

## Computer-aided design techniques for the graphical modelling of data from the prehistoric site of Runnymede, Berkshire

P.L.Main<sup>1</sup>, A.J.Spence<sup>2</sup> and T.Higgins<sup>1</sup>

<sup>1</sup> Department of Scientific Research, The British Museum, LONDON WC1B 3DG, UK

<sup>2</sup> Department of Prehistoric and Romano-British Antiquities, The British Museum, LONDON WC1B 3DG, UK

### 37.1 The archaeological background

The prehistoric site at Runnymede Bridge, Egham, is situated on a former island at the confluence of the Colnebrook and an ox-bow of the Thames (see Figure 37.1). The alluvial sequence has shown that the latter channel silted up in early post-medieval times, and the modern river channel formalised by the construction of Egham Lock some 150m upstream has truncated the northern edge of the site (Needham and Macklin 1992).

Archaeological material was first uncovered in 1959–1960 during the construction of the A30 bridge. However, proper investigation did not begin until 1975 when Barker and Longley located *in situ* material in the side of a recently scoured flood-way on the east side of the road. Early the following year a small excavation on the berm between this cut and the embankment produced evidence of Late Bronze Age settlement activity, including post-hole structures (Longley 1980).

In 1978 intermittent excavations on the site of the M25

bridge revealed timber piles lining a river bank and hard standings for river craft, indicating a well organised settlement of the early first millennium bc. Beneath the Late Bronze Age (LBA) sequence, evidence of the Middle Neolithic settlement was recovered with pottery, animal bone and more timbers bearing stone axe impressions. This material has all been published in the first of the British Museum Runnymede reports (Needham 1991).

Further attention was paid to the site in the 1980s. At the start of the decade two small trenches at the foot of the west side of the A30 embankment (Areas 9 and 10) revealed more *in situ* LBA and Neolithic material. The continuation of these deposits this side of the roadway led to the British Museum research excavations from 1984–1989 (see Figure 37.2). Initially a series of trenches determined the presence of archaeology over much of this area, and two zones were then examined by open excavation on a larger scale. On the north side the LBA deposits were uncovered, with significant quantities of settlement debris and many cut features on the top of a silted river channel bank and levee behind. On the

