University of Reading

Sharing and visualizing environmental data using Virtual Globes

Conference or Workshop Item

Accepted Version

Blower, J. D., Gemmell, A., Haines, K. ORCID: https://orcid.org/0000-0003-2768-2374, Kirsch, P., Cunningham, N., Fleming, A. and Lowry, R. (2007) Sharing and visualizing environmental data using Virtual Globes. In: UK e-Science All Hands Meeting 2007, 10-13 September 2007, Nottingham, UK, pp. 102-109. Available at https://centaur.reading.ac.uk/803/

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

Published version at: http://www.allhands.org.uk/2007/proceedings/papers/792.pdf

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the End User Agreement.

www.reading.ac.uk/centaur

University of Reading

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Sharing and visualizing environmental data using Virtual Globes

Jon Blower[1], Alastair Gemmell[1], Keith Haines[1], Peter Kirsch[2], Nathan Cunningham[2], Andrew Fleming[2], Roy Lowry[3]

[1] Reading e-Science Centre, Environmental Systems Science Centre, University of Reading, RG6 6AL, United Kingdom
[2] British Antarctic Survey, High Cross, Madingley Road, Cambridge, CB3 0ET, United Kingdom
[3] British Oceanographic Data Centre, Joseph Proudman Building, Liverpool, L3 5DA, United Kingdom

Abstract

Virtual globe technology holds many exciting possibilities for environmental science. These easy-to-use, intuitive systems provide means for simultaneously visualizing four-dimensional environmental data from many different sources, enabling the generation of new hypotheses and driving greater understanding of the Earth system. Through the use of simple markup languages, scientists can publish and consume data in interoperable formats without the need for technical assistance. In this paper we give, with examples from our own work, a number of scientific uses for virtual globes, demonstrating their particular advantages. We explain how we have used Web Services to connect virtual globes with diverse data sources and enable more sophisticated usage such as data analysis and collaborative visualization. We also discuss the current limitations of the technology, with particular regard to the visualization of subsurface data and vertical sections.

1. Introduction

E-Science is "connected science". Its philosophy is to connect scientists with data and compute resources, with each other and with non-academic partners in industry, government and the general public. Environmental science has a strong need for applications of the e-Science philosophy, being characterized by highly interdisciplinary work using large datasets and powerful computing resources, and by a strong need to engage stakeholders outside the scientific community. Public interest in environmental science has never been higher and so it is vital to find methods to communicate science effectively to non-specialists.

The broad range of sub-disciplines within environmental science require common languages to enable the exchange of data and ideas. Considerable effort is underway world-wide to solve this problem at a number of levels, from the low-level technical details of digital file formats [7, 17] and precise metadata [16] to the semantics of high-level Web Services for geospatial data access [19]. The effort required to conform to these standards, however, is currently considerable and beyond the technical reach of most scientists, who require assistance from experts such as dedicated data centres to ensure that the data they produce are interoperable. There is therefore a clear need for simple formats that all scientists can use to publish their data (or, at least, meaningful subsets thereof). Scientists must be able to see rapid benefits from the use of standards if they are to adopt them.

The process of "doing science" usually results in the production of some visualization of data for formal publication or informal discussion. Such is the plethora of visualization packages in common use by environmental scientists that visualizations are most commonly shared in basic "common denominator" image formats such as GIF and PNG. Unless these are properly scaled and georeferenced, images from different scientists cannot easily be overlain and intercompared. Such simultaneous visualization of data from different sources (often residing in different physical locations) has enormous potential to reveal new information and generate hypotheses to be tested in future research. Geographic Information Systems (GIS) are designed for precisely this purpose and have many sophisticated capabilities; however, they have historically been complex, expensive and mutually incompatible.

Recently there have been great advances in the development of "virtual globes": software

applications that display a three-dimensional representation of the entire Earth, usually based on satellite imagery, upon which new information can be superimposed. Virtual globes and associated technologies have great potential to support and promote environmental science and act as easy-to-use, lightweight GIS systems. There is a clear (and rapidly-growing) interest in this technology across the environmental sciences, as demonstrated by a recent workshop [23] at the National Institute for Environmental e-Science, from which much of the overview information in this paper was obtained[1].

