Calculating the inherent visual structure of a landscape (‘*total viewshed*’) using high-throughput computing

Marcos Llobera[1], David Wheatley[1], James Steele[1], Simon Cox[2], Oz Parchment[3]

[1] Archaeology, School of the Humanities, University of Southampton, Southampton SO17 1BF
[2] School of Engineering Sciences, University of Southampton, Southampton SO17 1BJ
[3] Information Systems Services, University of Southampton, Southampton SO17 1BJ

Abstract. This paper describes a method of calculating the inherent visibility at all locations in a landscape (‘total viewshed’) by making use of redundant computer cycles. This approach uses a simplified viewshed program that is suitable for use within a distributed environment, in this case managed by the *Condor* system. Distributing the calculation in this way reduced the calculation time of our example from an estimated 34 days to slightly over 25 hours using a cluster of 43 workstations. Finally, we discuss the example ‘total viewshed’ raster for the Avebury region, and briefly highlight some of its implications.

Key words : Archaeology, landscape, GIS, viewshed, total viewshed, inherent visibility, Avebury, neolithic

1 Introduction

The aim of this project was to explore the use of low-cost commodity computing facilities for intensive viewshed calculations, in particular, for the calculation of the *total viewshed* (sometimes referred to as the *inherent viewshed*[1]). This is a type of *visualscape* (see Llobera *in press*, 2003), similar to the *cumulative viewshed* (Wheatley 1995), that results after the viewshed from each cell in the DEM has been calculated and added together. Essentially, it provides a crude description of the pattern of visibility inherent within a landscape as a product of its topography. Unfortunately, the regular generation and use of these products at useful resolutions has proved impossible due to the very high computational intensity required to calculate them. This is because computation using conventional (naïve) algorithms requires, for each location, comparison of the target with all other locations in the elevation matrix. In a raster of dimensions X by Y each individual operation requires XY calculations, and thus X^2Y^2 calculations for an inherent visibility raster.

Despite this, *total viewsheds* are worth studying not only as a possible means for establishing the significance of *cumulative viewsheds* but as a source from which new information may derive (*e.g.* visual prominence). Where statistical investigation of visibility patterns among archaeological monuments is desired, then the total viewshed is particularly useful in that it provides the population against which samples (groups of archaeological sites) can be assessed, obviating the need for inferentially less powerful two-sample approaches that compare an archaeological sample with a second, randomly generated, sample of locations (see Wheatley 1995, 1996).

Being a natural extension of *cumulative viewsheds*, and ultimately of viewsheds themselves, *total viewsheds* of course suffer from the same sort of limitations as the former (see Fisher *et al.* 1997, Lake *et al.* 1998, Wheatley and Gillings 2000, Gillings and Wheatley 2001). Broadly these can be thought of as either methodological or theoretical, with methodological problems being further considered as either primary (related to computation) or secondary (concerning their application in the real world) issues.

For reasons of brevity, we would not wish to develop an argument here about the theoretical merits of quantitative approaches to visibility. While we acknowledge anthropological critiques of visibility studies which see it as promoting a particular (often ‘western’, ‘scientific’) form of understanding that privileges vision over other senses, we also subscribe to the view of Ingold that this critique tends to ‘*lay the ills of modern Western civilization at the door of its alleged obsession with vision*’ (2000:246). This critique encourages us to seek additional methods for investigation of other senses, and certainly argues for a more holistic approach to perception but it does not, in our view, present any reasoned case against development of substantive methods of investigating vision as a component of this wider theoretical goal. We also acknowledge the range of phenomenologically-inspired developing archaeological approaches to the visual structure of landscapes including, but we see these as complimentary, rather than alternative, approaches to understanding the visual structure of archaeological landscapes.

We would also not wish to rehearse wider methodological issues in detail here except to observe that many of the primary issues such as DEM altitude errors, edge effects and reciprocity may be now be regarded as resolved, and may be properly handled by following certain guidelines (Lake *et al.* 1998, see below). A number of secondary methodological concerns such as how to accommodate the effect of atmospheric attenuation or vegetation cover require further consideration.

