

AGI Technical SIG Open Source Workshop: 11th March 2009

Oxford Archaeology: “Going Open”

Jo Cook, Senior IT Support and Development Officer, Oxford Archaeology/OA Digital

Senior IT Support and Development Officer for Oxford Archaeology (<http://thehumanjourney.net>), and also a consultant for OA Digital (<http://oadigital.net>), our geospatial consultancy arm. I blog online about computing, GIS and archaeology at www.archaeogeek.com.

What is Archaeology?

Archaeology is the study of past societies through their material remains, be that buildings, artifacts, skeletal or environmental remains. The purpose of archaeology is not, however, simply to understand how our ancestors lived, but to provide insight into how we got to where we are today. For example, are there lessons to be learned from how our ancestors dealt with (and sometimes caused) environmental change, disease, or warfare?

Introduction to Commercial Archaeology

In the UK, archaeology tends to be split into two camps- academic archaeology, practiced in universities, grant-funded, with a particular research question in mind, and commercial archaeology, which is developer funded and is a legal requirement of the planning process.

Commercial archaeology has been primarily linked to development since the mid 1980s, and particularly since 1990 when Planning Policy Guidance Note 16 (PPG16) was implemented. Whenever any building work is undertaken, it is a legal requirement to undertake at least some level of archaeological investigation into the site. The level of archaeological work required, from a simple desk-based assessment of the historic record for a site to a full-scale excavation, is decided by the County Archaeologist, at an early stage in the planning process. Commercial archaeological companies then bid to do the work. If, during that work, something unexpected shows up, then the building work can be postponed, leading to massive on-costs for the developers due to time penalties, equipment hire and so on.

The tension in this system is clear- the developers wish to see the work done as quickly and cheaply as possible, so choose the lowest bid for the work. The commercial units must keep their costs down as low as possible, yet this badly affects how they can treat their staff. While archaeologists tend to have at least a university degree, and usually a Masters, they are amongst the lowest paid graduates in the UK. Even in a good economic climate, contracts are normally short-term, and terms and conditions of employment are poor.

Our software and hardware requirements can be quite demanding. We work within a development chain that includes many other contractors, and gather and use many different types of data. We are obliged to work to strict deadlines and produce work not only of the highest quality, but also that survives in perpetuity often as the only record of the archaeology that we have examined. However, the financial realities of our discipline mean that we are not that attractive to commercial software companies.

Oxford Archaeology: the Company

Oxford Archaeology is the largest commercial archaeological unit in the UK, possibly also in Europe. We have approx 300 staff in three UK offices, and two French. This gives us the resources to work on large-scale development such as motorway widening and airport expansion as well as smaller scale work such as housing developments. We are also an educational charity, with a board of governors to ensure that we meet our educational remit on an annual basis.

Our size has pros and cons. We are able to provide better terms and conditions for our staff, yet our overheads are considerably greater than some of our smaller competitors, pricing us out of some sectors of the market.

Our Open Approach

Approximately 3 years ago we implemented what we call our "open ethos". The reasons for going down this route were many, but they boil down to both a financial impetus, and a desire to do better archaeology.

The "open ethos" has three strands:

- Open access to our data,
- Open standards for our file formats,
- Open source for our software.

Part of our remit is to store data about the sites that we work on, in perpetuity- or preservation by record. Our data is often all that remains of a site once it is dug up, and even if there are shiny objects in the local museum they are useless without the context that the archaeological data gives them. Not only that, but there is always the possibility that future archaeological research will provide a different interpretation of our data, re-writing the history books. Consequently choosing open standards for our software, and allowing open access to our data make this whole process easier, and more sustainable in the long-term.

We believe that the choice to use open source software adds to the long-term sustainability of our data. Furthermore, as there are no "black-boxes", our analytical process is transparent and reproducible. There are obviously other considerations to the decision to choose open source software too. We believe that this approach gives us more control, as we are not subject to changes in software licensing, or sudden changes in file format between one version of a package and the next. Clearly there are other financial considerations too, which will be discussed later.