The purpose of this paper is to present our own work in this field and to review the current state of virtual globe technology. We start this paper (section 2) with a short review of some popular virtual globe applications and their key capabilities. In section 3 we discuss the advantages and limitations of Keyhole Markup Language (KML) [13], a common data format for virtual globes. In section 4 we describe some ways in which we are using virtual globes in our own scientific work. We follow this with an analysis of the important current issue of working with four-dimensional data (section 5). In section 6 we show how we use Web Services to enhance the capabilities of virtual globes.

2. Overview of Virtual Globes

Virtual Globes allow the presentation of data at a huge range of spatial scales, from the sub-metre to the planetary scale. Although capabilities vary from product to product, VGs provide:

- Support for simple file formats for data exchange.
- The capability to display multiple datasets simultaneously.

Virtual Globes are designed to be very easy and intuitive to use: in fact, many (notably Google Earth [10]) are aimed at the general public. This paper will investigate the extent to which current Virtual Globes are useful to professional environmental scientists who already have access to sophisticated systems for data analysis and visualization.

At the time of writing there are at least 30 Virtual Globe (VG) applications in existence. A description of the relative merits of all of these applications is well beyond the scope of this paper, so we shall focus our attention on three applications: Google Earth, NASA World Wind and ArcGIS Explorer. These were the main VGs that were covered in the aforementioned workshop.

Google Earth is perhaps the most well-known VG application currently available. It is a free but closed-source application that is available for Windows, Linux and Mac OSX. The primary method for visualizing data in Google Earth is to create KML files, although there is limited support for other methods. It is aimed at the general public, primarily as a search tool, but has attracted a very large community of people who have used the application for a very wide range of purposes.

NASA World Wind [14] is an open-source VG application that is written in .NET and is therefore only supports Windows operating systems. Future releases will be Java-based and hence cross-platform. World Wind provides access to a wide range of NASA satellite imagery. Data can be imported through tile servers, OGC Web Services and there is limited support for KML. Its focus is toward scientific users, so World Wind has a more specialist community than that of Google Earth.

ArcGIS Explorer [1] can be used as a standalone (free, Windows-only) VG but it is best seen as a lightweight client to the (non-free) ArcGIS Server. It can import data in a very wide range of GIS formats (including KML) and can perform data analysis on the client (using plug-ins and the associated .NET SDK) or the server (through the interface to ArcWeb Services). At the time of writing this is a very new product and so has not had time to build a large community.

In this paper we shall focus on the use of Google Earth in environmental e-Science, although we will continue to discuss other VGs where appropriate. This is because (1) Google Earth is the only VG of the above shortlist that currently works across platforms and thus has the widest potential scientific audience, (2) Google Earth currently possesses the largest community, and (3) Google Earth supports the latest features of KML, many of which are particularly relevant for scientific data.

3. Keyhole Markup Language

KML is a simple XML language that is supported by many virtual globes. Although historically tightly connected with Google Earth, discussions are underway to standardize KML the Open Geospatial Consortium (OGC) [19]. KML is supported by a number of virtual globes and other GIS systems and is therefore already becoming a *de facto* standard. (Note that KML files can be combined with other

[1] Also see http://www.scispace.net/geobrowsers

supporting files such as imagery in a zip archive, producing a KMZ file. In this paper we intend "KML" to represent both KML and KMZ for clarity.)

From an environmental science point of view, KML is a somewhat limited language. It describes only simple geometric shapes on the globe (points, lines and polygons) and is not extensible. It is, in many respects, analogous to Geography Markup Language 2.0 [7]. GML 3.x is much more sophisticated and allows the rich description of geospatial features such as weather fronts and radiosonde profiles [22].

Computer science purists might also regard KML as a rather imperfect standard, primarily because it conflates information concerning geographic positioning with style information (icons etc) and virtual camera control in the same file, in the same way that badly-written web pages conflate structure (HTML) with style (CSS) and dynamic code (Javascript). Annotations of KML features are intended to be human-readable plain text or simple HTML, not machine-readable XML.

For the above reasons, KML is currently not suitable as a fully-featured, general-purpose environmental data exchange format. Its rapidly-growing adoption by scientists, however, shows that it is finding an important niche. It is important to remember that virtual globes (and KML) do not attempt to replace more sophisticated systems. They make it easy for non-technical scientists to share and visualize simple geospatial information, which can then be manipulated in other applications if required. The objections raised in the previous two paragraphs can actually be seen as advantages from the point of view of usability because they make it easier for users to visualize their data quickly using a single, simple data file. KML is useful because it spans a gap between very simple (e.g. GeoRSS [8]) and more complex (GML 3.x) formats. It is a much more useful interchange format than imagery alone since it holds georeferencing information and allows for the inclusion of links to related information (e.g. the full data archive).

