The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2005) *Short Presentations*
M. Mudge, N. Ryan, R. Scopigno (Editors)

Wandering the House of the Birds: reconstruction and perception at Roman Italica

G. P. Earl

Archaeology, University of Southampton, UK

Abstract
This paper describes a spatial analysis methodology derived from a combination of GIS and CAD/VR approaches, specifically using ArcGIS and 3ds Max. Texture viewsheds are defined and explored as a means for augmenting existing methods appropriate for analysing and presenting three-dimensional built spaces and landscapes. Within the context of a wider interest in alternative modes of spatial understanding, and specifically considering a domestic building from Roman Italica in Andalucia, Spain, the paper considers the potential for using the surface data accessible for all CAD models to form interactive and statistical archaeological and architectural interpretations.

Categories and Subject Descriptors (according to ACM CCS): J.5 [Arts and Humanities]: Architecture

1. Introduction

The site of Italica lies on a pair of plateaux leading down to a tributary of the Rio Guadalquivir near to Seville in Andalucia, southern Spain. It was occupied from at least the fifth century BC to the end of the Visigothic period, with its major surviving Roman elements deriving from the Trajanic period when a considerable enlargement of the town took place. The site includes a number of impressive public buildings including an amphitheatre and Traianeum, in addition to the remains of a number of substantial domestic buildings. This paper pays particular attention to one particular domus – the House of the Birds – and in particular its central peristyle. Ongoing research is considering the relationship of this space to others within the house.

Much of the visible architecture at Italica has been partially physically reconstructed (Figure 1), giving an impression of the porticoed streets, bath complexes and temple of the area now described as the Nova Urbs or new town. The House of the Birds is an impressive example of a Roman domus, with a substantial atrium and peristyle, and mosaic floored dining and other rooms. A number of computer aided design techniques have been applied to the town as a whole, including a series of animations produced for many buildings and an archive of three-dimensional AutoCAD drawings. These are gradually being supplemented by further models produced by the author in a research project ongoing since 1997 directed specifically at the experience of Roman domestic space at the site. This research is drawing together a variety of interpretative approaches and representative critiques to form an understanding of Italica's past derived equally as much from experience and perception as from statistical and graphical summaries.

Figure 1: *House of the Birds*

Interpretations of built spaces, and indeed of landscapes and in particular cityscapes, may draw upon an extensive range of analytical techniques. Thus, archaeology's appropriation and reorientation of phenomenological approaches has been widespread whilst, increasingly, agent-based and other simulations offer alternative analytical products. Similarly techniques such as isovists and access analysis provide interesting summaries of space, making accessible to statistical and network analysis routines the form and interrelationship of complex real and wholly artificial places. In addition, visibility approaches drawn from GIS are becoming increasingly sophisticated, taking in to account issues such as directionality, intention and habitual action in addition to broader perceptual concerns. Thus, in ongoing work by the author and colleagues at Southampton and in Spain, visibility between and around towns has been studied. Indeed it is now considered of key importance in the interpretation of the development of Roman urbanism in the south of Spain. The built environment forms a key element to this, since any modification of the broader landscape by the addition of structural forms may impact patterns of, for example, visibility, prominence, and perceptual dominance. In turn, the structure of internal space created perhaps in relationship to the surrounding (urban) landscape, to structural or practical factors, and to social relationships, is of considerable interest.

Given the growing employment of highly sophisticated computer modelling tools in archaeology, both as presentational and interpretative devices, many sites now exist at least in part as fully three-dimensional geometric databases. Elsewhere [Earl04; Earl05] I have discussed my position in the complex issues implicit both in the construction and subsequent consumption of such artificial past worlds. In the context of this paper it is sufficient to suggest that the production of such environments is a highly creative process, with competing, contingent theoretical influences. It is important to confront these influences and to demonstrate, where possible, a subset of the complex influencing factors. These might include correlating data sources, structural constraints or survey data, all incorporated into the three-dimensional models produced. This need not follow the treacherous path towards an attempt at objectification of the modelling process but instead draws upon our human perceptual and social interpretative strengths.