2 Methodology

It is evident that any improvement that the viewshed routine may undergo will improve the calculation of *cumulative* and *total viewsheds* in general. Taking in consideration that *total viewsheds* result from the exhaustive use of the viewshed routine on the entire DEM, it is possible to introduce some simple improvements in their calculation. One of these is to store at each viewpoint cell the number of cells that are ‘in-sight’ after the viewshed has been computed (ibid.) rather than to use *Map Algebra* (Tomlin 1990) to add each viewshed one at a time. This

---

[1] This is the terminology used by Llobera (2003), previously *inherent viewshed* was used by Lee and Stucky (1998) as two separate terms *viewgrid* and *dominance viewgrid*.

allows us to save time, and memory, that otherwise would have been spent updating the *total viewshed* at each iteration, and it simplifies the way in which the calculation may be distributed to various computers (see below).

In this occasion, special care was taken to avoid any edge effects while calculating each viewshed. These effects appear anytime a *focal operation*, i.e. an operation that requires defining a 'neighborhood' around each cell (see Tomlin 1990) is executed. The way to handle this effect varies: if the size of the radius is small, it is possible to 'wrap around' the image, so that values that are missing in the neighborhood are read from the opposite side of the raster. This technique is commonly used in lattice models found in physics, but would not have much sense, except for some cases, in a geographical context especially in relation to viewsheds. Another technique is to normalize the result by the maximum number of cells in the neighborhood as it is found at each cell in the raster, *i.e.* at the corner of the raster we would have one quarter of the number of cells that we would find if the neighborhood was actually centered at the center of the raster, for given a certain radius (Llobera 2003). This technique allows us to compensate for possible low values towards the edges and corner of the raster but does not remove our uncertainty surrounding the final value that derives from the fact that visibility beyond the edge has not been tested. The best technique (not always possible) is to use an extended DEM, to use a DEM that is larger than the one we are originally interested. How much larger will depend on the search radius that we want to use during the viewshed calculation. Essentially we want to enlarge our original DEM so that even when we are calculating the viewshed for cells that are at the edge or corner of the original DEM, we are maintaining the same neighborhood size (*i.e.* we are checking against the same number of cells). This was the technique used here.

The study area for which we were interested in calculating the *total viewshed* was made by a DEM of 400x400 cells (at 50m cell resolution) roughly centered on the Avebury prehistoric monument complex in northern Wiltshire, England. The radius selected for each viewshed was of 20Km, *i.e.* ~400 cells (from observation this is approximately the maximum distance over which visibility is possible in ideal conditions in the Wessex chalklands) hence we used an extended DEM of 1200x1200 cells in which our study area was centered in the middle. Besides a search radius of 20km, each viewshed calculation used an altitude offset of 2.0m at the source location (representing the approximate height of a monument at that location) and 1.7m for the target point representing a viewer. The scenario simulated for each cell is therefore of a viewer looking towards a monument of up to 2m in height.

3 Condor: scavenging for idle computer cycles

As mentioned earlier, several strategies can be used to improve the calculation of the *total viewshed* notably optimising the basic viewshed algorithm. Ideally, we would have liked to combine several of these together but because of time constraints this was not possible and we concentrated on the single strategy of distributing the calculation of viewsheds for different viewpoints among a pool of computers.

To this end, we opted to use a cluster computing technique through the use of CONDOR HIGH-THROUGHPUT computing. CONDOR is a specialized workload management system for compute-intensive jobs freely available from the Computer Science department at the University of Wisconsin, USA. It may be employed with Unix, Linux and Windows computers networks, in this case, we employed the Windows NT version 6.4. Like other full-featured batch systems, it provides a job queuing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. Users submit their serial or parallel jobs to a central computer administrator. This computer places the jobs submitted into a queue, and chooses when and where to run them based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion. One of the main benefits of CONDOR is that it will work with a non-dedicated pool of computers. The pool may be part of a local area network (*e.g.* within a department). This makes CONDOR potentially very attractive for departments that do not have the benefit of a dedicated set of computing resources, frequently the case among Arts and Humanities departments. At its most basic level, CONDOR requires the submission of an executable file together with any data files (in our case a DEM) and set of arguments in order to execute.