The most recent version of KML (version 2.1) contains many new features that are particularly relevant to scientific data, such as large data support and the ability to timestamp features and hence create animations. Google Earth is currently the only VG that fully supports all of these features. Detailed documentation and tutorials on all KML's capabilities can be found on the web [13] and will not be reproduced here.

4. Scientific use cases

We shall now discuss a number of scientific use cases for virtual globes, illustrated by concrete examples from our own work and from the environmental science community.

4.1 Simultaneous visualization of data from different sources

A key strength of virtual globes is that it is very easy to download data over the Web from a number of different providers and visualize them simultaneously to look for relationships that can be investigated later in a more quantitative fashion. Environmental data are now available in KML format from a number of projects and data centres, including the Reading e-Science Centre, the DAMOCLES project [3], the National Snow and Ice Data Centre (NSIDC) and many more.

Figure 1 shows the results of one such combination of two independent datasets. The first dataset is a timeseries of images representing the sea surface temperature as daily snapshots (red colours are warm, blue colours are cold) from the UK Met Office's global FOAM ocean model [5]. The KML file was generated automatically by the Reading e-Science Centre's Godiva2 website [9], which provides interactive visualization of a large number of environmental datasets, with an option to export imagery in KML format. The second dataset is a timeseries of Hurricane Katrina's location and intensity and is from a TRACK [11] analysis of data from the European Centre for Medium-Range Weather Forecasts' operational analysis system. The TRACK analysis is due to Lizzie Froude and Kevin Hodges of the University of Reading (the TRACK output was converted to KML using a simple Python script). The hurricane's intensity at each timestep is given by the colour of the icons representing Katrina's location (red colours are most intense). The screenshot in Figure 1 clearly shows that the passage of Katrina causes the ocean to cool on the right-hand side of the cyclonic storm track[2]. This is due to the strong winds causing upwelling of colder subsurface water.

There are countless other potential applications for such simple combinations of datasets. For example, authors at the British Antarctic Survey use Google Earth to

[2] Video (linked from New Scientist article): http://www.newscientisttech.com/article/dn11773-virtual-earths-let-researchers-mash-up-data-.html

simultaneously visualize (in near real-time) the position of predators such as penguins in Antarctic waters alongside satellite measurements of chlorophyll concentration. This clearly shows a relationship between the penguins' movements and the location of high nutrient concentrations, which in turn control the locations of the penguins' prey (fish). This information is invaluable for directing scientific missions.

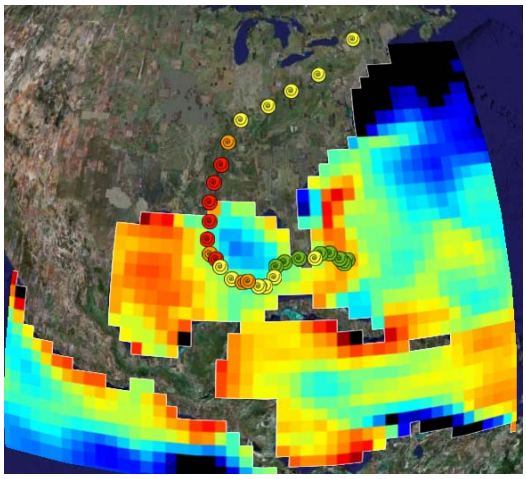

Figure 1: Google Earth screenshot showing the cooling of the surface waters of the Gulf of Mexico due to the passage of Hurricane Katrina in August 2005. The storm track data (points) and temperature data (image) were obtained from independent providers.

Although the above pieces of work could also be achieved using a number of other systems, KML and Google Earth provide a very simple means of exchanging and visualizing data that requires little experience to produce useful results. The adherence to common formats is key to such intercomparisons.

4.2 Diagnosis of numerical models

In the environmental sciences it is common to run large numerical models of the Earth system, into which various observations have been assimilated. It is important to be able to diagnose the numerical model by checking that the model does not depart too far from reality. It is also important to check that the observations are "sane". Since many of these models are global in scale, it is natural to use a virtual globe to display the model output alongside the observations.

