“Observing the Game”: What can Access Statistics Really Tell Us?

William Kilbride

Archaeology Data Service, Department of Archaeology, University of York

King’s Manor, York YO1 7EP, United Kingdom

e-mail: wgk1@york.ac.uk

Judith Winters

Internet Archaeology, Department of Archaeology, University of York

King’s Manor, York YO1 7EP, United Kingdom

e-mail: judith@minerva.york.ac.uk

Abstract

An ubiquitous feature of the “information age” is the hype that surrounds user statistics. Web sites are becoming an increasingly important and increasingly sophisticated marketing tool, and the success or failure of commercial web sites depends on meeting important thresholds of hits per day. Archaeology, at least within education, is shielded from these pressures, though we take pleasure in noting the more obscure locations or are disappointed when there seems to be a dip in demand for our all-important web servers. User statistics can provide an enormous wealth of information about who is visiting our web sites, where they come from, what interests them and when they work. Such issues can be discussed in theory, though in practice interpretation is hampered by some basic problems.

Using access statistics gathered by the Archaeology Data Service and Internet Archaeology, this paper explores the trends represented in them. This analysis, informed by a discussion on the place of the Internet in wider social agenda, suggests that much can be gained, but only if questions are formulated with care. It will be argued that the user populations engage with the digital content of web sites, partly in response to the nature of that content, and partly on the basis of existing working practices. It is proposed that virtual heritage services are not a threat to conventional engagements with the archaeological record, but that they supplement that engagement for particular audiences.

Key words: Internet, accessibility, Access, server statistics, usage patterns, interactivity, readership, online resources, heritage services, working patterns

1. Introduction

If anyone should understand computers and the impact of the Internet, archaeologists should. The impact of technology on social reproduction, and the reproduction of technology through social practice lie at the heart of much archaeological research, though admittedly these are often written in the past tense. There are a number of readily available tools that may allow us to investigate the social impact of the Internet, and the impact of society on computers, but the most frequently quoted of these come in the form of user statistics. Web sites are important in commercial marketing, and their success or failure depends on meeting important thresholds every day.

As archaeologists, we are largely insulated from these pressures, though we take pleasure in noting the apparently insatiable public appetite, or the relatively unlikely locations of our audiences. The world, we are told is getting smaller. Numerous national and international policy initiatives are establishing mechanisms to open information sources and to distribute electronically data that otherwise would have been unavailable. These strategies are presented in context with wider policies of social inclusion, using the strengths of the cultural and heritage sectors to foster participation in the body politic. More than a decade of research shows that the Internet and computing technologies can indeed give rise to specific forms of community participation, specific forms of interaction and practice which are impossible or impractical by other means. The user statistics we generate provide a wealth of information on who is visiting our web sites, when and to an extent how they obtain (or fail to obtain) cultural heritage information. To a certain extent, these figures give us an insight into the way that computing and information technologies, in particular the Internet, inform the development of archaeological research, and the relationship that archaeology has with its public. But interpretation of these numbers is problematic, and the conclusions dependent upon any number of problems. This paper will explore the issues associated with evolving information policies and the community building exercises they seek to support. In particular, we shall highlight the strengths and weaknesses of the interpretation of figures accumulated at the Archaeology Data Service (ADS) (http://ads.ahds.ac.uk/). The interpretation of these statistics is facilitated by comparison with similar figures for the electronic journal, Internet Archaeology (http://intarch.ac.uk). As will become apparent for technical reasons, Internet Archaeology provides the only reliable and comparable statistics for this analysis.

2. ADS access statistics

The ADS, like many web servers, retains a certain amount of information about the users who visit the web site. This data can be used to generate some fairly crude statistics, giving a rapid, if perhaps unreliable insight into the use of the ADS’s online catalogue, ArchSearch (http://ads.ahds.ac.uk/catalogue/), and the various documents that are associated with it. These figures reveal almost constant growth, but with enough dips and peaks to show

that there are other variables underlying the growth (see figure 1). We can identify certain peaks with particular events in the history of the organisation, such as the launch of ArchSearch, the online catalogue in September 1998, and the addition of individual elements to the catalogue. The launch of the ADS Library, with back issues of the Council for British Archaeology's Research Reports and Occasional Papers in March 2000 brought an unprecedented demand on the server. These figures are calculated on a daily, weekly, monthly, and annual basis, giving slightly different insights into the demand. Thus, the year from April 1999 to March 2000 brought just over 764,000 successful requests to the server, while the very last month of that year, March was a busiest single month, with around 90,000 requests. The week between 19th and 26th March was the busiest single week on record, with over 26,000 hits. Not surprisingly, the quietest week was from 26th December, with fewer than 6,000 hits. Saturdays and Sundays are relatively quiet: the weekdays proving more popular.