The main idea of using the CONDOR system is to submit repeatedly what may be considered an 'atomic' operation to the designated CONDOR administrator which will then send each of the submissions to any available machine within the pool. In this occasion this operation would involve the calculation of one or several viewsheds. With this in mind, we initially considered the possibility of recycling some extant binary executable. Unfortunately this was not possible for CONDOR would not allow early 8-16 bit applications to execute for security reasons. Hence a new routine had to be written in C++ that could be submitted to the CONDOR system.

This routine was constructed around a single class, called *raster*, which provided all the required methods (*e.g.* I/O, getters and setters, line-of-sight or *los* and *viewshed*) needed to generate the application. A very simple *viewshed* routine was designed based on a variant of a well known algorithm, the *Bresenham*'s algorithm (Foley *et al.* 1990). Table 1 provides some information on the performance of the *viewshed* method.

This *viewshed* method was used to generate a kind of *cumulative viewshed* routine that could repeatedly be submitted to different nodes (*i.e.* computers) in the CONDOR pool. The routine operates in the following way, the user specifies the name of DEM file (an ASCII file), a window size within the DEM (which could be the same as the entire DEM), a set of continuous viewpoints within the window, a target and an observer's offset, a search radius, and an output filename, it then returns an ASCII file containing a list the number of locations (*i.e.* cells) 'in-sight' for each of the original viewpoints.

The pool of computers used for the project was one made up by 43 computers (13 of which were P4-1.5GHz -512Mb/1Gb RAM and 30, P4-2GHz-1Gb RAM) dedicated for the sole purpose of intensive computation. This pool, however, was shared among other users and could shrink (and in principle also grow) in size at any one time. This is precisely what happened during the calculation of the Avebury *total viewshed*, where several computers (up to 20) were physically removed from the pool during a short period of time and re-connected at a latter stage. While this event makes the overall processing time hard to estimate, it is a good illustration of the robustness of the system; in particular, how it is possible to shrink and expand the pool of computers at any time without compromising the system, and the overall execution of the jobs submitted.

To determine the number of jobs that needed to be submitted to the pool (*i.e.* the size of the jobs) to calculate the entire *total viewshed* several elements needed to be considered (*e.g.* size of DEM, search radius, etc). Amongst all of these, the time limit set by the administrator of the pool on the execution of any job submitted was crucial. Given that the pool might be shared by other users it is necessary to specify some time limit within which it is guaranteed that no job will be preempted from the pool. The time limit was of 11 hours in this case. Beyond this period of time, a job may continue to execute provided the machine is not required by another user, in which case the job would be cancelled (other versions of the Condor system, not windows NT, may migrate the execution to another node). Some estimates were made to determine how many viewpoints could be computed within this time period in order to guarantee that no cancellation would occur; given the size of the DEM and the radius of search for this example this number turn out to be ~1000. This meant that 160 jobs were submitted to the pool.

To submit jobs to the Condor system is necessary to generate a simple ascii file, called a submission file, that enlists the parameters used in each job. A small routine, called *generateSub*, was written to help generate this file.

4 Results

The total computation time allocated to the calculation of the Avebury *total viewshed* was equivalent to 34 days 1h and 24minutes, a figure which represents how much time a single computer (with similar specifications to those of the pool), given the same DEM, parameters and routine, would have taken to generate the same result. The average computation time per computer was of slightly over 5 hours (~5h 6min). In real time, the entire computation took slightly over a day (25h 13m) though it could have taken a minimum of a bit more than 19h. This minimum was not achieved because of the re-shuffling of computers that took place during the computation, I/O operations, and because of administrative tasks that the CONDOR undertakes often putting on hold the computation at some node/s.

Table 1. Worst case scenario (Intel PIII 1.2 MHz, 256MB Ram). DEM with the same altitude value for all cells is used. Viewpoint is located in middle of the raster so no LOS is interrupted until it reaches each of the target cells within the search radius. Estimated time for reading a 1200x1200 DEM is 78 secs

Radius (cells)	Computing time (in seconds)
100	2
250	32
500	505
1000	1950

The output of the entire exercise generated in this case 160 ASCII separate files. The name of each file is numbered in such a way that allows it to be read in order and parsed together into a single file using an additional routine, called *assembler*. With the proper header the file is ready to be exported into a GIS.