Scientists in the National Centre for Ocean Forecasting (NCOF [15]) have a need to diagnose the new NEMO ocean model [18] and compare its results with observations from moored buoys, ARGO floats and ships. Even simple visualization tasks are technically challenging because of the large number of observations (around 9,000 sets of around 150 measurements each per month) and the irregular numerical grid that NEMO employs. We have solved these problems by (1) creating a Web Service that allows the user to select subsets of observations in KML format, and (2) using our Godiva2 website [9] to generate images of the model results in a standard geographic map projection, also in KML format.

The scientists need to calculate the goodness-of-fit between the observations and the models, which may in general be held at different physical locations. Google Earth is essentially a visualization tool and cannot itself perform such data analysis tasks (see section 6 below). Therefore we perform the goodness-of-fit calculations in the Web Service that produces the KML representing the observations. The placemarks that indicate the location of each observation are coloured according to the fit of the model and that observation: good matches are represented by green pushpins and bad matches by red pushpins.

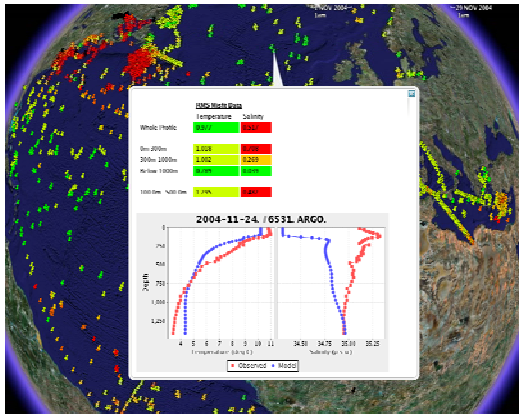

Figure 2: Diagnosis of a NEMO model run by comparing the model output with its assimilated observations. The left graph represents temperature, the right graph represents salinity. For each graph, the red trace shows the observations (in this case from an ARGO float) and the blue trace shows the model results. Note the large concentration of red pushpins (representing bad fits of model to data) in the highly-variable Gulf Stream region (top left of picture).

When the scientist clicks on a particular placemark a balloon appears that shows the temperature and salinity profiles of that observation and the model results at that location (Figure 2). The Web Service allows

the contents of this information balloon to be tailored by the user. Due to the potentially-large number of placemarks this information cannot be stored in the KML file for every point: instead this information is generated dynamically by our Web Service and linked from the KML file. The link is activated (and the Web Service is called) when the scientist clicks on a pushpin.

A virtual globe is an appropriate platform for this activity because it allows the scientist to view the data at a huge range of scales. Sometimes a global view is appropriate to gain an overview and sometimes the scientist needs to zoom in tightly to a region in which the observations are very closely spaced.

4.3 Delivery of real-time data

Real-time data are very important in the environmental sciences. Many operational centres maintain networks of observing systems that must be continuously monitored. Real-time forecasts from numerical weather prediction models are of immense value to government and industry for myriad applications including healthcare and marine emergency response (the DEWS project [4] is developing a system to deliver such information in a secure environment).

(Near) real-time data can be displayed in Google Earth using the NetworkLink facility in KML. This facility allows all or part of a KML dataset to be refreshed at a fixed rate or based on the expiration time given in the HTTP header. A user can therefore download a fixed KML file containing the NetworkLink and Google Earth will automatically refresh the link, ensuring that the user always sees the latest information.

This technique was used by the authors at the British Antarctic Survey (BAS) to monitor and direct two scientific cruises by the BAS ship James Clarke Ross in 2006. Multiple data streams including ship location, sea surface and air temperature, ocean salinity and conductivity and wind speed were streamed from the ship to a server that was polled at regular intervals from Google Earth via the NetworkLink. These up-to-date local measurements were combined with information about the wider environment to allow timely decisions to be made by scientists both on and off the ship.

4.4 Monitoring of observing systems

Authors at the British Oceanographic Data Centre have adopted Google Earth as a spatial metadata quality control tool and as an access portal for detailed metadata documenting the holdings in the National Oceanographic Database. "Light" metadata (position, date/time, data type) plus a URL to detailed metadata are exported using a simple application into a KML file. The interactive browse functionality of Google Earth is then used to quality control the spatial metadata through visualization, such as identifying oceanographic data erroneously located on land. Previously this function had been done by automated comparison against a 1-minute land/sea gridded dataset and the inadequacies of this method are currently being exposed through a Google Earth-based audit. Google Earth provides additional spatial quality assurance through the display of spatial patterns in related data such as the line formed by the individual profiles of a ship track. The KML file (i.e. the "light" metadata) contains a link to the metadata server, enabling much more detailed metadata to be viewed in Google Earth's embedded web browser. This makes Google Earth an ideal metadata browsing framework when making scientific assessments of data collections.