Overview of the open source geospatial stack

From the GIS perspective we are trying to build a fully functional open source software stack that gives us the maximum control over our data, yet flexibility within our workflow. This is best achieved, we believe, by a modular approach, where each stage is addressed with a different software package. Each package can be chosen for suitability for a particular task, and can be replaced as necessary without jeopardising the rest of the stack. This is in contrast to the monolithic approach where one package attempts to do everything.

Data Layer

In our stack the data layer tends to comprise traditional file-based storage along with the geospatial databases PostgreSQL (<http://www.postgresql.org>), with the PostGIS spatial extension (<http://postgis.refractor.net>), and SQLite (<http://www.sqlite.org>), with the Spatialite extension (<http://www.gaia-gis.it/spatialite>). PostgreSQL is server-based and provides our main data store, whereas SQLite can run on smaller devices such as hand-helds or netbooks. We use a cross-platform, software-agnostic database synchroniser called SQLSync (<http://silvercoders.com/index.php?page=sqlsync>) to get data from the hand-helds to our main servers. This data can be accessed directly using desktop GIS or database packages such as Open Office Base (or even Microsoft Access).

Mid Layer: Map Servers

These convert data to web protocols such as WMS/WFS, but can also be used to style the data and convert to different spatial reference systems. We use geoserver (<http://geoserver.org>) as our map server, because of its excellent styling capabilities, but it also provides added extras such as allowing google to index the data. We serve base data such as background mapping to desktop and web-based GIS using WMS, and active data layers using WFS.

Desktop GIS

We use a combination of desktop GIS packages, depending on the nature of the work. The most common package that we use is GvSIG (<http://www.gvsig.gva.es/index.php?id=gvsig&L=2>), which was created by the Valencia Regional Council for Infrastructures and Transportation in Spain in 2003 as an ArcView 3.3 replacement. This has developed into a very solid and fully-featured package, which meets the majority of our GIS needs. Alternatively we use Quantum GIS, or QGIS (<http://www.qgis.org>), which is simpler, but has a plugin to allow people to use the GRASS GIS software (<http://grass.osgeo.org/>). We do occasionally use GRASS directly, but the steep learning curve is a problem.

Web-based GIS

Web-based GIS packages can be roughly categorised into those that require additional software to be installed on the web server and those that don't. Previously it was necessary to use the first category to get really good functionality, but this has changed. We are using two packages, depending on the nature of the work that we are doing- MapGuide Open Source, and OpenLayers. MapGuide Open Source (<http://mapguide.osgeo.org>) is an Autodesk product that was forked into an open source and a proprietary package in 2006. It does require additional software to be installed on the web server, but provides a management suite for managing the data and creating web layouts as part of the package. Development can be in a range of different languages, from php to java and .net. OpenLayers (<http://openlayers.org>) is, on the other hand, entirely javascript, and requires no additional software to be installed. All of the work is done within the web page itself, which means that configuration is simple, but requires some knowledge of javascript. Recently the mapfish framework (<http://mapfish.org>) has been developed, which uses extjs (<http://extjs.com>) to provide additional functionality around the OpenLayers map in a variety of programming languages. This is our package of choice for simple outward-facing mapping.

Against...

It would be foolish to try and convince people that there were no downsides to our approach. Convincing hearts and minds amongst our staff, who have little personal interest in the ethics of software licensing has occasionally been difficult, particularly when they see open source as being free and therefore "a bit dodgy". Convincing clients that our data will still be compatible with their software has also been a problem- particularly amongst people who believe that only ESRI software can create and use shapefiles.

TCO considerations have to be taken into account. Rolling out any new software is not a cost-free process, and staff have needed training in using packages that often look quite different to those that they are used to using. We have needed to invest heavily in IT staff to research, manage, and troubleshoot the new packages, and have needed to upgrade some hardware to cope with the change. However, we believe that most of these issues would have come up regardless of the software we had chosen.

