

Using Keyhole Markup Language to create a spatial interface to South African water resource data through Google Earth

Michael Silberbauer¹, Willie Geldenhuys²

Resource Quality Services, Department of Water Affairs and Forestry, Private Bag X313, Pretoria, South Africa 0001, ¹SilberbauerM@dwaf.gov.za ²GeldenhuysW@dwaf.gov.za

Abstract

The rainfall in South Africa is in the low to moderate range (around 500 mm per year) and the runoff generated is less than 50 mm per year, with evaporation exceeding 2000 mm per year in most of the country. By the middle of the twentieth century, urbanisation, industrial development and mineral exploitation were generating so much effluent that the parched hydrological system had no further capacity for dilution or assimilation of contaminants before the water arrived at the next user downstream. In order to monitor and manage this situation, which becomes critical during dry spells, the government institution currently known as the Department of Water Affairs and Forestry established a nationwide network of water quality monitoring sites, many of which have now been operating for thirty years or longer. In the early years of the programme, the network layout and operation were part of the corporate consciousness, with only occasional hard copy reports of the entire system to jog people's memories. Most of those involved in the design and early operation of the network have long retired, and new users of the data have difficulty in appreciating the extent and interactions of monitoring programmes and sites. Soon after the department produced its last monumental hand-drawn water quality map in 1984, geographical information systems (GIS) became available and greatly simplified data presentation. The next fundamental shift in spatial data methods occurred in 2005, when Google launched Google Earth on the Internet, with the Keyhole Markup Language (KML) as a method for users to overlay data. We used these tools to set up a publicly accessible spatial inventory on the departmental Internet web site <URL <http://www.dwaf.gov.za/iwqs/wms/data/000key.htm>>, giving rapid access to several thousand surface water and groundwater monitoring site locations. Using the KML popup balloon function, users can select pre-generated time-series graphs of water quality, flat data files and flow data records. Maucha ionic symbols show the major ions and salinity at sites where such data are available. Date tags for the opening and closing of monitoring sites enable a user to show monitoring activities for selected periods and the hierarchical folder layout permits the selection of sites related to particular types of monitoring, such as sewage works or rivers. The advantages over a conventional GIS system are Google Earth's ability to zoom rapidly from national to local scale, the frequent update of the background layers and fly-through perspective views where topography emphasises the inter-relationship of pollution sources, monitoring sites and water users. To help orientate users, we have provided supplementary background data, including South African river names and drainage region codes in KML format. Limitations of the system include the need for a more-than-entry-level PC with an Internet connection and the sparking of unreasonable user

expectations regarding elevation resolution, real-time satellite images and real-time water quality data.

Integrating the water quality inventory with Google Earth has provided users of water quality data with a new way of visualising the South African data network and of sifting the inventory for useful data. Despite the limited potential audience, the site logs more than 100 visits per month. Future developments could include a tighter integration with corporate systems, including the live water quality database and the use of time-enabled markers to show how water quality at each site measures up to the applicable standards and guidelines from year to year.

Introduction

To appreciate the importance of South Africa's water quality monitoring network, and the need for a spatial inventory, some background information on the region's climate, hydrology and economic development is necessary.

Rainfall over South Africa is in the low to moderate range with an average of about 480 mm per year (Figure 1a). Annual evaporation rates increase from the southeast to northwest and mostly exceed 2000 mm (Figure 1b). The runoff generated from rainfall is generally less than 50 mm per year, representing a yield of 9% against the world average of 35% (Allanson 2004; Schultz 2006; Schultz & Lynch 2006). Rainfall is also very seasonal and river flow is often low or non-existent in the dry season, even in major rivers such as the Limpopo.

