at Klinglberg are recorded to scale in plan and profile only. However, the volume of material extracted from each feature is also recorded, and photographs are also taken. This information is not sufficient to enable an accurate three-dimensional model of the feature to be built, but it can be used to investigate our question.

The method we have devised is simple in conception, if somewhat more complicated to implement. Each planned outline of a cut in the area of interest is digitised, and these digitised outlines are then extruded to form prisms. Extrusion is achieved by passing a plain through each adjacent pair of digitised points on the outline. This is not a trivial task, posing as it does some interesting problems when re-entrant features are encountered. The colour of each prism will be dependent on the property under study (e.g., average sherd size for example). These prisms are then unioned with the modelled box contexts from the area of interest, which are also shaded according to the same set of colour conventions (Fig. 17.1).

Slices of the excavation area may then be progressively removed by running a clipping plane across the model of the excavated area, thereby enabling the researcher to see how the distribution of properties in the features compares with those in the larger layers overlying them (Fig. 17.2). This approach was used in the second animation which was produced to help us examine how the size of the bronze age pottery sherds varied between the layers and features in an area in the centre of the main excavations. In this instance, no obvious patterning was observed.

17.7 Recording pictures onto videotape

Transferring the computer-generated scenes to videotape requires some specialist hardware and software. After computing our images, a simple program is used to specify what pictures to record, for how long, where on video-tape they must go, and so on. This information and the pictures are then passed on to a host-attached IBM personal computer which controls a Spaceward Supernova Framestore/Animation Controller. The Supernova has 24-bit/pixel capability, and a resolution of 768 × 576 pixels.[1]

Although in this pilot study we used it in 8 bits/pixel mode, in future we hope to exploit the higher quality offered by the 24-bit options in both WINSOM and the Supernova now that we are confident with the method. Spaceward software on the PC controls a recorder attached to the framestore. Currently, we are using a Sony Betacam SP recorder for high-quality work, which may then be edited professionally, copied to Umatic/VHS, or used to master a video disk, but other recorders may be used if wanted. Pictures are converted into Supernova-format by a host program, down-loaded to the PC using an IBM internal use only terminal-emulator package called MYTE, in blocks of up to 10 images at a time. Due mainly to the slow download speed to the PC, it takes about one minute to record one picture onto tape, so it is a slow process, suitable for running overnight.

17.8 Outlook

The two short video sequences we produced last for about three-and-a-half minutes, but required in the order of 5000 1Mb images to be computed. This is obviously a high price to pay at the moment. However, the analytical advantage of using solid modelling techniques will be vastly improved once real-time interaction is possible, and it is no longer necessary to plan or choreograph tours around the data. We want to manipulate such models in real-time so that we can pursue observations immediately. Moreover, if the hardware can handle the necessary transformations for interaction with these models, there will be no need to generate and store thousands of separate views of the same object.

Unfortunately, there is a limit to the speed of operation which is inherent in the use of current hardware technologies. Real-time interaction with fully rendered solid models, requiring sub-second response times, cannot be achieved with the processing power available in our present workstations. Even with the use of supercomputers, there is still a shortfall in processing speed. One approach to solving the problem of providing sufficient processing power is to link together a network of relatively simple processors such that each one can work, concurrently with the others, on one part of the total task to be done.

IBM UK Scientific Centre has a research group directed to investigate ways of using a network of processors to implement image synthesis algorithms. As a vehicle for

[1] The resolution of the computed scene is 1024 × 1024, or 1Mb. The object of interest in the scene has to be positioned and scaled in such a way as to fill the video display, which involves truncating the initial picture (Fig. 17.3).

this research the team has chosen the INMOS^{TM} transputer as the processor building block, and Occam^{TM} as the parallel-processing language. For example, Fig. 17.4 is a snap-shot showing a lighting editor at work. The model on the left has been strongly lit from the right-hand side. The image being synthesised on the left is the same model, but instead of bright directed side illumination the overall ambient lighting is being increased.

The whole process can be achieved in a matter of seconds, and improvements continue to be made. As each processor is concerned with rendering only a small part of the model at any given time, the synthesised image has a patch-work quality about it until the entire image is finished. The histogram in the bottom right hand side indicates the current work load of each processor as the new image is being synthesised. When a processor becomes available it is assigned another task.

17.9 Conclusions

The exercise of making such models has several benefits to the archaeological analyst. Apart from their aesthetic appeal and the intellectual curiosity they provoke about the technology employed in their construction, they encourage the user to think more deeply about how the various archaeological components relate to one another. In developing such a model one becomes more familiar with the particular properties of the data. When modelling archaeological formations one becomes acutely aware of the deficiencies in the data upon which the model is to be based. The modeller must consider and precisely define each individual component used in the reconstruction explicitly but, even more importantly, modelling encourages structural analysis of how the components relate to one another. Frequently the most simple insights help the archaeologist to isolate and to examine more closely those areas where evidence lurks unnoticed to resolve particular problems. The archaeologist can then think about other ways of testing the theory or confirming the observation.

There are also major implications for archaeological recording methods. In the past, field-workers were justified in employing the traditional techniques of drawing plans and profiles of excavated features, since three-dimensional recording was prohibitively expensive and there were severe limits to how such data could be manipulated. Now a wide range of three-dimensional recording instruments with computer loggers are available. Moreover, computer systems already allow us to handle enormous sets of data. Now we are beginning to model and manipulate complex shapes in practicable time spans. In other words we are fast approaching the reality of creating virtual excavations which embody the full three-dimensional properties of features and artefact/ecofact distributions. We can therefore expect major advances to occur in the study of formation processes, their effects on the archaeological record, and consequently our interpretations of this material. As custodians of our archaeological heritage, many excavators feel obliged to preserve the most accurate record possible of the maximum amount of data: data they invariably destroy in the course of their work. The implication of all this is that archaeologists will have to change their recording procedures to make use of the full capabilities of current information processing technology. Clearly, such modelling also has a great role to play in presenting archaeology to the public.

Bibliography

EITELJORG II, H. 1988. *Computer-Assisted Drafting and Design: New Technologies for Old Problems.* Center for the Study of Architecture, Bryn Mawr, PA, USA.

REILLY, P. 1988. "Data Visualisation: Recent Advances in the Application of Graphic Systems to Archaeology", Report 185, IBM UK Scientific Centre, Winchester.

REILLY, P., A. M. LOCKER, & S. J. SHENNAN 1988. "Pattern Recognition in Sub-Surface Artefact Distributions". *in* Rahtz, S. P. Q., (ed.), *Computer and Quantitative Methods in Archaeology 1988*, International Series 446, pp. 265–94. British Archaeological Reports, Oxford.

WOODWARK, J. R. 1986. *Computing Shape.* Butterworths, London.