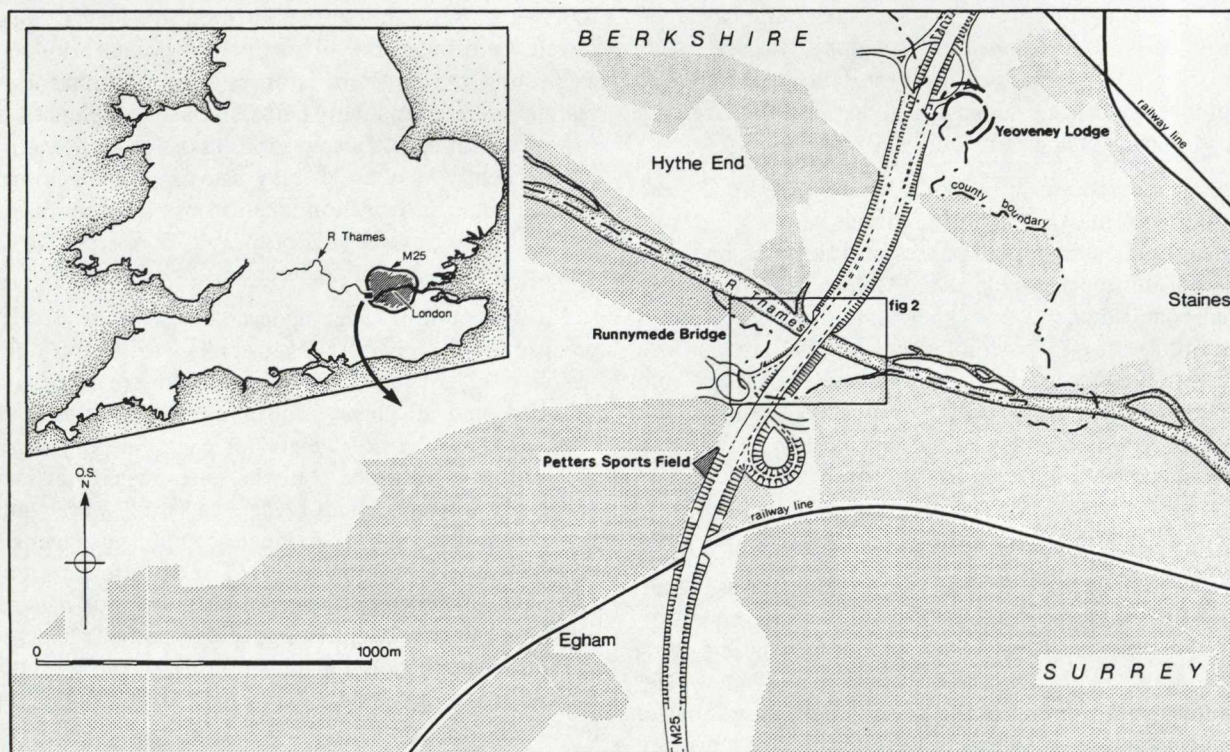


Figure 37.1: Location of the Runnymede site (Illus. S.Crummy).



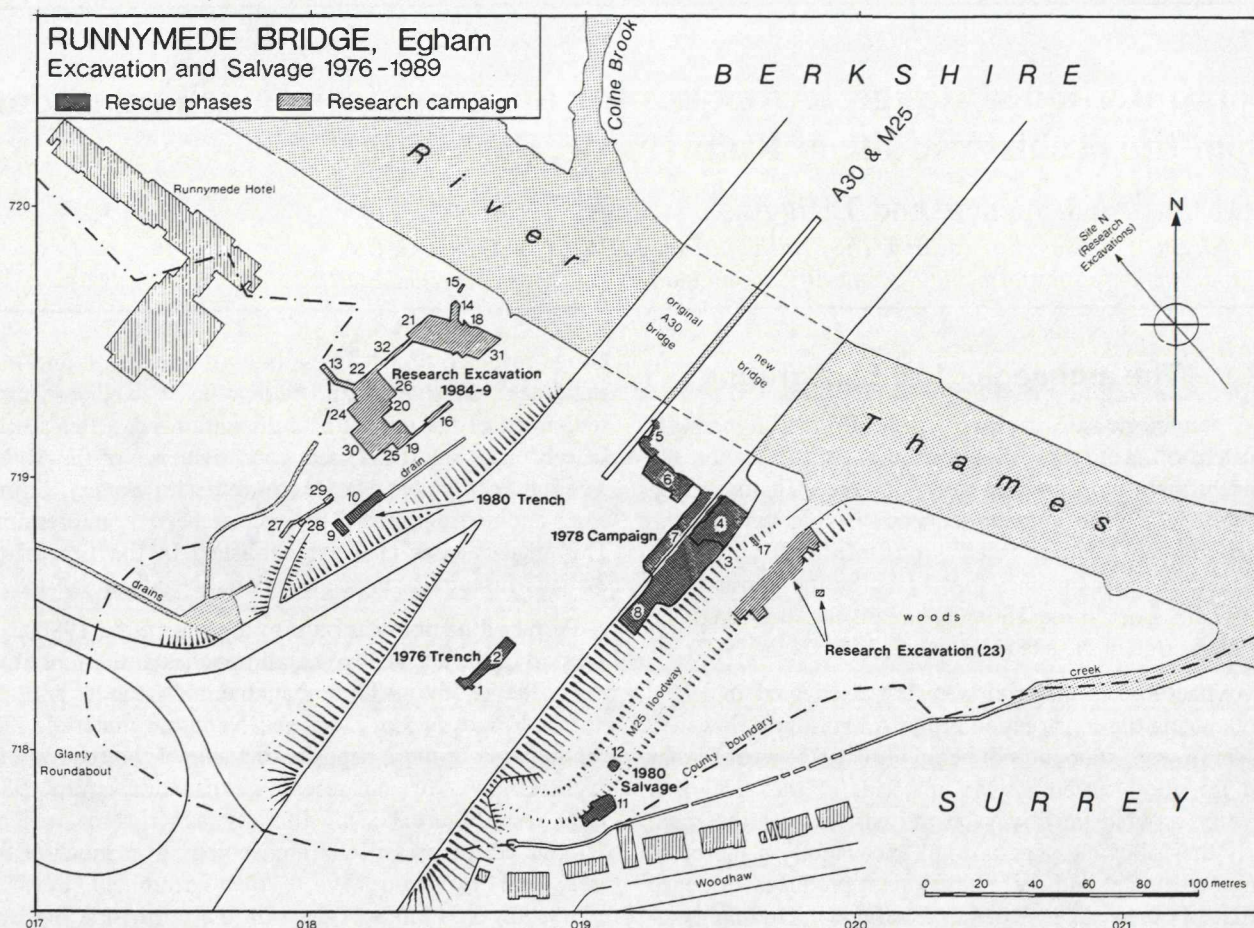


Figure 37.2: Areas of excavation and salvage at Runnymede during 1976-1989 (Illus. S.Crummy).