5 The total viewshed for the Avebury region

As an illustration of the utility of the total viewshed, we imported our result into ESRI ArcGIS™, where an existing database of archaeological and environmental data was already in place. Our discussion here is confined to visual inspection of the products and a few cursory remarks about their significance. A fuller reinterpretation of the distribution of the Avebury long mounds, including these results, is in progress and will be published in due course.

This region is archaeologically and historically interesting for many reasons (*e.g.* Pollard & Reynolds 2002[2]) Most obviously, it contains an extraordinary range of large ceremonial monuments built during the third millennium BC (later Neolithic). These include the massive stone and earthwork monument of Avebury itself and its associated stone Avenues, Silbury Hill and the nearby West Kennet palisaded enclosures (Whittle 1997) and the more diminutive Beckhampton Enclosure (Gillings *et al* 2002).

Wheatley (1994, 1995, 1996) applied cumulative viewshed analysis to the earlier Neolithic long mounds from this region and used it to investigate whether these monuments were non-randomly located at places with higher landscape visibilty, and/or higher intervisibility with each other. Although the distribution did show both of these tendencies, in neither case was the association significant at the 0.1 level. A similar experiment with the monuments around Stonehenge produced a similar pattern of association, but in this case both were found to be significantly (at the 0.1 level). This difference may be attributed either to archaeological differences between the regions, or (more prosaicly) to the fact that the sample size is smaller in the Avebury region and so it is correspondingly more difficult to demonstrate significant associations.

With the benefit of the *total viewshed*, it is now possible to compare the visibility of the barrows directly with the distribution of general visibility within the landscape. Figures 2 and 3 show respectively the *cumulative viewshed*[3] for the long mounds, and the *total viewshed* calculated as described above. Visual inspection of these two side-by-side immediately draws attention to the following characteristics:

1. The area of chalk ridge that runs north-south between Monkton Down and Old Chapel (formed by Monkton Down, Overton Down and Avebury Down) exhibits far greater potential visibility for long mounds than its inherent visibility properties suggest. This effect is caused by the 'ring' of sites that surround it (Shelving Stone, Monkton Down, Old Chapel, Temple Bottom, Manton Down and Devils Den).

---

[2] There is a huge literature on the archaeology of the Avebury region. For brevity, references cited here are generally to the most recent major work on the subject which the reader should probably consult first, and where references to previous work can be found.

[3] Note that the sample of long barrows used here differs slightly from that used by Wheatley 1994, 1995, 1996. Further field observations and other research led to the rejection of three sites included in the earlier dataset which are now considered unlikely to have been EN long mounds, and the addition of one additional site identified here as the 'South Downs' barrow after Barker (1985). For this reason, the cumulative viewshed presented here should not be compared directly with that previously published.

2. Windmill Hill causewayed enclosure (between Milbarrow and Horslip long barrow) is remarkably non-prominent in the total viewshed, but is far more prominent with respect to potential long mound visibility
3. The opposite is true of the Adams Grave long barrow, and the causewayed enclosure opposite (Knap Hill). Here, the sites appear on a highly prominent part of the landscape, but their locations do not afford views of any other long mounds.

These effects are interesting although, of course, difficult to attribute with certainty to the intentions of prehistoric people. It is quite possible that some of these result from complex interactions of factors which were originally informing the placement of successive long mounds. Nonetheless they may be telling us something about the choice of locations for monumental construction during the earlier Neolithic of this region.

Whether or not the visual emphasising of the ridgeway is intentional (in the sense that it was an intended consequence of selecting particular locations for monument construction) it is still an interesting effect. The ridgeway is an historic trackway that follows a natural corridor of movement through the region and the concentration of (later) Bronze Age round mounds along it suggests that at least parts of the route were used as a trackway during the bronze age. There is no such concentration of earlier Neolithic mounds and so if the high level of monument visibility is important, then this may suggest that its use actually dates from the earliest occupation of the region.