The models underlying the research described here offer archaeology a series of interpretative routes. Having constructed a digital simulacrum, the House of the Birds being not reborn but rather repainted, the creator is afforded the opportunity interactively to explore its constructed spaces (Figure 2). Putting aside debates as to the relationship between virtual and real pasts [Bateman2000; Gillings99] one may recognise the potential value of moving around a volumetrically ‘accurate’ equivalent to a given Roman space. Similarly, the pattern of light and shade seen in renders of the built environment might give an impression of its emotional impact. Taking alone the plan of the House of the Birds (Figure 3) one might also provide an insight into its potential use and experience.

However, the data upon which such analytical approaches and indeed a more interactive, subjective appreciation of the constructed environment are based also offer further analytical potential. We may begin to ask how we may extend the spatial summaries described so frequently in two or two and a half dimensional space (for example viewsheds or access depth) to fully three-dimensional numeric and visually interactive summaries of artificial environments.

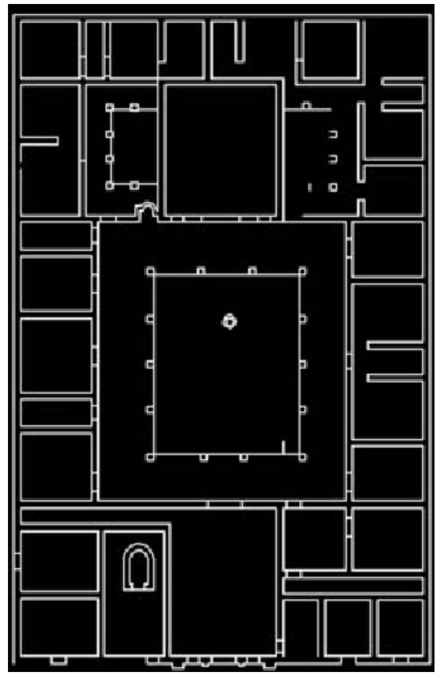

Figure 3: *Ground plan – House of the Birds*

2. Texture viewsheds

Every element of a computer model can be considered, in essence, to be comprised of planar faces. Each one of these in turn acts as a record of the complex interaction of light and shade in the environment of which it forms a part. In addition, at its most fundamental, it may record the visibility of given subsets of the environment relative to given observation points. The texture viewshed advanced here as an analytical tool provides a raster summary of complex three-dimensional surface parameters within a computer model by unwrapping the pattern of visibility (and optionally other surface characteristics) from the many planar faces in a model. The surface parameters include the potential visibility of an area given a range of viewing locations, the angle at which elements of a scene tend to be viewed, the geometric relationships defined between spatial elements, and conventional textural data such as surface treatment and illumination.

The data underlying the texture viewshed derives from the light mapping technique common to a variety of software approaches optimised for game environments. In 3ds Max this technique is known as texture baking and can be used automatically to produce a series of two-dimensional maps demonstrating the surface appearance of given objects in a scene. These maps are automatically associated with the model’s structural geometry via conventional UVW mapping co-ordinates – in the case of 3ds Max produced via a UVW Automatic Unwrap modifier. In the texture viewsheds methodology production of a large number of these unwrap maps has been automated, such that the model geometry can be mapped and remapped according to changing input parameters such as viewer location. Using the ArcGIS geographic information system these products are combined, manipulated and analysed in ways identical to conventional raster-based GIS visibility studies and to other image processing techniques.

Unwrapped textures once processed are remapped to the original geometry providing an intuitively accessible spatial summary that is fully three-dimensional, and fully view independent. The input parameters for viewers are controlled by producing point light sources throughout the area to be investigated (Figure 4). These are positioned to reflect likely viewer configurations and as such constitute a significant factor in the results produced.

Each point light represents a single monocular viewer with full rotational potential in all three axes. The system designed allows for a regular three-dimensional matrix of input points, where the XY and XZ intervals are equal. However, in order to optimise the approach it is also possible to reduce the number of viewers in the Z plane to a number thought sufficient to represent likely viewer heights. Thus, in the analyses undertaken to date only four viewer heights are considered, approximating roughly to individuals lying and sitting, and standing to two different heights at each of the locations in the building. Optimising viewer numbers is crucial since every additional viewer adds a considerable computational load since a texture viewshed is required for each object for each viewer location. The viewers are automatically animated such that in any given frame only one viewer contributes (i.e. is an active light source) to the scene. As a consequence a series of texture viewsheds can be produced by rendering and unwrapping the geometry for each successive frame, leading to a total number of texture viewsheds equating to the number of scene objects multiplied by the number of viewers. An example of the texture baked onto a single object is shown in Figure 5.