central-west side of the island the Neolithic deposits, dating to approximately 3700–3400 cal BC, were on average only 0.3m beneath the Late Bronze Age. Radiocarbon determinations from the latter, supported by artefactual evidence, suggest activity lasting about a century starting at approximately 900 cal BC.

The research excavations have reinforced the importance of the site. The inundating silts and gravels that marked the end of both phases of occupation had only disturbed the uppermost layers. Many of the potsherds, animal bone fragments and other finds therefore relate to the multitude of pits, post-holes and original land surface features of the respective phases. Spatial analyses of both finds and features are potentially far more revealing here than on sites where centuries of dislocation by plough and natural agencies have occurred.

### 37.2 Nature of the data collected on site

The research excavation recording system was based around the single context, with a few enhancements to assist with the large quantities of artefactual and soil-related data. Although the excavation had no on-site computer facilities, even in the field it was realised that in post-excavation (levels III and IV) this would be a priority and the recording system has since proved readily adaptable to computerisation.

The basic excavation system was by trench, subdivided by layer or spit, and further divided into metre squares. Levels were taken on all metre grid points prior to the excavation of any soil layer or spit. Archaeological features were all routinely half-sectioned and planned both before and after full excavation, and the section strings independently levelled, thereby providing a check on the gridded data. Information taken from plans, profiles and context sheets gave details of the position, type, diameter and depth of individual features.

Finds recording was undertaken at three levels of precision. The coarsest level applied to the bulk finds, which were attributed to their layer metre squares and recorded on 1:20 plans. The number of fragments and, more usefully, the dry weight for each category of finds material were recorded. In the cases where distinctive clusters of artefacts could be detected these were planned in at 1:5 and the individual components numbered consecutively. These are referred to as 'artefact groups', and have also been included in the bulk finds databases. The final category comprised special (or small) finds, such as bronzes, bone tools and amber beads, which were three-dimensionally recorded *in situ* whenever possible, and have been recorded in separate 'special finds' databases.

The amount of data generated by the above regime of recording is of course huge, and computerisation of the





Figure 37.3: Exposed cut features from Area 19 at Runnymede (Photo. S.Needham).

datasets was seen as a very necessary process in post-excavation. Four types of database are now being produced, designed in part to allow computer graphics reconstruction, and have been taken directly from the sets of data recorded manually on site.

The first type relates to grid levels data, and contains the 3D coordinates and the upper and lower contexts at each metre point. By using four corner points this effectively fixes a block of soil in 3D space, and also enables a measurement of soil volume to be calculated for any given metre square, and, by amalgamation, for layers.

Bulk finds databases have now been completed and stand at 12,000 records for squares and feature fills. As well as simple analysis by weight and number of fragments, further statistical analysis linking soil volume to finds will be carried out, to permit study of deposition rates and promote a better characterisation of disparate soil bodies. Changes both horizontally and vertically can then be examined, using distributions represented as dot density plots.

The third type of finds-related database is for special finds. As with the bulk finds, the primary purpose of computerisation is to enable speedier manipulation of the data, with the added benefit of providing core data for the British Museum's computer-based registration records. However, as post-excavation progresses, special finds have

also become an important part of the spatial analysis programme. As many were recorded *in situ* with three dimensional coordinates they are readily amenable to graphical representation. They have greater precision than the bulk finds records, and should serve to confirm the detail of any observed pattern in bulk finds distribution, although it is important to be aware that there may be different factors governing the deposition of the two finds categories.

The final database type is concerned with cut features, where a slightly different problem presents itself. Runnymede has produced an abundance of pits, post- and stake-holes, as well as the occasional natural feature, all of which have tended to mask the outline of any particular structure (see Figure 37.3). For example, the Neolithic layers in Area 19 have about 350 features covering 56 square metres. In confirming wall lines and boundaries it will be useful to filter the data on the basis of fields such as the layer cut into, to ensure contemporaneity, or the category of feature, which may be of architectural significance.