The similar ‘visual emphasis’ of Windmill Hill is equally interesting. This was clearly a major focus for the prehistoric community during much of the fourth millennium BC. Although its chronological relationship to the long mounds is not entirely clear, it is likely that its main phase of use was preceded by some long mounds (Horslip and Windmill Hill barrows, for example) while others are more or less contemporary with it (Whittle *et al* 1999). As such, it could be that the selection of Windmill hill for construction of an enclosure was partly informed by the visibility of nearby long mounds, while the construction of other long mounds may in turn have been influenced by the enclosure. The site is on a low hill formed by middle chalk, and affords good views in all directions although it is clearly not as naturally prominent as, for example, Oldbury Hill or even Waden Hill to the south of Avebury itself.

The inverse of these effects can be seen at Adam’s Grave and Knap hill and, to a lesser extent, at West Wood. Here the inherent visibility of the location is far higher than their intervisibility with other monuments, and this must suggest that (if visibility is a factor at all) then the inherent properties of the landscape were implicated in the decision to construct a monument at these places, rather than intervisibility. In fact neither Adam’s Grave or West Wood offer views of any other known earlier Neolithic mounds.

These are preliminary observations only, and need to be pursued in more depth and with more related lines of evidence. However, if they reveal anything at all it is that it may not be profitable to pursue a single explanation for the locations of all these sites. Generalisations about the visual properties of the sites as a group (such as those offered by Wheatley 1994, 1995, 1996) are useful to a point, in that they enable broad statements about, for example, whether or not visibility is a factor in the location of the monuments to be made. However, to progress beyond these simple ‘case to answer’ questions the monuments are so variable both in archaeological terms, and in terms of their visual characteristics that it is necessary to consider each location separately, and then to seek to characterise their locations according to different visibility ‘signatures’ before more sustained analysis can take place.

Table 2. List of the long mounds used here

Grid reference (OSGB36)			Site name
406680	,	164800	Kitchen Barrow
415200	,	169650	Devils Den
401050	,	165990	Kings Play Down
406660	,	167730	Beckhampton Road
408710	,	169150	Longstones
409000	,	167150	South Downs
411630	,	172300	Monkton Down
403870	,	166100	Shepherds Shore
410460	,	167740	West Kennet
414860	,	172510	Temple Bottom
415690	,	165630	West Wood
404690	,	169310	Oldbury Hill
406370	,	166100	Easton Down
405480	,	165760	Roughridge
409020	,	169280	South Street
410370	,	171560	Shelving Stone
414780	,	171350	Manton Down
411240	,	163390	Adams Grave
411630	,	166860	East Kennet
408600	,	170520	Horslip
409430	,	172210	Millbarrow
401490	,	164830	Roundway
407680	,	165800	Horton Down
412900	,	172900	Old Chapel

6 Conclusions

This project has demonstrated that significant performance gains can be obtained using high-throughput computing systems such as these, without extensive reprogramming of original source code. The success of these systems derives from their ability to ‘steal’ processor cycles from workstations which would otherwise not be used and, as such, this can be regarded to all intents and purposes as ‘free’ computational power. It seems extremely probable that this approach could work for the calculation of a number of other landscape indices and ‘visualscapes’. While this paper should not be taken as a full reanalysis of the Avebury landscape, we have also briefly demonstrated that the inherent visibility product holds significant potential to shed further light on the visual structure of this and other archaeological landscapes.

Acknowledgements

The authors gratefully acknowledge the University of Southampton for providing the funding for this research in the form of a University Annual Grant in 2002.

References

Ashbee, P., Smith, I.F. and Evans, J.G. 1979 Excavation of three long barrows near Avebury, Wiltshire. *Proceedings of the Prehistoric Society*

Fisher, P., Farrelly, C., Maddocks, A. and Ruggles, C., 1997, Spatial Analysis of Visible Areas from the Bronze Age Cairns of Mull. *Journal of Archaeological Science* **25**, 581-592

Barker, C.T. 1985 The long mounds of the Avebury region. *Wiltshire Archaeology and Natural History Magazine* 79, 7-38.