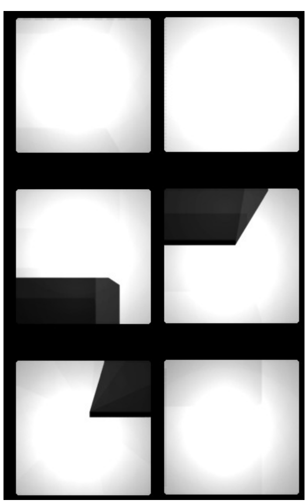

Figure 5: *Baked textures from a simple cubic room divided internally by a rectangular wall*

Here a simple room has been constructed from a single box primitive, with the addition of a dividing wall approximately three quarters of the total room height. When a viewer is placed in the centre of the room – i.e. when a single point light is used to illuminate the space – the texture unwrapped from the box interior appears as shown. The shadows equate to areas that cannot be seen, assuming a monocular view, whilst the changing angles of incidence for the light (which has no attenuation) indicate the relative angle of viewer to surface.

3. Visibility, distance and spatial character

Given a sufficiently large sample of viewer locations it has been possible accurately to record the varying patterns of visibility in the building studied. This interpretation relies upon the generation of two versions of a standard GIS product – namely the cumulative viewshed. In the first version each unwrapped texture is automatically converted to a binary value via a GIS raster reclassification operation, where any area receiving direct illumination is valued as 1 and the remainder as 0. The product of these reclassified raster datasets provides a texture viewshed for all elements of the model relative to all viewer locations and indicates very clearly which areas are most prominent purely in terms of potential visibility. This equates to a conventional viewshed in GIS where each cell can either be visible or invisible to a given observer.

The second cumulative viewshed uses all direct illumination data extracted. This will be a function of the angle between the light source and the surface normals for the given unwrapped object. It is standardised by ensuring that all surfaces have an identical matt white surface with linear shading between normals. In this way the texture viewshed demonstrates not only what areas are visible but also the angle at which they are viewed, itself a potential indicator of prominence (Figure 6). By comparing the latter to the former it becomes possible to extract in an automated fashion the changing patterns of prominence around any given modelled space.

It is important to stress that the outputs of this process are fully three-dimensional when mapped back to their originating geometry. As a consequence analysis can proceed via numeric summaries or alternatively through interactive exploration of the modelled environment. In the study of the House of the Birds, following normalisation of the various texture viewsheds for individual objects according to the objects' surface areas, the numeric summaries indicate clear areas of prominence – in particular the major doorways defining the key axes of the peristyle. This is supported by interactive exploration of the model and will be further explored as this initial simple case study is developed in the future.

Visibility does not constitute the only summary property provided by this methodology. Indeed, visibility alone is insufficient to describe the complex perceptual stimuli developed by a given model scenario. Distance from viewer may be similarly important. For example, without distance two parallel surfaces at offset distance relative to the viewer would appear identical in terms of visibility. The distance of each mapped surface from each viewer can be calculated by comparing pairs of texture viewshed maps produced using a point light source with no attenuation and with linear fall-off. Thus, given a knowledge of the range of the linear attenuation variation in pixel output directly corresponds to a three dimensional offset relative to the viewer location.

Combining the visibility and distance values provides a map of the spatial character of the spaces investigated which may provide a useful correlate to other methods. For example, enclosed corridor areas have distinctive relative distance values compared to open areas and stairwells. Visibility statistics may thus take account both of angle of view and distance from viewer, in a way identical to reclassification of viewsheds according to perceptual norms. It might be argued that a plan provides the opportunity to derive similar statistics. Is a corridor any different, really, when understood in an extra dimension ? It could be argued that although as archaeologists we often represent data in a plan form we do not envisage the spaces thus defined as operating in a two dimensional fashion. As a consequence the production of three dimensional summaries provides a more realistic impression of our understanding of the space. It could be argued that it is rare indeed that an archaeologically attested space can be considered regular in terms of horizontal profile.