The importance of Runnymede lies in the depth and *in situ* nature of the deposits. Spatial analysis of the finds and features is the key to understanding how the site functioned. Where the buildings are, what happened within and between them, and how material came to be



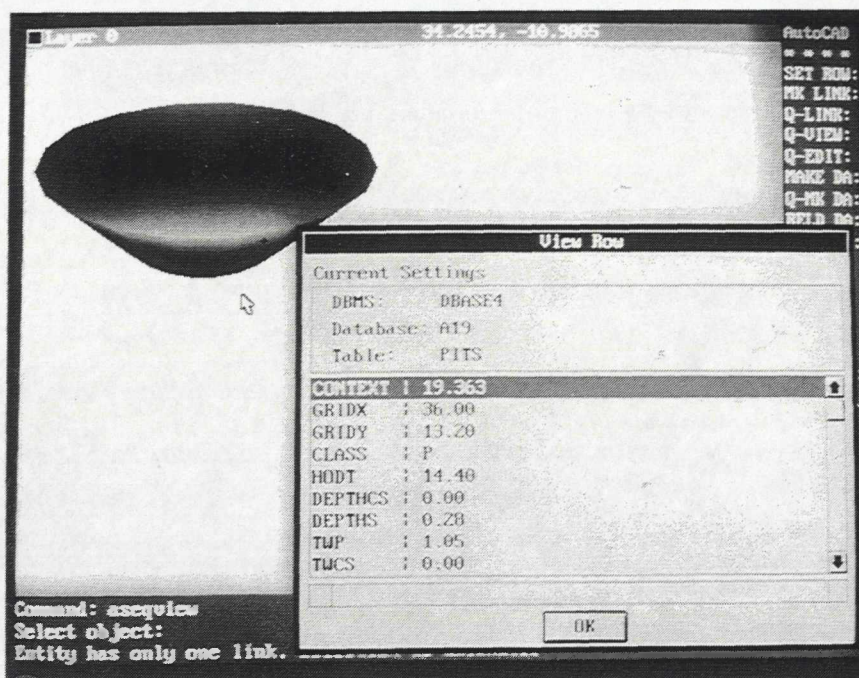


Figure 37.4: A modelled cut feature and its associated SQL data.

deposited, are all questions that need to be addressed. The site plans can provide some of this information, but computer generated reconstructions offer the possibility of comparing and contrasting datasets more effectively and, crucially, in three dimensions.

### 37.3 The software environment

The requirements of the previous section can be summarised as software for: validating and sorting sets of three-dimensional coordinates; calculating soil volumes; and carrying out schematic (as opposed to visually realistic) geometric reconstructions of features, ground surfaces and artefact distributions, providing visual aids to interpretation. Many of the calculations and reconstructive geometries needed are not sufficiently general in nature to be immediately available in off-the-shelf software; a programming environment was therefore required. It was felt to be highly advantageous if the software could directly access existing data held within dBASE IV files.

AutoCAD Version 12 was chosen for the following reasons:

- It was already available at the British Museum and had already proved its worth in an earlier project (Main *et al.* 1994).
- It has powerful two- and three-dimensional graphical modelling capabilities.
- The ADS (AutoCAD Development System) module provides a C programming environment.
- The ASE (AutoCAD SQL Extension) module, which is new with Release 12, provides embedded SQL routines and drivers to link with a number of common

database management systems, including dBASE IV. If need arose in the future, the software written for the Runnymede project could be used for data held within other DBMSs, such as Paradox, by altering only one line of code within each application.

Applications were written in C to read data from appropriate dBASE IV files, and call routines from the ADS library to generate corresponding graphical entities. This required the dBASE IV files to be converted to SQL tables, making the data accessible both to the standard dBASE IV menus, and to C routines running under AutoCAD. This conversion is straightforward using a utility that is well documented in the dBASE IV manuals. Although the SQL link has been used mainly in read-only mode to import filtered and sorted data into AutoCAD, data can also be written back to the tables under software control. This facility was used to write soil volumes information to the bulk finds database, as described in Section 37.4 below.

A particularly valuable feature of the ASE software is that it is possible to incorporate within the drawing file itself information which links each graphical entity back to the row of the SQL table used to generate it. By picking an entity with the mouse, a window can be overlaid on the drawing which allows the user to view and modify the data within that row (see Figure 37.4). Data maintenance can thus be integrated with graphical modelling, entirely within the AutoCAD environment.