In addition to identifying the peaks and dips in demand, the figures also give an insight into the users of that data. Thus, the log files for the same period identify different users and their contexts. Not surprisingly, the UK is the largest single user community, responsible for around 40% of the data downloaded. The ADS is based in the education sector, and the work we do is of particular relevance to the UK higher education sector. This is evident in our statistics also, with something like 38% of the data going to UK universities. Other countries are in evidence in a variety of guises, with the US and Europe taking up most of the rest of the demand. Many Internet service providers operate without an identifiable national domain name or with unresolved numerical addresses. Many of these are based in the United States, though they are more widely dispersed, representing almost 35% of the demand. A further 10% is served to the US education community.

In addition to giving us a broad idea of who uses the service, and when, it is, of course, possible to identify what particular resources are in demand, and how those users came to be referred to the ADS. We can identify the number of requests for individual directories and pages within these directories, showing which of the ADS's Guide to Good Practice Series are more popular, and the requests which users make on the catalogue. We can also identify the last page which the user visited, and thus get an insight into how users came to the catalogue. For example, two hundred users came to the ADS through the BBC's news server when it published an article about the ADS on their online news service.

3. Access 1: The Internet and the community (in the UK)

The increasing demand evident in these figures is hardly surprising given the rapid growth of Internet services and the numbers of people connecting to the Internet. Increasing numbers of consumer services are migrating to electronic formats, with the financial and retail sectors developing strategies for electronic commerce, encouraging, and in some cases forcing traditional customers to adopt novel practices that depend upon the Internet and allied technologies.

Access to information is one of the corner stones of political rhetoric in the UK and EU, a drive that may have particular implications for the heritage sector and for archaeology. Such policy ini-

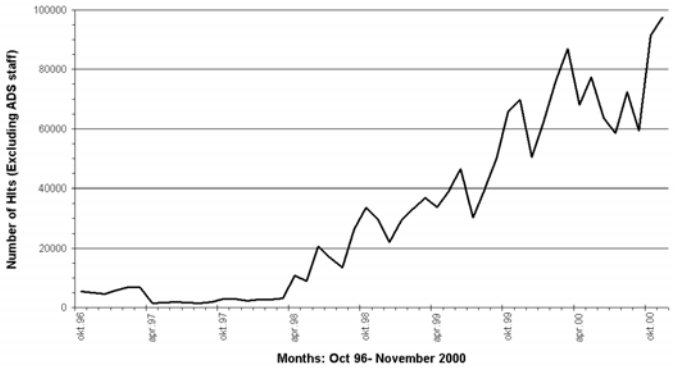

Figure 1: Hits on the ADS website Oct 96-Nov 2000.

tiatives depend on the rapid development of physical infrastructure and content, both of which are becoming priorities of government (e.g. EU 2000, Cabinet Office 2000). These initiatives are not formulated in isolation, but invariably form parts of wider social and economic policies. The United Kingdom government now has an evolving policy on information, aiming to use this investment not simply to disseminate relevant government documents, but, we are told:

> "The Government has a mission to modernise - renewing our country for the new millennium. We are modernising our schools, our hospitals, our economy and our criminal justice system. We are modernising our democratic framework" (Tony Blair, foreword to HMSO 1999).

Allied to this, ambitious policies to combat "social exclusion" have moved the accent from the conventional concerns of health and welfare, to wider issues of social and cultural regeneration, with some emphasis being placed on the cultural sector as a means of engaging actively with groups disenfranchised by conventional social and economic policies (e.g. SEU 2000). The move towards engagement for social regeneration is paralleled by a recognition of the importance of these new technologies in economic developments, represented by exercises such as the Department of Trade and Industry's *Information Society Initiative* (http://www.isi.gov.uk/isi/). Investing in communications infrastructure to facilitate social and cultural inclusion will provide the foundations for a wider economic engagement, by raising standards in universal IT skills, and providing a platform by which business can expand its use of communications technology. For many years governments have supported industry with the infrastructure of roads, railways and ports. In the "information age" the "knowledge economy" will receive a similar sort of investment. This, at least is the political theory of community building on the Internet.