Particularly in the desktop GIS market, the near monopoly of a small number of packages has meant that the open source competitors are still catching up to some extent. The production of high quality cartographic output is still difficult with many of the packages, and often additional illustration packages are needed to produce the quality of output that we need. However, as we can be more heavily involved in the development process, we can ensure that these issues are addressed in later versions.

For...

There have been many benefits to our approach. Since the software is free, we don't have to control who has a copy, or indeed who installs it on their computer at home. Also, we have been able to develop training courses in using the new packages. The increased investment in IT staff has meant that we have a core group of highly skilled people who know the packages intimately. The nature of the open source software world, with its communities of developers and users has meant that our profile has been raised significantly to the extent that we are often asked for advice on moving over to open source. Consequently we have formed a consultancy arm called OA Digital (<http://oadigital.net>), to take full advantage of this.

Hardware Diversification

One of the things OA Digital is investigating is the use of low-cost digital devices for on-site recording. Archaeologists have been grappling with replacing the traditional, time-consuming paper record for decades, without success, due to the need for extremely robust devices, connectivity, and ease of data entry. Now however, all archaeologists are proficient with mobile phones, and it is possible to get good enough screen resolution to make data entry possible on the small screen.

We are field-testing openmoko (www.openmoko.org), which are linux-based mobile phones, with built-in GPS plus the ability to plug in an external aerial, accelerometers, plus blue-tooth, WIFI and GPRS. They are also USB hosts, which means it is possible to plug a variety of USB devices into them, including digital cameras, extra hard drives, and keyboards. We are also experimenting with Otter Boxes to provide extra weather proofing and bump-protection.

Software Diversification

The openmoko will run a large number of programmes, including a mobile version of the GvSIG desktop GIS package, and the database package SQLite. This is the combination that we are building "DigiSite" around- a spatial database built around a core that is designed to be as flexible as possible, working for other disciplines outside of archaeology. This work is at an early stage, but we are working directly with the developers of GvSIG and GvSIG mobile, as well as developing DigiSite in-house.

We are also involved in the main GvSIG development process as our staff have been instrumental in finishing the english translation of what was predominantly a spanish package, as well as reporting and fixing countless bugs. For us this is ideal, as we get to shape the development of the software to meet our needs as well as contributing to the community.

More generally, it is clear that many people are put off moving towards an open source solution because they can appear less polished than their proprietary counterparts. Installation and configuration can be difficult, and sometimes the documentation and support forums are geared towards people with considerable experience already. While products like live cds and DVDs have existed for some time, they do not deal with the problems of installation and cannot be used in conjunction with the user's more familiar operating system and data- they are simply a sandbox for demonstration and teaching purposes. I personally have experienced the frustration of this, and have come up with the Portable GIS concept (<http://www.archaeogeek.com/blog/portable-gis/>) that allows users to load a pre-configured software stack onto a USB stick or their hard drive, and have it "just work" with the bare minimum of setup (you simply have to tell it which drive letter you are using).

As well as being popular with new users, other uses for this have been suggested, from disaster response kits where no internet connection is available (the kit could be pre-loaded with data) to demonstration setups for people trying to persuade their bosses that open source is a good way to go without needing to invest time and money setting it up. Version 2 of the package will be out shortly, with the addition of paid-for support and customisation for people wishing to use it in a business environment.

Conclusions

Our attempt to "Go Open" is still ongoing, and it's fair to say that it's a learning process, and not easy. However, we think the long-term result will be a better software stack for our staff, more efficient and flexible workflows, as well as some interesting diversification opportunities, and hopefully a hugely successful consultancy arm!

Thanks and where to go for more information

This talk and notes are freely available under a Creative Commons License on my website (<http://www.archaeogeek.com/blog/talks/>). Please don't hesitate to get in touch with me personally, or with OA Digital if you have any questions or are interested in "going open" yourselves.

Joanne Cook
Senior IT Support and Development Officer
Oxford Archaeology/OA Digital
j.cook@thehumanjourney.net
+44 (0)1524 880212