This illustrates the value of using lightweight tools to access simple data and metadata, using links to allow the user to drill down to more complex information. The system just described was developed in under a week.

5. Four-dimensional data and Virtual Globes

A key requirement for environmental scientists is to be able to visualize four dimensional data (i.e. time-dependent three-dimensional data). Existing GIS systems have often been limited to "two and a half" dimensional data, i.e. static map data that is tied to the land surface.

Many virtual globes answer this need to some degree. Google Earth has good support for the time dimension (allowing animations to be created) and allows data points and horizontal images to be placed above the Earth's surface. However, there is currently no support for displaying images as vertical sections, or for displaying subsurface data; both are key limitations for oceanographers and geologists. Subsurface and vertical section data are topics of ongoing research in the NASA World Wind community. ArcGIS Explorer has full support for subsurface data but is a relatively new product with a currently-small community.

Due to the popularity and ubiquity of Google Earth there is very strong community

interest for full support for 4-D data. Through the NIEeS workshop [23] and other forums such as the American Geophysical Union the community is feeding requirements to Google but the future is still unclear in this respect, illustrating an important disadvantage of using closed-source software.

Some progress can be made, however, through workarounds to the problems of vertical section and subsurface data. Vertical section data can be displayed by taking advantage of Google Earth's ability to display 3-D Collada [2] models: it is possible to create a simple 3-D model consisting of a vertical wall and to project an image onto it [6]. However, this is not an ideal solution because (1) it is rather convoluted, and (2) the 3-D modelling environment uses Cartesian coordinates rather than geographic coordinates, meaning that depth sections do not wrap the globe properly.

Subsurface data can be displayed meaningfully by an ad-hoc change to the vertical datum: for example, for oceanographic data one pretends that the sea surface is actually the sea floor and displays imagery and data at an appropriate altitude *above* sea level. The sea surface can be draped with an image (or even a three-dimensional representation) of the sea floor to enhance the illusion. This means that Google Earth's measurement of altitude will be wrong, but it does give the user potentially greater control over the vertical exaggeration (limited to 0.5 to 3 in Google Earth, which is far too narrow a range for general scientific visualization). An example of this is shown in Figure 3, in which we illustrate the three-dimensional structure of the Gulf Stream.

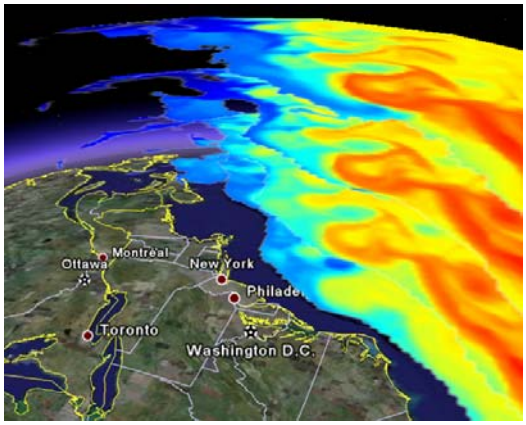

Figure 3: Three dimensional structure of the Gulf Stream, working around Google Earth's limitations regarding subsurface data by displaying four layers of ocean imagery above ground.

6. Virtual Globes and Web Services

The capabilities of virtual globes can be augmented through Web Services. The primary uses of Web Services in this context are to provide data, data analysis capabilities and interactive control. Few (if any) virtual globes have native support for the SOAP protocol [20] (or even the HTTP POST operation) and so such Web Services must usually be constructed with HTTP GET bindings so that the entire operation can be encoded in a URL. Different client and server systems have different limits for the maximum length of URL that they support so there is a limit on the richness of queries that can be performed using GET. Limitations such as this can be worked around by using the virtual globe in tandem with a website, as described in section 6.2 below.