On a more practical level, however, a number of important steps have been taken to realise these ambitions: at least to provide the tools that will make these theories a reality. The "open government" "learning society" and "knowledge economy" of political aspiration are supported by a number of concrete policy initiatives that will undoubtedly change the nature of archaeological practice, for better or worse. The most important of these developments, at least from an historical perspective, is the introduction of legislation on the freedom of information (inter alia House of Lords 2000). Knowledge acquired and maintained at public expense for the public good will soon be necessarily put into the public domain. The precise implications of this legislation will not be clear for some time, but plans are already afoot to exploit

the potential of the Internet to distribute such information. Parallel legislation on the management of public bodies to achieve "best value" are already in place and will have an extraordinary impact on local government, the planning process and the degree to which different agencies share expertise and information (HMSO 1999a). These and related projects provide legal substance to some of the rhetoric. In practical terms, a number of organisations exist which are either already delivering these objectives or are some way down the line in implementing them. In the Higher and Further Education sector, the Joint Information Systems Committee (JISC) has a formidable track record in providing technical assistance and support to large numbers of users in numerous locations, experience gained in working with "JANET", the academic network which connects the Higher and Further Education sectors to each other and the Internet (http://www.jisc.ac.uk). JISC has developed a coherent strategy for providing data to its networked community, for the benefit of research and learning with what it describes as a "Distributed National Electronic Resource". The Archaeology Data Service, located in a university but working for the whole research community is an example of the Distributed National Electronic Resource in practice. In the secondary and primary education sectors, the British Educational Communications Technology Agency (BECTA) is leading developments for a National Grid for Learning (http://www.ngfl.gov.uk/). In the museum sector, the Museums and Galleries Commission have been working to bring data about collections together from numerous museums and present it in a single portal operated by the Cornucopia project (http://www.cornucopia.org.uk/), while the electronic "Twenty Four Hour Museum" (http://www.24hourmuseum.org.uk/) has been recognised for its efforts by being designated as a National Museum. In different ways, these different initiatives are implementing and helping to inform strategies to put communities in touch with each other and with cultural and heritage datasets.

4. Access 2: The virtual community

Plans for the creation of "virtual communities" and "information societies" are far from new, and their analysis has become a regular feature of the academic world (see Hobart and Schiffman 1998 for an impressive historical insight into what might not be such an unusual transformation after all). For many years now, various researchers have attempted to understand and anticipate the impact of computers and communications technologies on social practices, and upon our own discipline also (inter alia: Harnad 1991, Huggett 1995, Holtorf 1999 see also Internet Archaeology 6 and Exon et al. 1999). Although experiments and studies have investigated all sorts of computer-mediated communication, these have been largely undertaken from a psychological perspective, and tend to be deterministic in nature, assuming that the social outcomes described derive from the material characteristics of the technology (Markus 1994). In that respect, they investigate the technology in question without making reference to the expectations or aspirations of the people operating the systems (Walther 1992). No good archaeologist would be satisfied with such an approach. Lab based studies suggest that a reduced number of lines of communication result in aggressive, task oriented behaviour, whereas real world anthropology suggests that the register of computer mediated communication reveals the development of "speech communities" with discursive norms and an awareness of shared or fragmented histories (Cherny 1999:247-295). It is these discursive practices more than anything that have been used to justify the claim that these are real and vibrant communities. Thus, the Internet allows for a form of purposeful social action, not previously possible.

Such purposeful action and interaction occurs right across the Internet, by email and by discussion groups, but historically has been the particular provenance of "Muds" and "Moos" (Cherny 1999:1-31). Muds – multiple user dungeons – developed as gaming tools in the late 1980s as a means of putting numerous, geographically isolated fantasy role playing gamers together into one playing environment. Thus, different players encountered each other in a form of social interaction, working out and developing new storylines for games that became as variable as the characters present there. Moreover, the introduction of programmed and specified objects into these dungeons gave users a consistent content with which players could interact, providing a foundation for shared experience and knowledge, and thus a subject matter for coherent social discourse.