The applications have been designed to allow the user to specify interactively the filter conditions for retrieving the data. In this way, any subset of the data can be represented graphically, and different subsets overlaid using a variety of colours or symbols. As this process



takes place under the control of a C program, it is straightforward to cope intelligently with exception conditions such as those arising from missing data or other idiosyncrasies of the data recording method.

All the applications described in later sections have a similar preamble where the user is asked to provide the following information:

- The name of the SQL database containing the data. The data have been subdivided according to the excavation areas distinguished on-site. All data tables contain coordinates relative to the overall site grid, so that it will eventually be possible to generate reconstructions of the complete site, either by amalgamating corresponding tables from the various databases, or by modifying the software to retrieve data from multiple databases.
- The name of the SQL table containing the data. Within each database there are four types of table, having structures appropriate for grid levels, cut features, special finds and bulk finds.
- A selection condition for retrieving data from the table. This is entered by the user as the WHERE clause of an SQL SELECT sentence. For example, the clause

`gridx > 30 and gridx < 40 and cuts = '16.501'`

applied to a cut features table would instruct the application to retrieve only data relating to a strip between 30 and 40 metres from the site origin, and where the features cut through context '16.501'. The complete SELECT sentence is constructed by the application, by prepending the list of table columns required to generate the reconstruction, and, in some cases, by appending an ORDER BY clause. (Pre-sorted data are necessary, for example, when checking grid level sequences as described in the following section.)

More detailed descriptions of the specific applications written to cope with the different classes of data appear in the following sections.

## 37.4 The grid levels application

This application comprises four principal routines.

### 37.4.1 Data retrieval

Each row of a grid levels table corresponds to one surveyed point in 3D space. Seven columns are retrieved from each selected row, representing the X, Y and Z coordinates of the surveyed point, the numbers of the contexts lying above and below the point, and the names of the stratigraphic phases lying above and below. The user is given the choice of discarding context information and dealing only with stratigraphic phases, in which case the application amalgamates the data from those contexts that are adjacent and lie within the same stratigraphic phase.

The data are accumulated in memory as linked lists of C structures, with pointers assigned to connect each survey point with its previous and following point in the X, Y and Z directions. This structure allows the data to be accessed either as vertical 'pipes' of survey points (appropriate for the data checking routine described below), or as horizontal spreads of points which make it easy for the display routine to reconstruct the surface of any stratigraphic unit.

### 37.4.2 Data checking

This routine supplies the user with tabulated details of each survey point, flagged where necessary with warnings or errors. These may relate, for example, to breaks or reversals in the sequence of stratigraphic phases. It is most important that inconsistencies in grid levels data are corrected before the data are used to reconstruct surfaces, otherwise some bizarre geometric anomalies may be evident. A feedback loop has therefore been introduced at this point, whereby the software flags problems, the archaeologist corrects the raw data, and the data are checked again. This is repeated as often as required until 'clean' data results.


### 37.4.3 Volume calculations

The grid levels tables contain data which define the upper and lower surfaces of each stratigraphic unit at all points on a regular one metre square grid. The three dimensional extent of the stratigraphic unit can therefore be regarded as comprising a spread of volume units which have square cross-section in the horizontal plane, and flat but non-parallel upper and lower surfaces. This is in fact an approximation to the truth, since the four points on the top or bottom will not be strictly coplanar. In practice, however, they will be nearly so since surface contours do not exhibit sharp local variations within a square metre. In geometric terms these units of volume are parallelepipeds, whose volumes can readily be calculated. The level of positional recording employed for bulk finds was, in effect, sufficient to tie each find to its enclosing parallelepiped, and aggregate weight figures for each class of bulk find were therefore known for each parallelepiped. These weights, together with the volumes calculated by this routine, allow finds densities to be calculated for use by the bulk finds dot-density display routine described in Section 37.6 below.

### 37.4.4 Surface reconstruction

This routine allows the user to select a stratigraphic unit and to reconstruct its ground surface graphically. This is done by taking each set of four adjacent survey points in turn and generating a Coons surface patch which interpolates the points. A Coons patch (Coons 1967) is a bicubic surface which interpolates any four edge-connected 3D curves, and is implemented as AutoCAD's EDGESURF command. Where only three adjacent points exist, as may happen at the edge of a stratigraphic unit, a triangular polyface mesh is generated instead with the RULESURF command.