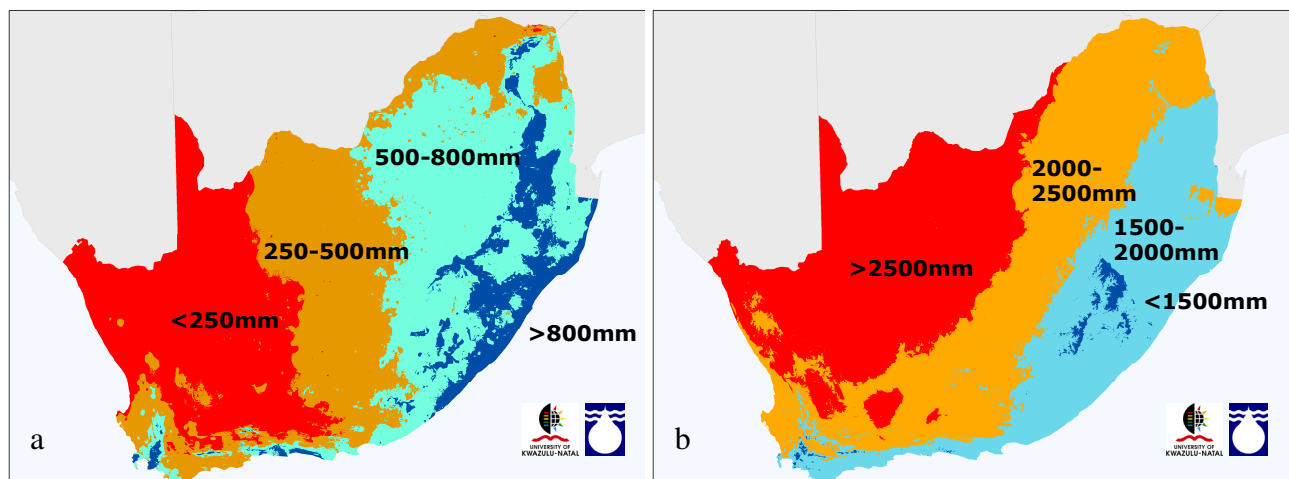


Figure 1. a) Mean annual rainfall for South Africa (data from Schulze & Lynch 2006); b) annual evaporation for South Africa (data from Schulze & Maharaj 2006).

This hydrological regime severely limits the capacity of South African surface waters to absorb the impact of the pollutants that are an inevitable consequence the country's tremendous growth in agriculture, mining, industry and urban population during the past century. Scientists sounded a warning at the start of the 1960s that the quality of the streams draining Johannesburg was already

poor and that the main impoundment to the north of the city, Hartbeespoort Dam, had become eutrophic because of high nutrient inputs (Allanson 1961; Allanson & Gieskes 1961).

In response to concerns about water quality, the government department that administers national water resources (now known as the Department of Water Affairs and Forestry, DWAF) stepped up the national sampling effort during the 1960s. The number of samples and the types of analyses soon became too much for a conventional wet chemistry laboratory, and DWAF established a partially automated laboratory in the mid-1970s, with a laboratory information management system feeding data to the national hydrological database. Procedures have since undergone many updates and the data now reside on the Informix-based Water Management System (WMS). The WMS includes tools for database administration, monitoring programme management and report generation (in the form of tables, flat files and maps created in ArcView 3.2 with Avenue). The national network now consists of many thousands of sites, some sampled routinely and others in response to events. DWAF's Resource Quality Services directorate administers the network under a number of national programmes dealing with water chemistry, eutrophication, environmental health and microbiology.

As the database grew, DWAF water quality managers saw the need for an easily accessible inventory of monitoring sites and data, together with a facility to produce tabular summaries, maps and data sets for water quality modelling. A team of aquatic scientists and technicians hand-drafted a comprehensive map of the Vaal River basin water quality from 1979-1983 (van Vliet & Nell 1986) and a massive paper inventory of 1652 sites using a combination of mainframe database and personal computer Pascal (Swart, van Veelen & Nell 1991). A parallel report on water quality at 683 sites also included the first GIS water quality maps (van Veelen, Nell & Geldenhuys 1990). During the 1990s, rapid developments in GIS, remote sensing, desktop computing power and data distribution via the Internet enabled the development of graphical interfaces to water quality data, such as the ArcInfo-based WaterMarque (Cobban & Silberbauer 1993). DWAF staff also took the first steps in publishing data on the web.