The statistics made accessible by this approach offer a potentially powerful comparative to methods based on two-dimensional plan data. The only limiting factor is the number of individual rasters produced. Clearly one stimulus for the production of three-dimensional archaeological graphics is to explore the potential configurations of lost spaces. This requires considerable variation in structural layout. However, where variations in structural layout are to be considered the processing required to produce the texture viewsheds may increase exponentially when extracting visibility, surface texture and shading, distance and other spatial characteristics. As an example, the grossly simplified model of the peristyle shown above contains approximately 150 separate objects. Given 100 viewer locations the total number of unwrapped textures is 60,000, or c. 250 billion pixels. This total has been greatly increased in ongoing analyses using an optimised version of the geometry used to produce Figure 2 above. Thankfully 3ds Max allows the texture baking process underpinning the methodology to be carried out on a render farm of up to 999 work stations. The ArcGIS raster processing stages can similarly be comprehensively automated. In addition, since the majority of the rasters produced form part of later visual and numeric summaries the required storage overload is only ever temporary. However, the storage and processing requirements are certainly non-trivial and would benefit from custom written algorithms.

4. Lighting studies

The methodology described is able to take account of changing patterns of illumination, where such information is available, albeit limited by the global illumination techniques available within 3ds Max. This may be very important. For example, where a given area seems highly prominent in terms purely of visibility potential and distance, it may be that the pattern of illumination throughout the day serves either to support or diminish this visibility. Furthermore, addition of artificial lighting might considerably alter the apparent spatial character of any given area. Previous analyses based on subjective assessment of such data are important but gain a great deal when considered in conjunction with the spatial summaries offered here. Indeed, unless a thorough understanding of light transport is obtained spurious conclusions may otherwise be drawn. This area in particular will benefit from increasingly accurate light modelling techniques.

For the House of the Birds case-study baked textures were created throughout the day at four times in the year. This in turn provided a background level of illumination against which to compare the other results. Initially such records of the light used a constant diffuse white texture acting as a control for the radiosity solutions produced (see for example Figures 7 and 8). However, given the clear influence of the likely textures present within the building a series of additional lighting studies were conducted in order to assess the relative luminance and illuminance values for the spaces studied.

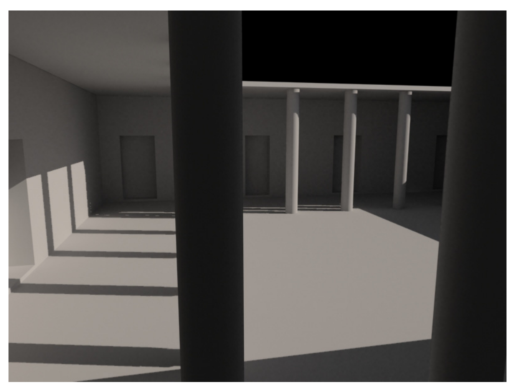

Figure 7: *The peristyle on a summer morning – summary of baked textures*

A problem with deriving these illumination maps from the standardised, constant textured geometric components has been that the light transport simulated does not reflect the transmittance and other properties inherent in a given material. Thus any such control study has been supplemented by assessments of lighting taking in to account potential building elements such as stucco and brick. The 3ds Max architectural textures used in combination with its radiosity advanced lighting seem to provide useful and apparently accurate solutions in these cases. However, considerable more work is required in order fully to assess the techniques' physical accuracy and indeed the significance of such accuracy in the terms phrased by this inquiry. Adoption of a more advanced suite of physically accurate modelling and rendering tools in the future could greatly enhance the overall efficacy of the texture viewsheds produced in the way described previously but have to date been little used.

5. Conclusions

The methodology described here is still very much under development. Considerable more work is required to achieve full automation of what can at times prove extremely laborious. However, the results are already very promising, and as archaeologists working with computer graphics we may benefit from the fact that so many key monuments already exist in a form ready for such analyses. Clearly by unwrapping the standardised textures from all objects in a given scene a range of new summary statistics becomes available. One might calculate the most prominent areas of an enclosed space or indeed a landscape in terms purely of visibility, taking into account angle of view, or indeed consider automated extraction of clear changes in angle. Similarly, the distance data that may be produced provide numeric equivalents to what in archaeology we might describe as liminal, encounter or transition spaces and crucially allow for direct standardised comparison between layouts or sites. Already in the House of the Birds it is clear that plan alone cannot define the complexity of potential spatial orientations. As we discuss and write about, for example, differential use of Roman domestic space as a means for mediating social status, we might now also consider how perceptually the access to and inhabitation of these spaces differs. This consideration may be performed not solely in terms of subjective assessment of modelled environments, or indeed in the rather rigid dislocation of space required by two dimensional analytical tools, but through a range of alternative methodologies.