Foley, J., D., van Dam, A., Feiner, K., S. and Hughes, J., F. (1990) *Computer Graphics: Principles and Practice*. Reading, Massachusetts: Addison-Wesley

Gillings, M., Pollard, J. & Wheatley, D. 2002. Excavations at the Beckhampton Enclosure, Avenue and Cove, Avebury: an interim report on the 2000 season. *Wiltshire Archaeological and Natural History Magazine* 95, 249-258

Gillings M & Wheatley DW 2001 Seeing is not believing: unresolved issues in archaeological visibility analysis, edited by Slapsak, B. *On the good use of geographic information systems in archaeological landscape studies*, COST Action G2 (European Union: EUR19708) : pp25-36.

Ingold, T., 2000 *The perception of the environment: essays in livelihood, dwelling and skill* . London: Routledge.

Lake, M. W., Woodman, P. E. and Mithen, S. J., 1998, Tailoring GIS Software for Archaeological Applications: An Example Concerning Viewshed Analysis. *Journal of Archaeological Science* **25**: 27-38

Llobera, M. (2003) Extending GIS based analysis: the concept of visualscape. In *International Journal of Geographic Information Science*, 17(1), pp. 25 - 49

Llobera (in press). What you see is what you get? Visualscapes, visual genesis and hierarchy. In *Digital Archaeology* (eds. Daly, P. Evans, T. and Trifkovic, V.). London: Routledge

Lee, J. and Stucky, D., 1998, On Applying Viewshed Analysis for Determining Least-Cost-Paths on Digital Elevation Models. *International Journal of Geographic Information Systems* **12**: 891-905

Piggott, S. 1962 *The West Kennet long barrow: excavations 1955-56* (London: HMSO).

Pollard, J. and Reynolds, A. 2002 *Avebury: a biography of a landscape*. (Oxford: Tempus).

Tomlin, D. C., 1990, *Geographic Information Systems and Cartographic Modeling*. (New Jersey: Prentice-Hall)

Wheatley, D., 1994, The application of Geographic Information Systems to Archaeology with case studies from Neolithic Wessex. PhD Thesis, University of Southampton.

Wheatley, D., 1995, Cumulative Viewshed Analysis: A GIS-Based Method for Investigating Intervisibility, and Its Archaeological Application. *Archaeology and Geographic Information Systems: A European Perspective*, edited by G. Lock and Z. Stancic (London: Taylor and Francis) pp. 171-186

Wheatley, D. 1996. The use of GIS to understand regional variation in Neolithic Wessex. *New methods, old problems: Geographic Information Systems in modern archaeological research.* edited by H. D. G. Maschner (Carbondale: Southern Illinois University at Carbondale) pp. 75-103.

Wheatley, D. and Gillings, M., 2000, Vision, Perception and GIS: Developing Enriched Approaches to the Study of Archaeological Visibility. *Beyond the Map* edited by G. Lock, (Amsterdam: IOS Press), pp. 1-28

Whittle, A 1994 Excavations at Millbarrow chambered tomb, Winterbourne Monkton, north Wiltshire. *Wiltshire Archaeological and Natural History Magazine*. 87 1-53.

Whittle, A. 1997 *Sacred Mound, Holy Rings. Silbury Hill and the West Kennet Palisade enclosures: a LaterNeolithic complex in north Wiltshire*. Cardiff Studies in Archaeology. Oxbow Monograph 74. (Oxford: Oxbow).

Whittle, A., Pollard, J. and Grigson, C. 1999 *The Harmony of Symbols: the Windmill Hill causewayed enclosure*. (Oxford: Oxbow).