The technology of the MUD is not, however, restricted to fantasy battles or make-believe monsters. The multiple-object oriented (MOO) technology that supported these adventures was soon abstracted into other contexts, with different types of interaction and new content. MOO's proliferated in the early 1990s. The most renowned of these, called LamdaMoo, hit the headlines in the early 1990s with many thousands of participants taking on many different characters in a hippy-style commune, imagined as a large house in which each participant was invited to build their own space, then encounter other participants in their own unique spaces. The population of LambdaMoo demonstrated consistent learned behaviour with the terms of social engagement being developed ad hoc. The infamous LambdaMoo "cyber-rape" and the vigorous protests which it provoked are evidence that the communities thus created developed within the constraints of the technology, but in directions that are the property of their populations (Cherny 1999:31).

MOO's and MUDS are still with us. The multiple-object oriented technologies of the early 1990s were often based on technologies that now seem to be anything but interactive. Combined with an expansion of bandwidth and more flexible communications protocols, they have given rise to new forms of engagement with digital content and populations. Specifically, the almost unbridled proliferation of web "Chat Rooms" and related chat relay services on the World Wide Web have inherited and consolidated social discourse on the Internet. Chat services have a thoroughly deserved reputation for absurd triviality (Arnaut 1998), yet the technology that facilitates such inanity is far from trivial. MOO configurations have been applied in a number of contexts, most usefully in the development of "virtual learning environments" (e.g. http://www.cose.staffs.ac.uk/). These teaching tools use multiple media, and many objects to bring students and teachers together with specific materials for discussion and experimentation. Focussed on specific learning outcomes with managed interactions, or released to open ended learning, these tools are being used to put increasingly large numbers of students in touch with each other and with increasingly hard pressed resources. Used in this way, these virtual environments have the benefit that in addition to the manifest subject matter, students and staff rehearse or develop transferable skills in computing and IT, thus enhancing their own digital creativity.

Of course, there are numerous other expressions of community life and social action on the Internet. In this respect all that is being suggested here is that the Internet is not used for its own sake, but for the accomplishment of given, and often very conventional, tasks. Online consumerism, on line teaching and learning, electronic protest, hero worship and romance are all prevalent. In these contexts, and demonstrably in others, the communities formed are supplementary to, not in competition with, other forms of community interaction. In this sense, it is the communities that drive content, as well as the content that forms communities.

Consequently, the relative value of access statistics may lie in estimating the scale and configuration of the communities that use and explore this content. If the aim is to establish a local community, then large numbers of external participants may represent a failure to achieve such an objective. Where the aim is to stimulate certain practices at certain times in the day or week, then this too may be represented in the relevant statistics. Moreover, constant levels of activity through the day or week, when contrasted with relatively localised populations may demonstrate the migration of working practices from the normal working day. Analysis of these figures is only one part of a wider and much more challenging anthropology of the communities in question.

5. What access statistics really *don't* tell us

We may wish to explore our digital communities from the comfort of our own offices, but like all anthropological analyses, the data we use has its own inconsistencies, flaws and inadequacies. It is as well to be aware of these before taking that analysis forward.

First and foremost, and perhaps most critically, the number gathered by our automated processes don't represent the real numbers of hits or user sessions on any given web-site. There are numerous problems that distort the figures. On one hand, the use of large numbers of files inserted into pages can, in some circumstances lead to a gross overestimate of the real numbers involved. This is particularly the case where graphics are used. In addition the constant use of the web site by its developers can count in the number of hits, but for uses that are hardly relevant to the exercise. Both can be overcome if the data gathering tools are configured adequately: graphics and local users can be filtered out. On the other hand, web and personal browser caches, established to reduce network traffic, can lead to a gross underestimate of the popularity of a site. Thus, instead of loading images directly across the Internet, a copy of the data saved locally will suffice, meaning that a genuine hit is not registered. In large networks, a combination of external cache, proxy and mirror services means that even when the pages are loaded "as new" the web server may not register a hit in the normal way. This problem is particularly evident in the ADS access statistics where the University of York is responsible for an inordinate proportion of use compared to other higher education institutions even though office terminals are specifically excluded from the figures. The reason for this is, in part, the configuration of national and local university cache and proxy services that are configured specifically to reduce the number of hits on the ADS server. There is no obvious way round this problem, other than to recognise that the numbers produced are interval numbers, not ratio ones (Shennan 1988:11-12).