**Figure 37.5:** Schematic representation of the technique for modelling cut features.

Each complete reconstructed stratigraphic unit is assigned to a separate AutoCAD 'layer' in the drawing, so that each can be switched on or off independently, or coloured differently. Since the surfaces generated are 3D polygon meshes, they will 'hide' objects they overlap if hidden line removal is applied. This gives added clarity and realism to 3D views of the surfaces. It is also possible to render the surfaces under chosen lighting conditions using AutoCAD's rendering facilities.

## 37.5 The cut features and special finds applications

### 37.5.1 Cut features generation

This routine models cut features (typically post-holes or pits) using measurements of position and size stored within a cut features SQL table. The measurements utilised to reconstruct the features are the X, Y and Z coordinates of the centre of the top of the feature, the top and bottom widths of the feature ( $w_1$  and  $w_2$ ), and its depth ( $d$ ). The measurements are those routinely taken on site and are sufficient to allow a simple stylised reconstruction in the form of a truncated cone.

Following the user's specification of database, table and selection conditions, this routine generates one truncated cone, correctly located in 3D space, for each row retrieved from the table. This is achieved by drawing a line corresponding to the central vertical axis of the feature, a polyline whose locus is the top, edge and bottom of the feature, and calling AutoCAD's REVSURF routine to generate a solid of revolution by rotating the polyline around the central axis (see Figure 37.5). The original line and polyline are then deleted from the drawing, and a link created between the solid of revolution and the row of the SQL table which was used to generate it.

The solids of revolution generated by this application are a particular type of polyface mesh, and as with reconstructed ground surfaces, can therefore have hidden line removal or rendering applied to give more realistic 3D views. Figure 37.6 shows an AutoCAD reconstruction of exposed pits and post-holes from Area 19.

### 37.5.2 Volume calculation

This routine calculates the volume of all features selected by the user, using a formula based on the volume of a cone, and writes the results back to a specified column within the table.

### 37.5.3 Special finds generation

This routine accesses a special finds SQL table and draws markers to represent the positions of the finds in 3D space. Each selection from the table can be associated with a particular type of marker, chosen from a file of visually distinctive and customisable symbols. Markers are drawn with AutoCAD's SHAPE command, which also allows markers to be drawn in different sizes if required. Further refinements to the display can be applied by using the layer allocation routine described below to assign different colours and layers to the marker groups.

### 37.5.4 Layer allocation

This routine accesses an existing AutoCAD drawing which already contains modelled cut features and/or special finds markers. It allows the user to select subsets of the entities and assign them to different AutoCAD layers, choosing an appropriate colour for each layer. The user can select the members of each subset manually (i.e. by pointing at the features and clicking the mouse) or by executing an SQL selection sentence on the table to which the entities are linked, or by a combination of both methods. Where manual selection is used, the user may optionally connect the centres of selected cut features with straight lines to delineate, for example, the edge of a putative structure. The alternative method of selection is potentially powerful since the selection condition can involve any data contained within the table, not only the measurements used to generate the reconstruction. Provided appropriate data have been recorded in the table, one could, for example, assign all post-holes that have been cut into a specified stratigraphic context to a single AutoCAD layer.

## 37.6 The bulk finds application

This application comprises two principal routines, the first of which draws a 3D grid of lines marking the boundaries of each volume unit with metre-square cross-section, and the second of which populates each volume unit with a random spread of AutoCAD POINT entities to represent a distribution of bulk finds.

### 37.6.1 Drawing the grid

The data utilised from the grid levels and bulk finds tables comprise the 3D coordinates of the vertices of each volume unit, and the density, within each unit, of each






Figure 37.6: AutoCAD reconstruction of selected cut features from Area 19 at Runnymede.

bulk finds type as derived from the volume calculations within the grid levels application. The grid is constructed by drawing lines between the corners giving the impression of truncated cubes. Since the context surfaces involved are relatively flat, provision has been made for the user to specify an exaggeration factor to enhance non-obvious features.

### 37.6.2 Populating the grid

The number of points plotted within a volume unit is related to the density of the finds type. Both linear and logarithmic mappings are provided to allow experimentation with the weighting of the high and low densities. An upper limit of 300 points per volume unit is, however, imposed so that AutoCAD is not overstretched during the drawing and regeneration procedures. Figure 37.7 shows part of the grid from Area 16 populated with points.