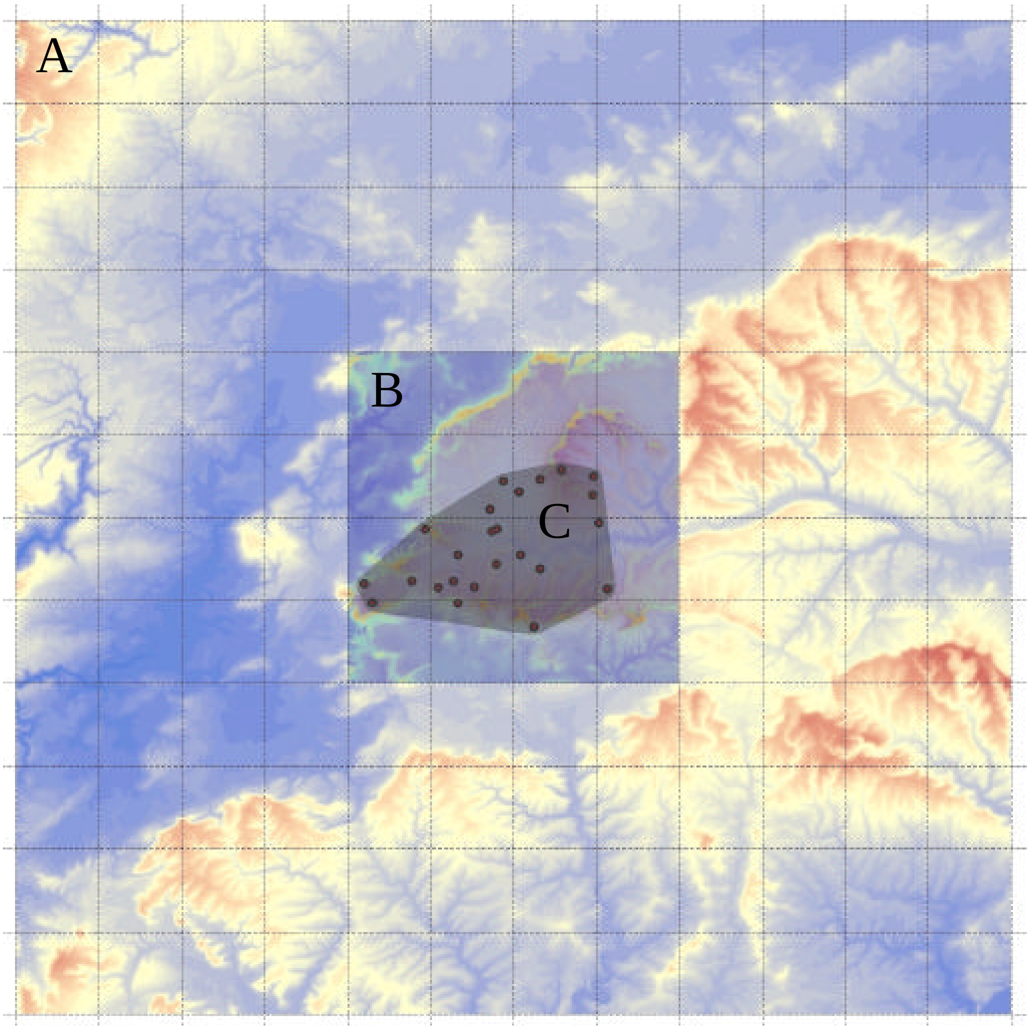

Figure 1. Showing (A) the source DEM of 60x60km (B) the calculated total viewshed of 20x20km and (C) the defined area of interest for the discussion, defined as the convex hull of the archaeological monuments buffered to 500m. Grid is 5km intervals on OSGB36 National Grid.