Then, there is the not inconsiderable problem of understanding. While we may be able to track down an individual user session, and know something about the pages which appear to interest them, user statistics cannot tell us whether the end user actually understood the content of the page. In this respect at least, web delivery is entirely consistent with conventional publication. Authors and publishers can know how many copies they print and sell, but not how many copies are read or understood. It is perhaps true that web delivered data will escape the ignominy of propping up tables or soaking up coffee, but in the same way, once the data is released over the Internet, the author and publisher have limited control of what people do with the data presented in this way. This problem is normally presented in concerns about copyright and the ownership of information, but from other perspectives, it is more serious than that. From the perspective of teaching and learning, the degree to which students have attained the level of expertise expected cannot be assumed simply because the students have visited the relevant web site, or completed the tasks set. To that extent, the success of web delivered teaching materials depends upon the success of a broader environment of assessment and reporting. We might have a mechanism to report what sorts of things people want from web servers, but often cannot tell what they do with that data.

Returning briefly to the cognitive perspective we abandoned at the outset, it is possible to anticipate, and thus steer users in a number of directions, predisposing them to certain forms of activity. Ideally, resources should not be more than three clicks from the home page. Established wisdom suggests that interest levels fall as the number of mouse clicks rise. Moreover, large files or obscure formats will restrict access to those competent enough to manage such resources, or with higher specification connections. It is normally supposed that interactivity enhances the users' experience of web server, encouraging their participation in data sets that are otherwise incomprehensible. Yet even here conventional wisdom is disputed. One psycho-analyst has pointed out that the state of activity so sought after in computing has its inevitable counterpart in the state of passivity (Žižek 1999). Inter-activity likewise has its counterpart in *inter-passivity*: the more we attempt to create interactive environments so the theory goes, the more we risk creating passive users.

> "More precisely, the term interactivity is currently used in two senses: (1) interacting with the medium - that is, not being just a passive consumer; (2) acting through another agent, so that my job is done while I sit back and remain passive, just observing the game." (Žižek 1999:105-106).

Though presented in scurrilous tones, this theory is not as far fetched as it might seem: resources may require some work to understand, and this may take some activity. But at least the small number of users that persist will be better informed about the limitations, complexity and efficacy of the objects presented to them. This is hardly an argument for poor design, but the point is made.

In addition to problems associated with the internal coherence of the statistics, and the use that people make of the data, there are serious problems when we attempt to map our own statistics onto other similar data sets. Server side software and the logs they draw upon are configured for different purposes. Servers themselves have different characteristics: some provide data really quickly to slow or busy networks, others run slowly on faster networks. Some forms of data – especially streaming video or sounds – require

time-consuming processes; other forms of data are supplied with trivial overheads. Servers crash from time to time, or are taken offline for maintenance, enhancement and repair. These all mean that, without the most rigorous constraints, the results produced for one server are not comparable with results from others, even where the logs and statistical software are implemented in the same way.

These and other limitations mean that, with the best will in the world, the user statistics we generate are not going to answer some of the more interesting and important questions about the digital communities we develop on the Internet. Moreover, as a means of monitoring the performance of the service we provide, they are perhaps the crudest of measures.

6. What access statistics really tell us

Readers may by now be under the impression that we do not rate access statistics very highly, or that we would wish to see more carefully controlled and more comparable statistics. In part this is true, but the goal of this paper is to explore what these numbers really do tell us. As will become clear, the trick is not so much knowing what questions they can answer, so much as what questions we may reasonably ask. In fact, this wealth of data provides a powerful analytical tool to answer quite specific questions about how archaeologists and others engage with the Internet. Specifically, they reveal patterns of how archaeologists, and to a lesser extent the pubic, work in relation to different types of service. They show us how archaeologists work through the day, how they work through the week, and how they work through the year. In each of these, but especially in the last two, they may be used to interrogate the relationship between virtual heritage services, and real, conventional heritage services. In order to achieve this sleight of hand, two elements will be introduced. Firstly, for good reasons, comparison with another data set will be proposed – a comparison that for unique reasons is comparable to the statistics gathered by the Archaeology Data Service. Secondly, even these comparable figures will be massaged to make them more amenable.