6.1 Obtaining data through standard Web Services

The Open Geospatial Consortium (OGC [19]) publishes a suite of specifications for standard Web Services for accessing geospatial data. There are three services of particular relevance here: the Web Map Service (WMS) serves map images, the Web Feature Service (WFS) serves feature data (in GML) and the Web Coverage Service (WCS) serves "coverages", which are usually georeferenced rasters of data. Of these, the WMS is perhaps the most relevant and is also the simplest since it serves fully-rendered images that do not have to be further interpreted. (Support for WFS and WCS in virtual globes is rare, since these services produce data that must be rendered on the client. It is likely, however, that future versions of the WFS specification will support KML as an output format for simple features.)

Support for WMS in virtual globes is variable: ArcGIS Explorer 9.2 supports WMS version 1.1.1, NASA World Wind supports WMS 1.1.1 and 1.3.0 (although it does not currently support the vertical dimension) and Google Earth possesses only very basic WMS support. However, a WMS request can be encoded in a single URL and so a request for a map image can be embedded in a KML document just as easily as a link to a static image. The disadvantage with such basic WMS support is that the user has to take responsibility for reading the WMS server's metadata in order to construct the URL (but see section 6.2 below to see how this can be improved). Currently Google Earth and most other virtual globes can only display images rendered in standard Plate Carrée (linear longitude-latitude) projection.

We have developed a Web Map Service implementation (ncWMS [21]) that serves dynamically-generated images from CF-NetCDF [16] data. This focuses on the fast generation of map images for interactive data exploration on geo-websites such as Godiva2 [9] and in virtual globes. It contains many enhancements for scientific data, such as the ability to create animations, the ability to query data values at individual pixels (using the WMS GetFeatureInfo operation) and the ability to create timeseries plots. This software supports KML as an output format, enabling the direct generation of georeferenced images that can be viewed in virtual globes without further intervention (in this case, the KML "wraps" the image and provides georeferencing information).

There is much ongoing research concerning the creation of workflows from OGC services. It is possible to create workflows in which data are read from a WFS or WCS, then converted to a virtual globe-friendly format (such as an image or a KML file) using another service such as a WMS. (Intermediate data processing can be handled by another emerging OGC specification, the Web Processing Service.) Among the current challenges is the problem of how to create such distributed workflows efficiently, given that data volumes can be very large (terabytes would not be uncommon).

6.2 Extending functionality using Web Services and interactive websites

Some virtual globes provide means to have their capabilities extended through custom plug-ins. ArcGIS Explorer can be extended through a .NET Software Development Kit (SDK). NASA World Wind is open source and future versions of NASA World Wind will be written in Java and based on an SDK framework, allowing easier customization and even embedding of the World Wind application in other pieces of software, including websites.

Google Earth is a closed-source application and does not currently have a mechanism for the community to develop plugins to add functionality beyond that provided by KML. However, as illustrated in section 4.2 above, it is possible to interface Google Earth with remote Web Services that perform calculations on data, returning results in KML (for simple geographic features) with embedded HTML (for display of further information and links).

The functionality of Google Earth can be enhanced still further by combining many of the ideas discussed previously in this paper. We have created a demonstration system that enhances the WMS support in Google Earth by combining a periodically-refreshing NetworkLink (see section 4.3 above) with an interactive website that is displayed in Google Earth's built-in web browser to effectively provide extra controls and displays. Although still an early prototype, this system has the potential to turn Google Earth into a fully-featured WMS client. The basic mechanism is as follows:

1. The user downloads a simple KML file that contains a NetworkLink. This link points to a fixed-name target KML file on a web server and reloads this target file periodically (say, once a second).
2. The initial KML file also contains a URL to a web application, whose user interface opens automatically in Google Earth's built-in web browser (or an external browser if preferred). This interface provides the input fields and buttons required to interact with remote Web Map Services.
3. In response to user interaction via the browser (e.g. the selection of a particular WMS layer to display), the web application updates the target KML file on the web server with new information (e.g. the name of the new WMS layer).
4. The display on the virtual globe itself is automatically updated with the new target KML file when the NetworkLink refreshes itself.

This provides a simple but powerful general mechanism for creating interactive web applications that use Google Earth as a display engine. The same idea could be extended to any graphical application (not just a web browser) that can write new KML files to a location that is accessible by Google Earth. Collaborative visualization is easy to achieve by this method, simply by downloading the initial KML file (containing the NetworkLink) to multiple Google Earth clients: each client will be automatically updated with the same information when the target KML file is changed. However, this method is currently limited to information refreshes at rates of no more than one update per second.

