

New Cartographies to Chart

Over the last decade or so there has been a phenomenal growth in the use and diversity of information and communications technologies and the conceptual 'space' they support: cyberspace. Understanding the growth of cyberspace, and its myriad of social, economic, and political consequences, as well as the practical tasks of navigating and comprehending the various types or domains of cyberspace (such as the Web, email, real-time chat and instant messaging, file sharing, or 3D virtual worlds) is no easy task. Concepts and techniques from geoinformatics can be usefully employed to promote our understanding and to aid analysis of cyberspace. In particular, maps and map-like interfaces are increasingly becoming useful in representing, and in some senses actually creating, cyberspace. Mapping is thus being recognised as a powerful tool in the visualization and analysis of cyberspace. Therefore in this article the authors of *Atlas of Cyberspace* discuss how the Internet is being mapped.

By Martin Dodge and Rob Kitchin

Geoinformatics - Making Sense of Cyberspace Form

Paradoxically, more and more of our time and our leisure and business activities are spent in virtual space and yet it is a space that is difficult to comprehend and mentally visualise. Moreover, it is a space in which it is easy to get lost and confused. This is because, with the exception of the supporting infrastructure – fibre-optic cables, servers, satellites and so on – cyberspace is composed of computer code with no material existence. As a consequence, many of the projects that seek to map cyberspace are using processes of spatialisation in order to make sense of its form and the transactions that take place there. Of particular significance to the geoinformatics community is that these projects are developing new mapping techniques that are pushing back the boundaries of cartography and how we interact with maps, creating new kinds of interactive and dynamic representations.

Mapping Network Infrastructure

Much of the effort to map cyberspace has historically focused on the measurement and representation of the supporting tech-

nology of telecommunications and computing infrastructure. At a basic level, it is relatively easy to map the locations of telecommunications infrastructure such as cables and computer servers onto real-world geography, at various scales, using typical GIS methodologies. Indeed much of

Cyberspace is difficult to comprehend and mentally visualise.

this infrastructure is planned, installed and maintained using AM/FM, CAD databases and cable management systems that are based around spatial databases. This type of approach has many merits, providing the necessary geographic inventory and census of where the Internet nodes are located, how the data networks interconnect them, and the traffic that flows between them. Well-designed infrastructure maps can clearly show how computers are physically wired together to create complex networks that link cities and countries across the globe. It also reveals the uneven geographic distribution of

infrastructure and those areas of the world that have poor access to cyberspace or are presently excluded altogether. An example of such a map is shown in Figure 1. It is a 3D, interactive map of the Internet Mbone network created by researchers Tamara Munzner, K. Claffy, Eric Hoffman and Bill Fenner. The Mbone comprises a special set of routes, known as 'tunnels' in technical jargon, which run on top of the ordinary Internet. Munzner and her colleagues map these tunnels as arcs on 3D model of the globe, which the user can manipulate to rotate and view from any angle. The line colour and thickness are used to show characteristics of the Mbone tunnels, while the height of the arcs above the surface of the globe is a function of distance between the end Mbone router nodes.

Geographic Maps of Internet Infrastructure

Geographic maps of Internet infrastructure are commonly employed by network owners for two particular tasks. First, they are useful to network engineers as interface tools in the monitoring and control of traffic flows and network performance. Second, they are used as promotional tools to demonstrate to potential customers how extensive and capable their networks are. The cartographic techniques employed can vary substantially, including interactive 3D globe representations and dynamic maps that update in real-time.

Network Weather Maps

The detailed network monitoring maps and tools are generally not made public for reasons of security and commercial confidentiality. However, some Internet networks, particularly those serving the research and education communities, do provide summary performance data using map interfaces. These interfaces are popularly referred to as network weather maps. They are in some senses useful marketing as well as public-spirited information dissemination tools, providing network customers (usually universities and labs) with useful information, especially to identify trouble spots. Below are two interesting examples of network weather maps – the Abilene network in the US (Figure 2) and NORDUnet serving

Cyberspace




Figure 1: 3D arcs on a globe representation of the Internet Mbone network.