The comparable data sets presented here are abstracted from the logs of the web servers of the Archaeology Data Service and the electronic journal Internet Archaeology, the first fully peer-reviewed electronic journal of archaeology. Though distinct entities, both these units operate out of the Department of Archaeology at the University of York. Thus, in a technical sense they are identical. They have identical profiles in terms of internal and external caches; they both connect to the Internet from the same network with the same number of network maintenance glitches; they both use the same server software and compile their figures in the same way. Indeed, not only do both services run across the same network, but their web servers are partitions of the *same* Unix server. Thus, server maintenance time and down time are exactly the same. For example, on the advice of security experts, the server was briefly disconnected from the Internet over the "Millennium" holiday for a five day period. Had Internet Archaeology and the ADS been working from different servers with different connections to the Internet, an interruption like these could render comparisons invalid. Even so, there is still diversity in the figures. In some respects, this diversity takes us to the core of the analysis, but there is sufficient difference in the figures to warrant a further qualification: the figures presented here are not the absolute values, but the proportion of total demand at certain times, expressed as a percentage of the total.

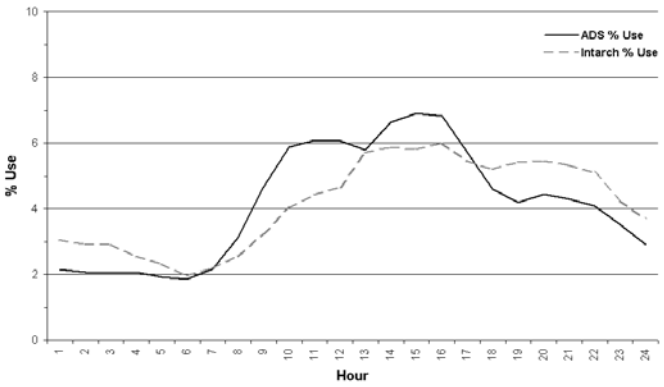

Figure 2: Working Patterns Through the Day.

This brief analysis looks at three different sets of figures based on time and the working patterns they represent. Many others forms of analysis are possible and more complex statistical modelling of these figures is certainly feasible. Moreover, an investigation of the geography underlying the use of these two web-servers might well prove revealing, and undoubtedly this informs the figures that are presented here. In this simple reconnaissance, however, we will investigate the relationship between Internet access and patterns of access to virtual heritage throughout the day, the same phenomenon through the week, and lastly, throughout the year. These figures were gathered over a twelve-month period from April 1999 to March 2000.

The levels of demand through the day are shown in figure 2. Perhaps the most striking feature of these figures is the way that the ADS profile varies so widely, while Internet Archaeology traffic fluctuates more gradually. It is clear that demand on the ADS site effectively starts around 0800 hrs, rising rapidly, and remaining constant throughout the day till around 1700hrs. The most remarkable element of this profile, however, is the clear dip in demand round lunchtime. This lunchtime dip at the ADS corresponds with a peak in demand for Internet Archaeology. Internet Archaeology remains more constant through the evening and even until quite late at night. There are, no doubt, numerous reasons for these particular profiles, but a number of explanations seem reasonable. Firstly, the content of the ADS catalogue and website is dominated by UK based resources. Though many of the documents are of general interest, it is the catalogue of over 300,000 archaeological sites and monuments in the UK that dominates demand. As we have already seen, the greatest demand for this information comes from the UK. What these figures demonstrate, however, is that users visit the site during the working day, using the catalogue for research or teaching from their desktops rather than from at home in their leisure time. It also suggests that, while the ADS catalogue may well have achieved a significant achievement in distributing resources widely across the UK, this change in the construction of space has not resulted in a change in the construction of time. Internet Archaeology, which has a more geographically diverse content and is more amenable to iterative browsing and recapitulation is clearly used in a very different way. Either, it has a more diverse audience, or an audience that is inclined to use the journal from the luxury of their own homes, and in their own spare time. A more thorough evaluation of the journal shows that it has both. (http://intarch.ac.uk/news/evaluation/index.html).