7. Discussion

The key strengths of virtual globe applications are their easy-to-use, intuitive nature and the ability to incorporate new data very easily. Generating a simple KML file, for example, is no more difficult than creating a simple web

page and so it is well within the technical capabilities of most scientists to share their data in this format. We anticipate that this will lead to much more sharing of data (beyond static images) and enable new and exciting science to be done.

Although the basic capabilities of virtual globes are relatively simple, we have shown that these capabilities can be enhanced through interfaces with Web Services: however we note that this is much more straightforward if those Web Services support simple HTTP GET interfaces. This illustrates that communities are best engaged by supporting the simplest and most widespread web technology, giving the greatest chance of stimulating new applications. Similarly, although Google Earth is not the most technically-sophisticated virtual globe available, the simplicity of its user interface and data format has attracted a large community which is working together to provide creative solutions to common problems.

We have not discussed the issue of security in this paper. Scientists will certainly have a need to use virtual globes to visualize data that is protected by some kind of access control. It is unlikely that virtual globes will support complex or bespoke security mechanisms in the near future. Therefore we have a strong need for security frameworks that are fit for purpose and compatible with the simplest web architectures (e.g. a client that only supports HTTP GET, without support for cookies). Research is clearly needed here, but perhaps emerging frameworks such as HTTPsec [12] will be helpful.

We have also not discussed the great potential virtual globes have for education and outreach. Simple, eye-catching tools that are available to the general public can help enormously to communicate scientific research to non-specialists. This will no doubt be the subject of much future work.

Acknowledgements

The authors would like to thank all the participants and organizers of the NIEeS workshop for stimulating discussions and inspiration. JB, AG and KH acknowledge support from the NERC Reading e-Science Centre grant, the DEWS project and the National Centre for Ocean Forecasting.

References

[1] ArcGIS Explorer, http://www.esri.com/arcgisexplorer/
[2] Collada, http://www.collada.org/
[3] DAMOCLES project, http://www.seaice.dk/damocles/
[4] Delivering Environmental Web Services (DEWS) project, http://www.dews.uk.net/
[5] Forecasting Ocean Assimilation Model (FOAM), http://www.metoffice.gov.uk/research/ncof/foam/index.html
[6] Geens, S., Proof of concept: visualizing vertical data in Google Earth, http://www.ogleearth.com/2006/07/proof_of_concep.html
[7] Geography Markup Language, http://www.opengis.net/gml/
[8] GeoRSS, http://www.georss.org
[9] Godiva2 website, http://lovejoy.nerc-essc.ac.uk:8080/ncWMS/godiva2.html
[10] Google Earth, http://earth.google.com
[11] Hodges, K.I., Adaptive Constraints for Feature Tracking, *Monthly Weather Review*, **127**, 1362-1373, 1999
[12] HTTPsec, Public key authentication for HTTP, http://www.httpsec.org/
[13] Keyhole Markup Language, http://earth.google.com/kml/
[14] NASA World Wind, http://worldwind.arc.nasa.gov/
[15] National Centre for Ocean Forecasting, http://www.ncof.gov.uk/
[16] NetCDF Climate and Forecast (CF) Metadata Conventions, http://www.cgd.ucar.edu/cms/eaton/cf-metadata/CF-1.0.html
[17] NetCDF, Network Common Data Form, http://www.unidata.ucar.edu/software/netcdf/
[18] Nucleus for European Modelling of the Ocean (NEMO), http://www.lodyc.jussieu.fr/NEMO/
[19] Open Geospatial Consortium, http://www.opengeospatial.org/
[20] SOAP, http://www.w3.org/TR/soap/
[21] Web Map Service for NetCDF data, http://www.resc.rdg.ac.uk/trac/ncWMS/
[22] Woolf, A., B. Lawrence, R. Lowry, K. Kleese van Dam, R. Cramer, M. Gutierrez, S. Kondapali, S. Latham, D. Lowe, K. O'Neill, A. Stephens, Data integration with the Climate Science Modelling Language, *Advances in Geosciences*, **8**, 83-90, 2006
[23] Workshop: "Google Earth and other geobrowsing tools in the environmental sciences", National Institute for Environmental e-Science, 2-3 April 2007, http://www.niees.ac.uk/events/GoogleEarth/index.shtml