The X and Y coordinates of each point are determined using randomly generated numbers between zero and one as offsets from the bottom left-hand corner of each context. Each point thus has coordinates  $(x_1+dx, y_1+dy)$  where  $dx$  and  $dy$  are the offsets. Calculating the Z coordinate is less straightforward. Although the limits to the X and Y coordinates are determined by the geometry of a square grid, if we include the Z coordinate we find that the corners are not coplanar. There are no obvious single upper or lower values that can be used as the Z coordinates at interior points of the volume units. If we assume that the lowest Z value is the Z coordinate of the lowest corner and similarly for the highest Z value then some points will lie outside the grid when plotted. If there

is another context immediately above or below the current context then there will be space common to two contexts which will be doubly populated. This will lead to 'false features' being displayed as bands of dense points.

To overcome this problem a 'minimum' surface passing through the corners is calculated. Thus, for a given  $dx$  and  $dy$  lying between zero and one, the Z coordinate is calculated as:

$$z = (1-dx)(1-dy)z_1 + dx(1-dy)z_2 + (1-dx)dyz_3 + dx dy z_4$$

This value is calculated for the current context, and for that immediately overlying it, with the final  $z$  value chosen to lie between these two. Although this procedure makes a particular assumption about the behaviour of contours between grid points, some assumption has to be made, and this one does ensure that the points populating adjacent volume units cannot overlap. Figure 37.8 shows a sectional dot-density view of Area 16 with clearly visible zoning of bulk finds.

## 37.7 Software assessment

Certain of the AutoCAD facilities we have exploited in developing these applications have proved to have shortcomings, or have been less easy to utilise than we would have wished. When considering these, however, it should be borne in mind that AutoCAD is a powerful, broadly based package designed for the wider needs of computer-aided design and was not specifically designed to cater for the often specialised and demanding requirements of archaeological graphics. It is a testimony to its power and flexibility that we have been able to achieve all that we have with very moderately sized






Figure 37.7: Three-dimensional dot-density representation of bulk finds from part of Area 16, overlaid with context grid.




Figure 37.8: Sectional view of Area 16 with clearly visible zoning of bulk finds.

segments of C program. Most of the lower level graphical routines were already available within the package; much of the code we have written has been to call these up in the correct sequence and to accommodate the particular data categories and coding employed on the Runnymede excavation.

The need to convert dBASE IV files to SQL tables to render them accessible to AutoCAD has been something of an irritation. Once these have been converted, the data can still be accessed and modified using dBASE's ASSIST menu system, but certain operations, such as modifying the file structure in place, become impossible. Furthermore, any dBASE 'memo' fields need to be deleted before the SQL tables are generated, otherwise any attempt to access the SQL table from AutoCAD fails with a 'corrupt table' error.

The ability to create links between graphical entities and rows of an SQL table has been of great value to us, particularly within the cut features and special finds applications. AutoCAD's ASE interface, however, has clearly been designed around creating these links interactively using 'point and click' methodology. It was essential for our purposes to be able to create the links programmatically, at the same time as the graphical entities were created. It has proved possible to do this using a combination of trial and error and intelligent guesswork to call the linking routine in a way that is not fully documented. The process of creating links in this way seems relatively slow, and it can take many minutes to generate a cut features drawing of any complexity. It is to be hoped that programmatic generation of entity/SQL

table links will be fully documented and supported in future releases of AutoCAD.

The use of AutoCAD 'shapes' to represent special find locations is not ideal. The result is only satisfactory in plan views, since shapes are two-dimensional and lie in the X-Y plane. When distributions of artefacts are viewed in section, they are visible only as lines and are effectively indistinguishable except by colour. Furthermore, being two-dimensional, they cannot be rendered and disappear altogether if rendering is applied. Ultimately we plan to develop a range of three-dimensional solids to use as icons for special finds, although colour will become the major discriminator since different types of three-dimensional solid are notoriously difficult to distinguish at small scales.

Although standard AutoCAD facilities allow only static views to be created, we see great potential for the user being able to interactively control his viewing position and 'walk' around and into rendered images. We have been able to achieve this with additional hardware and software, namely a Matrox MGA Impression graphics card and the supplied Dynaview driver for AutoCAD. Such facilities are not merely luxuries; they are important for examining the relationship between features, artefacts and their stratigraphic context. The 'right' view, which illuminates some aspect of the data, may be very difficult to find without exploring the data in this way.

Finally, it should be pointed out that we have not yet been in a position to fully exploit the software we have written. We have been working primarily with a small



core of test (though real) data relating to an LBA midden deposit in Area 16 (Needham and Spence, forthcoming), whose choice has been dictated largely by publication priorities. Although we also have cut features data for Area 19, we have not yet been able to generate composite reconstructions of all types of data on a single drawing. When this becomes possible, further modifications and improvements to the software will no doubt suggest themselves.