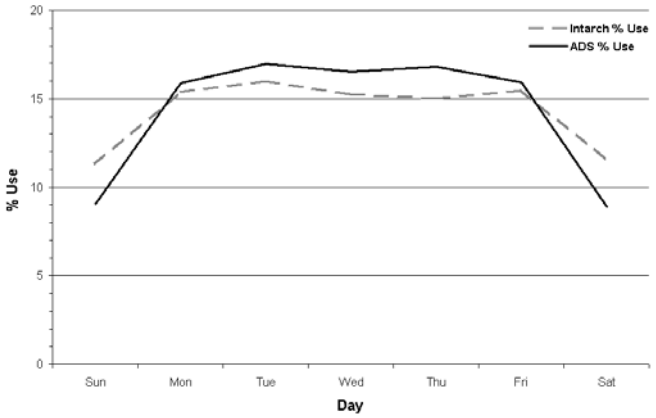

Figure 3: Working Patterns Through the Week.

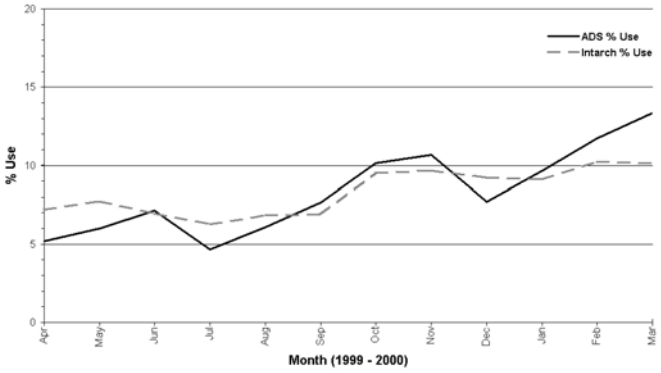

Figure 4: Working Patterns Through the Year.

The levels of demand through the week, shown in figure 3, tend to confirm this pattern. The demand on both is heavier during the week, as one might expect, but the journal tails off more gently than the catalogue at the weekends. Demand on the ADS drops to almost half of its weekday levels, whereas Internet Archaeology dips only slightly. Again, this suggests that different working practices are formed round different types of digital content. Moreover, by introducing other, conventional engagements in heritage, it would seem that the "Information Age" that we hear so many revolutionary things about, is supplementary to existing patterns of engagement. Museums and other conventional heritage sites have their largest visitor numbers at weekends: at precisely the same time as the ADS and Internet Archaeology see their least activity. Now the content of both these sites are academic in orientation, but there is sufficient public interest in both to merit a tentative conclusion about archaeology and its public. These figures suggest that virtual heritage information services really are supplementary to the public engagement in archaeology, and are not perceived by their user communities as alternatives.

This tentative conclusion is reinforced by analysis of the figures for use during the year, presented in figure 4. More than the previous sets, these figures have a clear drift, which could be eliminated by more sophisticated modelling: but this is not essential to make the points necessary here. Again, the ADS profile is more variable, the Internet Archaeology profile more constant. Again, both follow similar patterns of peaks and dips: more pronounced in one, less in the other. In both cases, the summer months of July and August are relatively flat, while the months of September, October and November, and then again January, February and March show almost constant increases. Internet Archaeology, however, remains noticeably stable throughout December, while the ADS figures drops away considerably. This pattern is strikingly different from conventional approaches to archaeological heritage, at least in the public sphere, where busiest periods are consistent with the holiday months of July and August. At least as far as the ADS figures are concerned, the opposite is the case. The peaks in demand are consistent with the peak times of the university term, so in that respect at least analysis of the figures demonstrates a congruence between the demand on the server and the ADS's principal target audience.