By 2005, the maturing and convergence of the technologies for GIS, image processing and the Internet had resulted in a number of global earth viewers becoming widely available, notably the Keyhole Earth Viewer system released as Google Earth. The convincing perspective relief view of satellite imagery draped over a digital elevation model of the surface of the Earth resembled a computer-game rather than a GIS and attracted the attention of a much wider audience than previously exposed to GIS. The facility for customising the interface with Keyhole Markup Language (KML) made the system useful to GIS specialists as well.

Methods

This section describes the general processes followed in using Google Earth's extensible markup language, KML, for the construction of a visual monitoring site inventory with metadata pertaining to location, data available and period of activity. For details of scripts and code, please see ###.

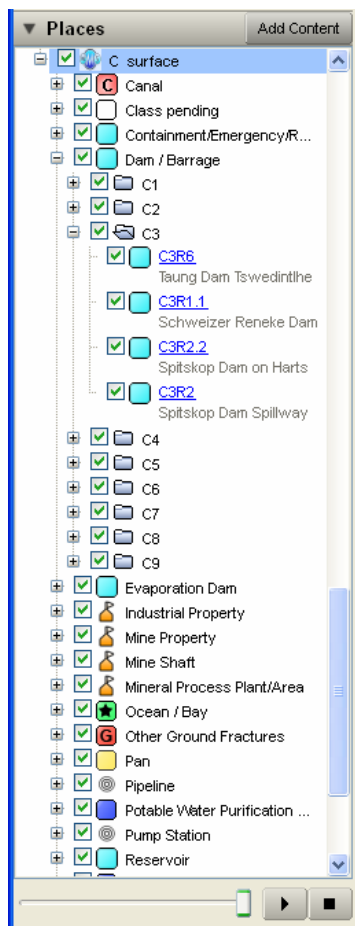
Data extraction and preparation

We have followed two approaches to displaying monitoring site data in Google Earth. The one is an inventory on a static web site, updated annually, where users can display available sites for regions of interest and bring up metadata, canned graphs and associated information for each site. The other method is now part of the WMS and gives database users the option to visualise selected sites in Google Earth.

The static inventory requires two flat files exported from the WMS: the first is a list of the more than 40 000 active sites with identifier, coordinates and description and the second is a list of the 500 000 sample results, one line per date, time and depth. Utilities adapted from the WaterMarque project mentioned above import the sites into a point coverage and the data into an INFO database.

Generating the monitoring site KML files

With no experience of KML, we found that a useful approach to get started was to set up a layout manually with a few sites in Google Earth itself, setting the required symbols, folder structure and descriptions.



Site class	Symbol	Symbol
normal		paddle/wht-circle-lv.png
unknown		paddle/wht-blank-lv.png
canal		pal5/icon50l.png
potable		paddle/blu-blank-lv.png
industry		pal2/icon10.png
hydro		paddle/blu-circle-lv.png
lake		paddle/ltblu-blank-lv.png
marine		paddle/grn-stars-lv.png
municipal		pal3/icon21.png
pipeline		shapes/target.png
sewage		paddle/ylw-blank.png
tunnel		shapes/donut.png
wetland		paddle/grn-blank-lv.png
pan		paddle/ylw-blank-lv.png
ground		paddle/g-lv.png
borehole		paddle/B-lv.png
borehole00		paddle/O-lv.png
borehole01		paddle/1-lv.png
borehole02		paddle/2-lv.png
borehole03		paddle/3-lv.png
borehole04		paddle/4-lv.png
borehole05		paddle/5-lv.png
borehole06		paddle/6-lv.png
borehole07		paddle/7-lv.png
borehole08		paddle/8-lv.png
borehole09		paddle/9-lv.png
borehole10		paddle/10-lv.png

Figure 2. a) Folder structure in Google Earth for DWAF's "C" drainage region, showing site types, secondary drainage regions ("C1", "C2", etc.) and site names; b) Complete list of icons used: most boreholes had 10 or fewer data records, so the number of records could be included in the icon.

When saved as a KML file, the code gives useful insights into the required syntax, enabling a more rapid learning process than starting from scratch with the on-line manuals and tutorials (e.g. Google Code 2008). Note that KMZ files such as those on our web site are simply zipped KML files. We used SecureZIP do the conversion (PKWARE 2007) but any zip programme would work.