(Courtesy of Tamara Munzner and IEEE, www-graphics.stanford.edu/papers/mbone/)

Scandinavia (Figure 3). The maps are updated frequently (for example the Abilene map is updated every five minutes) and allow outside users a 'peak inside' the network. The map is simply a summary of overall network performance with links colour coded by their traffic flows, but it also provides a visual interface to browse more detailed statistics available as tables and statistical charts. The NORDUnet example is also interesting because it uses a logical network diagram instead of a geographic link-node map to represent the infrastructure. There are many other cyberspace mapping efforts that also dispense with the framework of real-world geography. For example, researcher Stephen Coast at the Centre for Advanced Spatial Analysis has been undertaking large-scale measurement and mapping of the topology of IP routing. His mapping technique currently uses abstract graphs - for example, Figure 4 shows the topology of the large campus IP network of University College, London. Each cluster in the map is an individual university building or large department and the distance metric is based on speed of network links. So the long grey lines are generally network links between buildings. In total, some 4,600 nodes of the network are shown.

The Internet Is Its Own Space

Another example, at a larger scale is the Internet Mapping Project (<http://www.cs.bell-labs.com/~ches/map/index.html>) being undertaken by Bill Cheswick and Hal

Burch, at Lumeta Corporation (formerly at Bell Labs). Their project maps the topology of thousands of interconnected Internet networks to provide perhaps the best currently available large-scale overview of the core of the Internet in a single snapshot. They map the Internet in an abstract space, disregarding the actual location of nodes in physical space. As Cheswick argues, 'We don't try to lay out the Internet according to geography... The Internet is its own space, independent of geography.' Data is gathered by using the Internet to measure itself on a daily basis, surveying the routes to a large number of end-points (usually Web servers) from their base in New Jersey, USA. The resulting huge graph maps how the hundreds of networks and many thousands of nodes connect together to form the core of

Many cyberspace mapping efforts dispense with the framework of real-world geography.

Internet. The striking example (Figure 5) shows the structure of the Internet from December 2000, representing nearly 100,000 network nodes. This highly complex graph takes several hours to generate on a typical PC. The layout algorithm uses

simple rules, with forces of attraction and repulsion jostling the nodes into a stable, legible configuration. There are many permutations in the algorithm to generate different layouts and colour-codings of the links according to different criteria (such as network ownership, country). In the example shown, links have been colour-coded according to the ISP, seeking to highlight who 'owns' the largest sections of Internet topology. This project is ongoing and the data is archived and available to other researchers to utilise. Over time, it is hoped that the data will be useful for monitoring growth and changes in the structure of the Internet. Also, the experience gained in mapping the Internet is being applied commercially by Cheswick and Burch at Lumeta where they use the network scanning and visualization techniques to chart the structure of corporate intranets to identify security weaknesses and unauthorized nodes.

Mapping Information

The Internet, and the Web in particular, is a large and rapidly growing information resource. Literally billions of pieces of information are stored on networked computers around the globe and they can potentially be accessed in seconds. However, searching and navigating this information can be difficult, especially when it is composed of long lists of references. Finding useful and relevant information in a timely fashion can be very frustrating, as most Web users will be only too aware. One experimental solution has been to spatialise large information collections - that is, to summarize and char-



Figure 2: 'Weather map' of the traffic load on the core links of the Abilene network. (Courtesy of the Abilene Network Operations Center, Indiana University, <http://hydra.uits.iu.edu/~abilene/traffic/>)