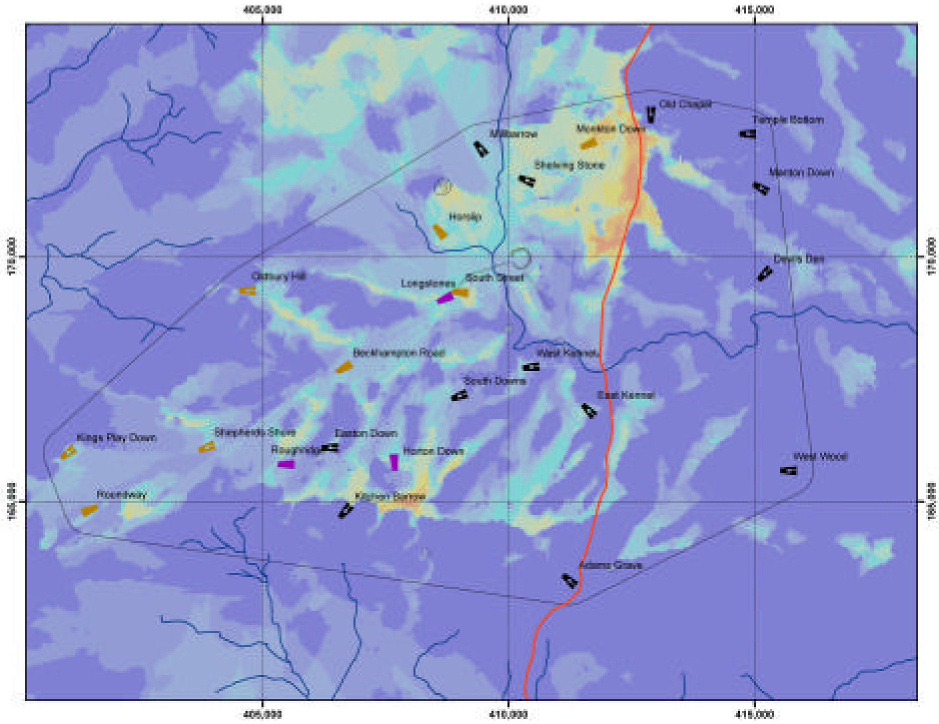

Figure 2. Cumulative viewshed for the 24 long mounds shown. Values represent number of potential barrows visible, ranging from 0 (dark blue) to a maximum of 16 (red).

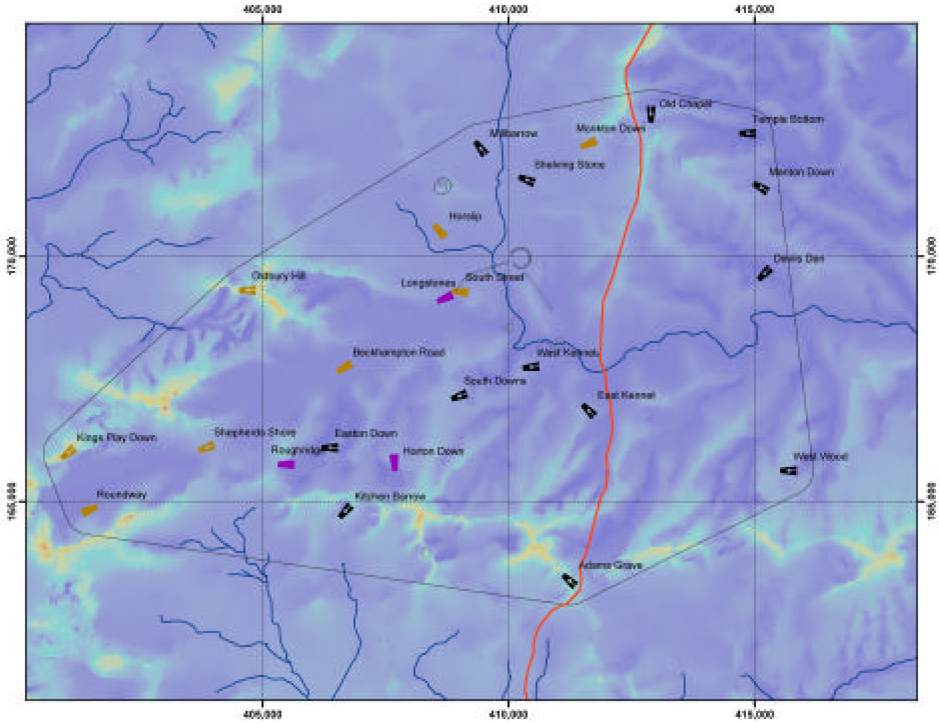

Figure 3. Total viewshed for the same region as in figure 2. Here, values represent the area from which a monument would be visible if built at each location. Lower values are in darker blue, with the highest values in red.