Using the gawk implementation of awk (Gnu 2008), and a sorted inventory list as input, we generated KML code for more than 41 000 sites within South Africa's 22 primary drainage regions, structured in a logical hierarchy of site type, secondary drainage region and site code (see Figure 2a for an example with the C1 secondary drainage region). An alternative hierarchy has the 19 administrative Water Management Areas at the top level. Standard map symbols for water use types are still under development, so we selected suitable icons from the extensive set supplied by Google (Figure 2b). Representative listings of KML code are in Figure 3.

In parallel with the KML production process, the awk script outputs a set of HTML site listings with data links. This is to provide a means for downloading .zip files and to give access to water quality information for users who do not have Google Earth.

The Delphi process in the WMS water quality database generates similar KML files to those discussed above. The main difference in operation is that a user working in the WMS environment selects sites of interest and then activates the process to generate a KML document for those sites. When users require pre-packaged fly-throughs of study areas for presentations we use FRAPS (Beepa 2008) to capture movie sequences and Windows Movie Maker (Microsoft 2004) to splice and edit (example at DWAF-RQS 2008).

We used the 'Export to KML' extension for ArcGIS 9.1 (Martin 2007) to create KMZ files of the hydrologically important rivers and drainage basin boundaries. The source file for rivers was the South African 1:500 000 coverage (Silberbauer, 2006) and we subdivided the KMZ river files by stream order to reduce their size. The option exists to show river names only, because the 10 000 river arcs slow down the performance of Google Earth considerably. For the same reason, drainage regions created from the DWAF layers are available in each of the four hierarchy levels separately, with the quaternary drainage region labels in a file on their own.

Water quality data summaries

While the inventory is useful for planning, what users ultimately need is the water quality information at each site. With limited skills available in Internet database access programming, the most cost-effective approach was to dust off two existing tools from the old WaterMarque project (Barcode and Maucha) and create a set of static data summaries, text files and salinity symbols.

Data summaries are in the form of Barcode graphs, which are one-pagers of the major water quality indicators for a site. Sixteen bar graphs of major water quality indicators and a bar graph of flow data (where available) are stacked above one another with a common time scale (Silberbauer 1997). Summary statistics (Minimum, median, 90th percentile and maximum) for each variable appear on the right-hand y-axis and locality maps of monitoring sites at primary, secondary, tertiary and quaternary scale ensure that the plot serves as stand-alone data summary. The application also writes the data to a text file that users can download and plot with their own spreadsheet or graphing

software. Metadata about the source data of the plot appear in the heading and a text box. The Barcode macro runs in ESRI's Arc Macro Language (AML) using data in the INFO database and takes about one minute per monitoring site in batch mode on Sun Fire V490 server (i.e. 28 days to populate the whole web site, barring server and power failures).

The first lines of a KML set the environment and top level folder name:

```
<?xml version="1.0" encoding="UTF-8"?>
<kml xmlns="http://earth.google.com/kml/2.2">
<Document>
<name>S surface</name>
```

<Timespan> links the visibility of each site to the time slider:

```
<TimeSpan id="ID">
<begin>1990-10-17</begin>
<end>2007-12-11</end>
</TimeSpan>
```

Marker and balloon styles are at the head of the script:

```
<Style id="hydroPlacemark">
<IconStyle><scale>0.4</scale>
<Icon>
<href>http://maps.google.com/
mapfiles/kml/paddle/blu-circle-iv.png</href>
</Icon>
</IconStyle>
<BalloonStyle>
<text>${description}</text>
<bgColor>ffffff</bgColor>
</BalloonStyle>
</Style>
```

Defining a site type folder (Rivers):

```
</Folder></Folder>
<Folder>
<name>Rivers</name>
<styleUrl>#hydroPlacemark</styleUrl>
<open>0</open>
<Folder>
<name>SI</name>
<open>0</open>
```

Geographic coordinates (longitude, then latitude):

```
<Point>
<coordinates>28.015556, -32.515278,0
</coordinates>
</Point>
```

A lowest-level folder entry, for a monitoring site:

```
<Placemark>
<name>S7H4</name>
<Snippet maxLines="1">At Area 8 Springs B
on Groot-Keirivier</Snippet>
```