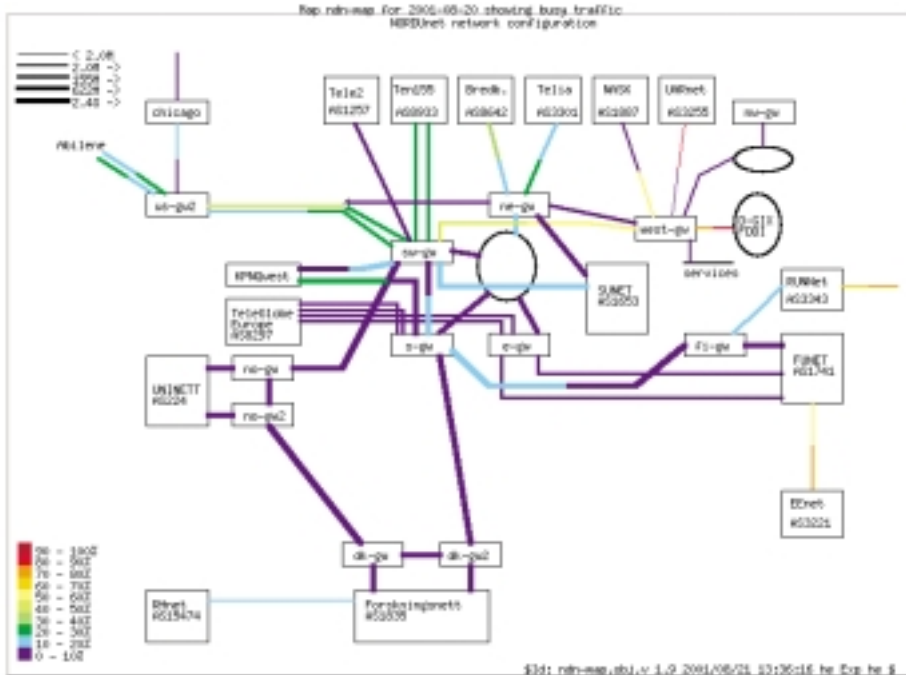


Figure 3: The 'weather map' of network load for the NORDUnet network. (Courtesy of NORDUnet, www.nordu.net)




Figure 4: The topology of the UCL campus network by Steve Coast. (Courtesy of Stephen Coast, Centre for Advanced Spatial Analysis, UCL, www.casa.ucl.ac.uk/steve/stuff/ipmap/)

acterize the information using a map-like representation where location, area, distance and proximity in the display represent key aspects of the data (e.g. the more related the information, the closer together it is drawn on the map). The result is an information map that summarizes thousands of pieces of information on a single screen and which can be navigated like a conventional map. Information mapping exploits the fact that people generally find it easier to process and understand visual displays than large volumes of written text or columns of numbers. This process is known as spatialisation and different algorithms can produce a variety of map forms, ranging from simple

2D maps to immersive 3D fly-through data-landscapes, and in scale from individual websites up to large sections of the Web. Examples of such spatialisations include Map.net, which maps a large directory of websites and Map of the Market, which shows the changing stock market.

Map.net

Shown in Figure 6 below is a small part of Map.net (<http://www.map.net/>), a huge spatialisation of the Open Directory, a classification of over 2 million websites into more than 350,000 categories. The map is a multi-

levelled hierarchy of categories that gets progressively more detailed in scale as the hierarchy is traversed. Browsing the map is achieved by the standard point-and-click navigation approach of the Web. The categories of websites are represented by colour coded, irregular shaped polygons whose size is in proportion to the number of websites they represent. Individual websites of particular significance in the category are highlighted by target symbols. The example below shows the third level in the information hierarchy, displaying a map of websites about board games. There are a range of more detailed categories for different types of board games available shown by the different sized and coloured polygon, arranged in alphabetical order across the map extent.

Map of the Market

Another excellent example of the potential power and usefulness of spatialisation for information summary and navigation is the Map of the Market (<http://www.smartmoney.com/map/>). Understanding the daily fluctuations in the stock market is a serious business for traders, analysts and investors and the Map of the Market is able to show the changing stock prices of over 500 publicly-traded companies on a single screen. Since its launch by SmartMoney.com at the end of 1998, the Map of the Market has become a firm favourite with users. This is due, in large part, to the fact that it presents large volumes of fast changing data in a very useful and usable format, providing people with answers to the basic question 'how is the market doing today?' at a single glance. It is probably the most useful exemplar of information mapping on the Web today. On one single map users can quickly gain a sense of the overall market conditions, yet still see many hundreds of individual data elements. It is fully interactive, allowing the user to access a great deal of information, statistics and news on the companies by clicking on their tiles.

Compact and Elegant Spatialisation

Map of the Market is a compact and elegant spatialisation using a display of coloured rectangular tiles to show the changing performance of the most important companies on the US stock market, updated every fifteen minutes. Figure 7 shows an example of Map of the Market from the 16th of July 2001. Each tile represents a single company, with the size of the tile being proportional to its market capitalization. So, the larger the tile,

Legend

live.net	8078
adobe.net	3957
apple.com	2479
att.net	2294
apple.net	2218
apple.net	2002
ams.net	1943
lzu.net	1848
teleplan.net	1438
quest.net	1243
holiba.net	1128
gsl.net	1128
veris.net	1058
hmc.net	807
deliaud.net	808
gbls.net	688
telelobe.net	506
qd.net	501
live2.net	328
swins.net	314
clgas.net	313
axobu.net	496
swat.net	431
lupew.net	422
stavis.net	315
ksa.net	376
hympanet.net	328
conf.net	307
lbi.net	305
Other ISPs	33271
Not an ISP	
none	