7. Conclusions

A number of tentative conclusions may be drawn from the preliminary analysis presented above. It is clear that moves are being made in the UK and the EU to invest more heavily in the infrastructure and content of the Internet in order to facilitate broader social goals. This is clearly consistent with existing practices on the Internet, where cultural interaction and purposeful social action are in evidence, not in isolation but as practices embedded within a wider social engagement. The means of reporting the role and impact of computing as a mechanism for facilitating broader social engagement exist in part through the analysis of access statistics, though these represent a complicated data set, which cannot be expected to supply cheap answers to expensive questions. The investment in information technology cannot therefore be audited purely by a facile appeal to access statistics as a portable indicator of performance. The figures, however, are internally consistent, and thus can be used to generate internally consistent data, where all that is required are values based on proportions of genuine numbers. Three simple analyses of two uniquely comparable sets of figures were presented here, showing that the populations engage with the digital content of web sites, partly in response to the nature of that content, and partly on the basis of existing working practices. It is clear that neither of these web servers, and by extension other similar servers, pose any immediate threat to conventional engagements with the archaeological record, but that they would appear to supplement that engagement for particular audiences.

Acknowledgements

This paper has benefited from the kind scrutiny of Tony Austin, Jo Clarke, Keith Westcott, Maureen Poulton, Julian Richards and Damian Robinson. The Archaeology Data Service is funded by The Joint Information Systems Committee and the Arts and Humanities Research Board.

References

ARNAULT, G., 1998. Internet chat rooms becoming a popular forum for business. In *New York Cyber Times*, January 26, http://www.nytimes.com/library/cyber/week/012698chat.html.

Cabinet Office 2000. *E-government: a Strategic Framework for Public Services in the Information Age*, http://www.cabinet-office.gov.uk/champions/Strategy.pdf, linked from http://www.iagchampions.gov.uk/Strategy.html.

CHERNY, L., 1999. *Conversation and Community: Chat in a Virtual World*, CSLI Publications, Stanford.

EU 2000. Communication from the Commission to the Council, the European Parliament, the Economic and Social Committee and the Committee of the Regions: *Towards a European Research Area*, http://europa.eu.int/comm/research/area/com2000-6-en.pdf, linked from http://europa.eu.int/comm/research/area.html.

EXON, S., DINGWALL, L., GAFFNEY, V., LAFLIN, S. AND Van LEUSEN, M., 1999. Archaeology in the Age of the Internet. *Computer Applications and Quantitative Methods in Archaeology* (CD-ROM).

HMSO 1999a. *Local Government Act 1999*, HMSO, London.

HMSO 1999b. *Modernising Government*. Presented to Parliament by the Prime Minister and the Minister for the Cabinet Office by Command of Her Majesty, HMSO, London.

HARNAD, S., 1991. "Post-Gutenberg Galaxy: the fourth revolution in the means of the production of knowledge", *Public Access Computer Systems Review* 2: 39-53.

HOBART, M.E. and SCHIFFMAN, Z.S., 1998. *Information Ages: Literacy, Numeracy and the Computer Revolution*, John Hopkins University Press, London and Baltimore.

HOLTORF, C., 1999. Is History going to be on my side? On the experience of writing and submitting a hypermedia. Ph.D. thesis. In *Internet Archaeology* 6, http://intarch.ac.uk/journal/issue6/holtorf_index.html.

House of Lords 2000. *Freedom of Information Bill*, House of Lords, London.

HUGGETT, J.W., 1995. Democracy, data and archaeological knowledge. In Huggett, J. and Ryan, N. (eds.), *Computer Applications and Quantitative Methods in Archaeology 1994*, BAR Int. Ser. 600: 23-26.

MARKUS, M.L., 1994. Finding a happy medium: explaining the negative effects of electronic communication on social life at work, *ACM Transactions on Information Systems* 12: 119-149.

SEU 2000. *National Strategy for Neighbourhood Renewal: a framework for consultation*, the Social Exclusion Unit, http://www.cabinet-office.gov.uk/seu/2000/Nat_Strat_Cons/default.htm, linked from http://www.cabinet-office.gov.uk/seu/index/publishe.htm.

SHENNAN, S.J., 1988. *Quantifying Archaeology*, Edinburgh University Press, Edinburgh.

WALTHER, J.B., 1992. Interpersonal effects in computer mediated interactions: a relational perspective, *Communication Research* 19: 52-90.

ŽIŽEK, S., 1999. Is it possible to traverse the fantasy in cyberspace. In Wright, E. and Wright, E. (eds.): *The Žižek Reader*, Blackwell, Oxford: 102-124.