The details within a popup balloon are coded in HTML, between <![CDATA [and]]:

```
<description><![CDATA[

<b>WMS S70_102568</b><br />
S7H004<br />
At Area 8 Springs B on Groot-Keirivier<br />
<i>Rivers</i> samples: <b>325</b><br />
1990-10-17 to 2007-12-11<br />
<a href="http://www.dwaf.gov.za/iwqs/wms/data/s70/s70_102568.pdf">graph</a>|
<a href="http://www.dwaf.gov.za/iwqs/wms/data/S_reg_WMS_nobor.htm#102568">data</a>|
<a href="http://www.dwaf.gov.za/hydrology/cgi-bin/his/cgihis.exe/StationInfo?Station=S7H004">flow</a>|
<a href="http://www.dwaf.gov.za/iwqs/wms/data/000key.htm">home</a>|
<a href="http://www.dwaf.gov.za/iwqs/gis_apps/maucha.pdf">Maucha key</a>|
]]></description> <styleUrl>#hydroPlacemark</styleUrl>
```

Figure 3. Examples of KML code performing the main functions of the Google Earth water quality inventory. Extract full code listings by unzipping KMZ files on <http://www.dwaf.gov.za/iwqs/wms/data/000key.htm>.

Maucha symbols portray the major dissolved ions in a water sample, showing at a glance the ionic balance and dominant ions (Maucha 1932; Day & King 1995). Designed to indicate habitat type, their radial format and scalability makes them a better choice for mapping than the more commonly used Stiff diagram, developed for the petrochemical industry (Stiff 1951).

Results

On the Department of Water Affairs and Forestry web site, internal and external users can now access a 3D globe interface to the WMS water quality inventory grouped by hydrological or management region, and can view flow gauging and dam level monitoring sites (Figure 4). Users can decide whether to show surface water sites or groundwater sites and whether to symbolise sites by monitoring type icon or by Maucha salinity symbol.

At each site, the user can click on the symbol and display a balloon with metadata (site description, number of results, time period sampled, ionic diagram) and links to further information (time-series graphs of major ions, subset of raw data, flow data).

Within the national water quality database, the WMS, users can quickly visualise the localities of a site or group of sites in Google Earth. With the tools at our disposal, we can also quickly develop specific views of sites on request, such as topographical map overlays, estuaries or environmental monitoring sites.