Having summarised spatial character relative to given viewer locations it becomes possible to reuse this information in simulations. Thus, taking a given starting point one might derive pathways optimising view of a particular area or maintaining a given degree of open-ness. This can take advance of techniques such as the 3ds Max look-at controller. More generally, by taking account of orientation of view, increased awareness of particular surfaces through habitual viewing also becomes possible, adding an analogue to memory to the space modelled and interrogated. This requires that, as a given viewer interrogates the remapped spatial summaries, the camera orientation is recorded. Subsequently, recalculation of the texture viewsheds, using a directed spotlight exactly matching the parameters of the recorded field of view, produces output tied in to a particular engagement with the space, rather than as an abstract summary. Conversely, by plotting a predefined direction through a modelled space one might extract the changing spatial characteristics associated with such motion, again as has been explored in terms of landscape viewsheds in recent years [Llobera2003].

The limitations to such further analyses remain those of space and processing power. However, given the rapid development in both of these areas it is hoped that texture viewshed analysis might become a regularly performed addition to extant spatial interpretations in years to come.

References

[Bateman2000] BATEMAN, J.: Immediate Realities: an anthropology of computer visualisation in archaeology. *Internet Archaeology* 8 (2000)

[Earl05] EARL, G.P.: At the edges of the lens: photography, graphical constructions and cinematography. In Evans, T. *Digital Archaeology*. Routledge, London, Forthcoming 2005

[Earl04] EARL, G.P.: Video killed engaging VR. In Smiles, S. and Moser, S. *Envisoning the Past.* Blackwell, 2004, pp. 204-22

[Gillings99] GILLINGS, M.: Engaging place: a framework for the integration and realisation of virtual-reality approaches in archaeology. In Dingwall, L., Exon, S., Gaffney, V., Laflin, S. and van Leusen, M. *Computer applications and quantitative methods in archaeology 97,* Archaeopress, 1999, pp. 247-54

[Llobera2003] LLOBERA, M.: Extending GIS based analysis: the concept of visualscape. *International Journal of Geographic Information Science* 1(17) pp. 1-25

Figure 2: *Simple reconstruction model of the peristyle*

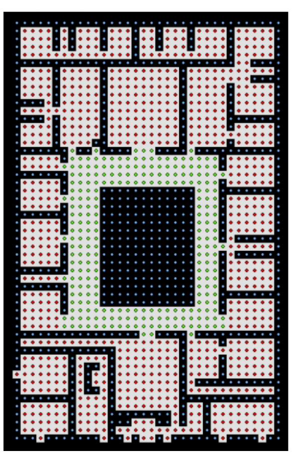

Figure 4: *One configuration of 2D viewer locations (green: viewers in peristyle; red: other viewers)*

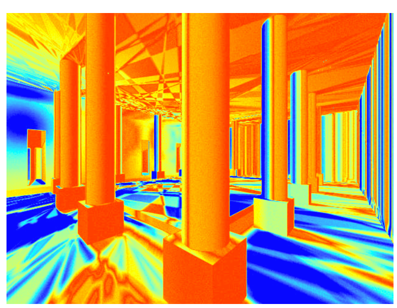

Figure 6: *Cumulative texture viewshed of a simplified peristyle model, derived from 100 standing viewers (blue: more visible)*

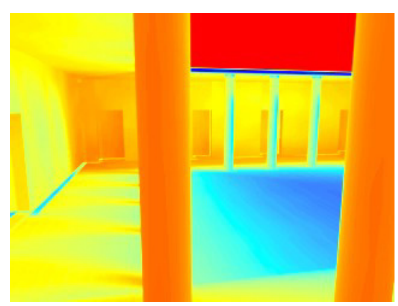

Figure 8: *The simple peristyle throughout the day – summary of baked textures (blue: greater illuminance)*