### 37.8 Archaeological assessment

The aims of graphical computerisation for the Runnymede project are to provide ways of representing bulk find, special find and cut feature information which aid visual interpretation. Initially much of the analysis can be carried out using sections and plans, but there are times when three-dimensional studies are essential. Additional benefit has been gained by the ability to study finds as densities as well as aggregated weights.

Some problems remain, however. The difficulties in accurately determining soil volumes have already been outlined, and these apply particularly to the smaller soil bodies. To a certain extent errors can be reduced by amalgamating contexts into stratigraphic phases, although this will then mask detail. In analyzing the Area 16 East midden deposits the individual spit information was gathered into phases to spot long term trends, but local variation was detected from studies of individual categories. It is also inevitable that cut features from more than one spit will be used when searching for alignments, as during excavation there were problems at times locating the 'real' tops of the features, resulting from the similarity of the fills to the surrounding soil matrix.

In representing the data graphically, there is also an important consideration to be taken into account. The cut feature data can be accurately represented in three dimensional space, even if only stylistically. Bulk finds data can only show distributions schematically within each metre square, while special finds fall between these two categories. For those special finds identified immediately on excavation, precise three-dimensional pinpointing is possible, whereas those recovered in finds processing can only be shown as being located 'somewhere' within their metre square spit. Given that the interpretation of much of the site detail will require very careful analysis of distribution, a distinction will need to be drawn between the precise and the approximate plots. It is probable that the computer graphics will point the way towards significant boundaries and distributions, but the final conclusions will need to be confirmed from the site plans, which contain detail not available within the computer records.

There is a temptation to regard the excavated data as objectively recovered, whereas perceived spatial patterning may merely reflect the excavation method, or even the abilities of the individual excavators. On-site practices such as sieving 'control columns' to recover all finds from certain squares, or using different teams on neighbouring trenches helps to reduce this human factor. Confidence in the interpretation can be increased by checking results against another dataset. Given that most archaeological processes are far from straightforward, it is the ability to try many different permutations of data that makes computerisation worthwhile. As a spin-off, the archaeologist becomes more explicit in formulating hypotheses; in asking these questions the limitations of the data often become apparent.

The use of the AutoCAD system will also allow the comparison of Runnymede data with results from elsewhere. So far there has been little opportunity to do this, but it is hoped to compare the Neolithic house ground plan in Area 19 with other excavated examples from Britain, and inputting these other plans will also provide a useful resource database. Perhaps the ultimate objective for the computerisation of Runnymede is to attempt a partial reconstruction of the site. While such graphical reconstructions must make many assumptions, they undoubtedly concentrate the excavator's mind on how the site may have looked and functioned. Ultimately, this must surely be the purpose of archaeological excavation.

### Acknowledgement

The authors would like to thank Dr Stuart Needham, director of the Runnymede excavations, for his participation and support throughout this project.

### References

- COONS, S. A. 1967. *Surfaces for Computer Aided Design of Space Forms*, MIT Project Mac, TR-41. MIT, Cambridge MA.
- LONGLEY, D. 1980. *Runnymede Bridge 1976: Excavations on the site of a Late Bronze Age Settlement*. Surrey Archaeological Society Research Volume No.6. Surrey Archaeological Society, Guildford.
- MAIN, P. L., HIGGINS, T. F. WALTER, A. ROBERTS, A. J. & LEESE, M. N. (1994). 'Using a three-dimensional digitiser and CAD software to record and reconstruct a Bronze Age fissure burial', in K. Lockyear & J. Wilcock (eds.) *Computer Applications and Quantitative Methods in Archaeology 1993* British Archaeological Reports, Tempus Reparatum, Oxford.
- NEEDHAM, S. P. 1991. *Runnymede Bridge 1978: Excavations and Salvage on a Bronze Age Waterfront*. British Museum Press, London.
- NEEDHAM, S. P. & MACKLIN, M. G. (eds.). 1992. *Alluvial Archaeology in Britain*. Oxbow Monograph 27. Oxbow Press, Oxford.
- NEEDHAM, S. P. & SPENCE, A. J. (forthcoming). *Runnymede Bridge Research Excavations, Volume 2, Area 16 East: A Late Bronze Age Midden and Underlying Neolithic Remains*. British Museum Press, London.