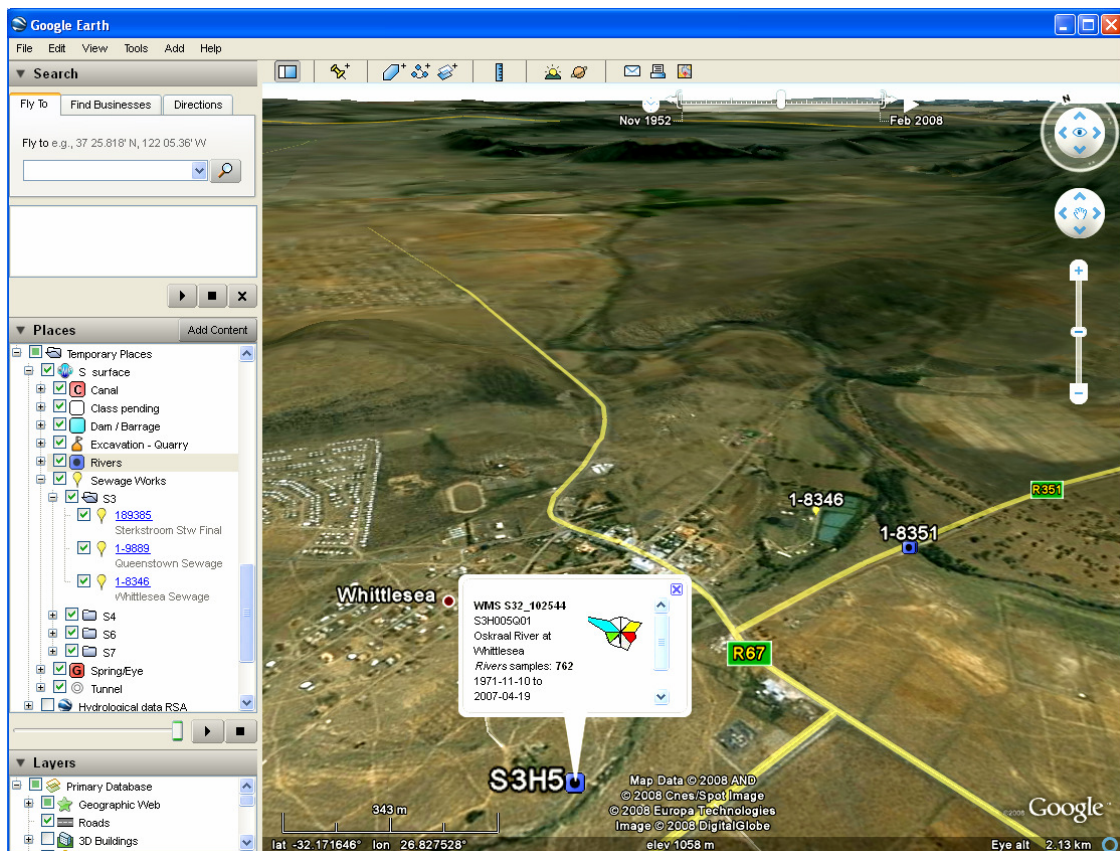


Figure 4. Water quality monitoring sites viewed in Google Earth, with a balloon at a river site including the Maucha symbol (main ions Na^+ , Mg^{++} and HCO_3^-) and links to further information. URL <<http://www.dwaf.gov.za/iwqs/wms/data/000key.htm>>.

Discussion

Customising Google Earth

Customisation is straightforward and the choice of KML generation method is not important. We simply used scripting tools lying on our 'workbench'. User requests are straightforward to implement, for example the requirement for a hierarchical folder structure and the separate display of borehole and hydrological sites from general water quality sites. The ability to create hyperlinks in HTML balloons allows for integration with many other data sources. However, certain links such as <.zip> and <mailto:> did not function as expected. The static data summaries (Barcode PDF plots) on the web site worked well, though users have to download the data from a separate HTML table if they need to change date ranges or show guideline levels.

Using Google Earth for a spatial inventory

The Google Earth interface is easy to use and allows rapid changes of scale from global to local and back. The inventory has been extremely useful for finding available sites in a study area, for seeing where sites are located in relation to one another, and often to detect sites where coordinates are in error. Preliminary usage tracking has been difficult to interpret because internal testing and internal usage have the same network location, but personal communications suggest that users are satisfied with the way the inventory works.

Google's massive centralised spatial database keeps the background data up to date with no effort on the part of the user: dam construction sites, pumps, sewage works are all there on the satellite images. The road network for South Africa displayed by Google Earth, in conjunction with the rivers and drainage regions from DWAF, ensure that hydrologists have no difficulty navigating through their study areas. During the past three years, the percentage of Africa covered by high-resolution images in Google Earth has increased greatly, giving an armchair view of land cover and surface structures not previously available so conveniently and on such a wide scale.

The built in functions and data are very welcome in a spatial analysis development environment where we have insufficient skills and resources to develop geoportal systems from the ground up. For example, something as deceptively simple as the Google Earth 'spread' function that displays overlapping points by offsetting them with leader lines, is extremely useful in identifying lazy coordinate capture, such as where all the sites for an impoundment are incorrectly located at one place. However, this is not easy to do on a conventional GIS.

Future developments

Eventually, we hope to develop a complete and stable set of water use icons for use in the WMS, GIS and KML environments. For future applications, an enterprise-based system might be more versatile for internal users, perhaps enabling live database queries from links in the popup balloons and the use of the time slider to view water quality trends at sites.

Other ways to create KML files

Off-the-shelf KML generation methods that do not involve developing specialised scripts are more practical for smaller, once-off data sets. One application that we have found to be useful for scientists without GIS training who need to view their own sets of coordinates in Google Earth is CSV2KML (CSV2KML 2007), an efficient, functional and free converter using spreadsheet data as input. Google Earth also provides KML creation tools on their site. Users find Google Earth less intimidating and simpler to configure than GIS, and are soon able to continue on their own. Those who are more comfortable in the GIS environment can create geographic layers and export them as KML files (e.g. Martin 2007, for ArcGIS).