Figure 5: Map of the Internet topology by Hal Burch and Bill Cheswick. (Courtesy of Peacock Maps Inc., www.peacockmaps.com)

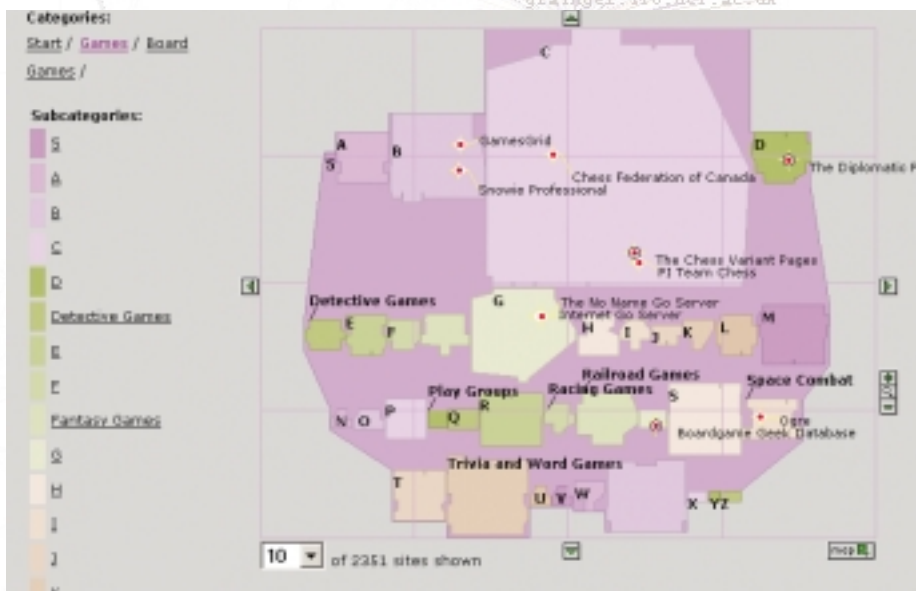


Figure 6: The Map.net map of a small section of Open Directory showing board game websites. (Courtesy of Antarcti.ca Systems Inc., www.map.net)



Figure 7: The Map of the Market. (Courtesy of SmartMoney.com, www.smartmoney.com)

the greater is the value of the company. This allows the value of different companies to be assessed by simply comparing the size of their tiles. The colour of the tile encodes the change in the company's stock price over a set time period. In the default colour settings, red represents a declining stock price, while green shows positive growth. Black indicates no change. The stronger the saturation of colour, the stronger is the percentage stock price change (either negative or positive). The position of the tiles in the map also conveys useful information. Firstly, companies are arranged into eleven major sectors and then further grouped by industry type. The particular spatial layout of tiles within the industry block tries to cluster companies as close together as possible, based on historically similar stock price movements - the idea being to create neighbourhoods in the map which contain similarly-performing companies. So, on the Map of the Market you will often see distinctive spatial patterns of light and dark tiles that can reveal significant large-scale trends in stock prices.

Conclusion

In this article we have sought to provide a brief introduction to the ways in which academic and commercial researchers are seeking to map and make sense of cyberspace. The maps they are creating are, we would argue, pushing back the barriers of cartographic technique and design and should therefore be of interest to everyone interested in geographic visualization.

Martin Dodge works as a computer technician and researcher in the Centre for Advanced Spatial Analysis (CASA), at University College London. He maintains the Cyber-Geography Research web site at <http://www.cybergeography.org>, which includes the Atlas of Cyberspaces. With co-author Rob Kitchin, he has also written the books *Mapping Cyberspace* (Routledge, 2000) and *Atlas of Cyberspace* (Addison-Wesley, 2001).

Rob Kitchin is a senior lecturer in human geography at the National University of Ireland, Maynooth. He is the author of *Cyberspace* (Wiley, 1998) and the co-author of *Mapping Cyberspace* (Routledge, 2000) and *Atlas of Cyberspace* (Addison-Wesley, 2001). He has published three other books and is the general editor of the journal *Social and Cultural Geography*.

Editors Note:

In depth and detail a wide range of such projects are presented in the book "Atlas of Cyberspace", written by the authors of this article Martin Dodge & Rob Kitchin and published by Addison-Wesley in August 2001 (HB, 268 pages, ISBN: 0201745755).