Other KML viewers

Many software developers, including the traditional GIS companies, have developed 3D earth viewers. ArcGIS Explorer (build 480) reads KML and KMZ files but opens HTML text for a balloon in the default browser instead of inside the balloon. It does not spread overlapping symbols, scale labels, or interpret time data and has a more rigid GIS feel to it than Google Earth. World Wind 1.4 is supposed to have a KML interpreter but we could not get it to work. The Microsoft Virtual Earth site was not accessible through our firewall.

Problems

One limiting factor in deploying this type of application more widely is the digital communications infrastructure in Africa. Although very ingenious algorithms optimise Google Earth's download rate, the slow bit rate is still frustrating for the dial-up Internet user. In 2006, fewer than one in eight South Africans had any access to the Internet at all and by 2007 less than 5% of connections were 'broadband' (Miniwatts 2008). While we verified that Google Earth could operate over dialup thanks to its caching ability, some of the immediacy that makes it so effective is lost.

Users not skilled in spatial analysis understandably have difficulty in distinguishing between Google Earth and the WMS front-end, delivering unwarranted praise or derision depending on how they experienced the system. Some concerns include the perceived 'age' and 'fuzziness' of the images, the occasional mis-registration of a scene and illogical system errors (for example, failure by the operating system to recognise KMZ files after certain system updates).

Conclusion

Google Earth provides a remarkable platform for a practical monitoring site visualisation system, creating a convincing, immersive view of the 3D surface hydrology network for many types of data users, including aquatic scientists, anglers and students. The background satellite and road data sets are far more up to date than any small organisation with a limited budget could hope to achieve. Properties of the sampling network that have become lost as the corporate memory fades, such as gaps or duplicate monitoring, become obvious when presented in a scalable interactive system.

At a more fundamental level, this kind of convincing graphical representation of a visual environment helps to convey information directly to our well-developed perception systems. The system creates ‘an actively constructed, meaningful model of the environment that allows perceivers to predict what will happen in the future so that they can take appropriate action and thereby increase their chances of survival’ (Palmer 1999).

Acknowledgements

The Department of Water Affairs and Forestry’s Resource Quality Services directorate provided the data sets and the environment in which to set up the spatial inventory. The staff of the Department of Water Affairs and Forestry’s Water Management System, especially Mrs Marica Erasmus, assisted in exporting water quality data into flat files, Mr Kobus Myburgh was very patient in dealing with site location queries and Ms Bets Davies has developed a prototype set of site user icons.

References

- Allanson, BR 1961, ‘Investigations into the ecology of polluted inland waters in the Transvaal, Part I: The physical, chemical and biological conditions in the Jukskei-Crocodile River System’, *Hydrobiologia*, vol. 18, no. 1-2, pp. 2-76.
- Allanson, BR 2004, ‘Limnology in South Africa: Past and present status and future needs’, in B Gopal & RW Wetzel (eds.), *Limnology in Developing Countries*, vol. 4, pp. 1-116.
- Allanson, BR & Gieskes, JMTM 1961, ‘Investigations into the ecology of polluted inland waters in the Transvaal, Part II: An introduction to the limnology of Hartbeespoort Dam with special reference to the effect of industrial and domestic pollution’, *Hydrobiologia*, vol. 18, no. 1-2, pp. 77-94.
- Beepa 2008, ‘FRAPS real-time video capture and benchmarking’, viewed 2008-08-10, <<http://www.fraps.com>>.
- Cobban, DA & Silberbauer, MJ 1993, ‘Water quality decision-making facilitated through the development of an interface between a geographic information system and a water quality database,’ in SA Lorenz, SW Kienzle & MC Dent (eds.), *Hydrology in Developing Regions... the Road Ahead, Proceedings of the Sixth South African National Hydrological Symposium*, vol. II, pp. 523-530.
- CSV2KML 2007, ‘csv2kml 0.4b compiled in October 2007’, viewed 2008-08-09, <<http://www.tchartdev.com/csv2kml.htm>>.
- Day, JA & King, JM 1995, ‘Geographical patterns, and their origins, in the dominance of major ions in South African rivers’, *South African Journal of Science*, vol. 91, pp. 299-306.
- DWAF-RQS 2008, ‘Resource Quality Services water quality data exploration supplementary information’, viewed 2008-08-10, <<http://www.dwaf.gov.za/iwqs/wms/data/000key2.htm>>.
- Fuchs, C & Horak, E 2008, ‘Africa and the digital divide’, *Telematics and Informatics*, vol. 25 pp. 99-116.

- Google Code 2008, 'KML Reference', viewed 2008-08-10, <http://code.google.com/apis/kml/documentation/kml_tags_21.html>.
- Gnu 2008, 'Gawk', viewed 2008-08-10, <<http://www.gnu.org/software/gawk/>>.
- Martin, K 2007, *Export to KML Extension Version 2.4 documentation*, Bureau of Planning, City of Portland, viewed 2008-08-09, <http://outreach.cast.uark.edu/east/east/resources/Export_to_KML_documentation.pdf>.
- Microsoft 2004, 'Windows Movie Maker 2.1', viewed 2008-08-10, <<http://www.microsoft.com/windowsxp/downloads/updates/moviemaker2.msp>>.
- Miniwatts Marketing Group 2008, *Africa Internet usage and population statistics*, viewed 2008-08-09, <<http://www.internetworldstats.com/africa.htm>>.
- Palmer, SE 1999, *Vision science—photons to phenomenology*, The MIT Press, Cambridge, Massachusetts.
- PKWARE 2007, 'Data security software', <<http://www.pkware.com/>>.
- Schulze, RE 2006. 'Water Resources Complexities and Hydrological Needs in South Africa: Background', in RE Schulze (ed.), *South African Atlas of Climatology and Agrohydrology*, Water Research Commission, Pretoria, South Africa, WRC Report 1489/1/06, Section 21.1.
- Schulze, RE & Lynch, SD 2006, *Annual Precipitation*, in RE Schulze (ed.), *South African Atlas of Climatology and Agrohydrology*, Water Research Commission, Pretoria, South Africa, WRC Report 1489/1/06, Section 6.2.
- Schulze, RE & Maharaj, M 2006, 'A-Pan Equivalent Reference Potential Evaporation', in RE Schulze (ed.), *South African Atlas of Climatology and Agrohydrology*, Water Research Commission, Pretoria, South Africa, WRC Report 1489/1/06, Section 13.2.
- Silberbauer, MJ 1997, 'The application of geographic information systems to water quality monitoring', in *Remote Sensing and Geographic Information Systems for Design and Operation of Water Resources Systems* (Proceedings of Rabat Symposium S3, April 1997), IAHS Publication no. 242, pp. 189-195.
- Silberbauer, MJ 2006, 'The construction of a hydrologically-correct, annotated 1:500 000 spatial dataset of the rivers of South Africa and contiguous basins', Report number N/000/00/REH/0701, DWAF: Resource Quality Services, Pretoria, South Africa, viewed 2008-08-09, <http://www.dwaf.gov.za/iwqs/gis_data/river/River_Report_01.pdf>.
- Stiff, HA. Jr. 1951, 'The interpretation of chemical water analysis by means of patterns', *Journal of Petroleum Technology*, vol. 3, no. 10, section 1, pp. 15-16 & section 2, p3.
- Swart, SJ, van Veelen, M & Nell, U 1991, 'Water quality data inventory, Volume 1: Drainage regions A, B, C, D, E, F', Technical Report TR146, Hydrological Research Institute, Department of Water Affairs, Pretoria, ~400pp.
- van Veelen, M, Nell, U & Geldenhuys, WF 1990, 'Surface water quality of South Africa 1979-1988. Volume 2: Drainage Region A and B', Technical Report TR145, Hydrological Research Institute, Department of Water Affairs, Pretoria. ~100pp.

van Vliet, HR & Nell, U 1986, 'Surface water quality of South Africa, the Vaal River catchment 1979-1983', Technical Report TR131, Hydrological Research Institute, Department of Water Affairs and Forestry, Pretoria, 208 pp.