Mapping Cyberspace

Martin Dodge and Rob Kitchin

"... maps of cyberspace are almost as rare as 16th century portalans. [Cyberspace] explorers practice their trade without maps. ... Few among this frontier fraternity have both the navigational and drafting skills of a Ferdinand Magellan or a James Cook. Even for those that do, the challenge of mapping cyberspace is in some ways more formidable than that faced by the sea captains of the past." (Staple 1995, page 66)

Introduction

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. Surveying concepts and visualisation techniques from cartography, we would argue, can be usefully employed to promote our understanding and to aid analysis of cyberspace. After all cartographer have long experience in producing legible representations of complex, multidimensional space. 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. In this short article we review some of different aspects of cyberspace that can be mapped and outline in brief seven projects that are attempting to survey and map the Internet (for more in-depth coverage see Dodge and Kitchin 2000, Dodge 2002).

These 'maps' aim to address one of the paradoxes of the Internet, namely that 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. 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 (using visualization based on mapping principles) in order to make sense of its form and the transactions that take place there. Of particular significance to the members of the Society of Cartographers 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.

__Martin Dodge__, Centre for Advanced Spatial Analysis, University College London, 1-19 Torrington Place, Gower Street, London. Telephone: 020 7679 1255, Fax: 020 7813 2843, Email: m.dodge@ucl.ac.uk
__Rob Kitchin__, Department of Geography and NIRSA, National University of Ireland, Maynooth, County Kildare, Ireland. Telephone: +353 1 708 3372, Fax +353 1 708 3573, Email: rob.kitchin@may.ie
An earlier version of this article appeared in GeoInformatics magazine, and a version was given as a paper at the Society of Cartographers conference at Liverpool John Moores University on 3rd Sept 2002.

Mapping Network Infrastructure

Much of the effort to map cyberspace has historically focused on the measurement and representation of the supporting technology 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 this infrastructure is planned, installed and maintained using AM/FM, CAD databases and cable management systems which 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. This has allowed geographers and others to analyse the material geography of ICT infrastructure, and information production and consumption (see Kellerman 2002 for recent review).

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 (1996). 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. It is also possible to interactively query each of the links by clicking on them; in the case of Figure 1 a link between Mountain View, California and

Moscow has been chosen with details displayed at the bottom of the window.

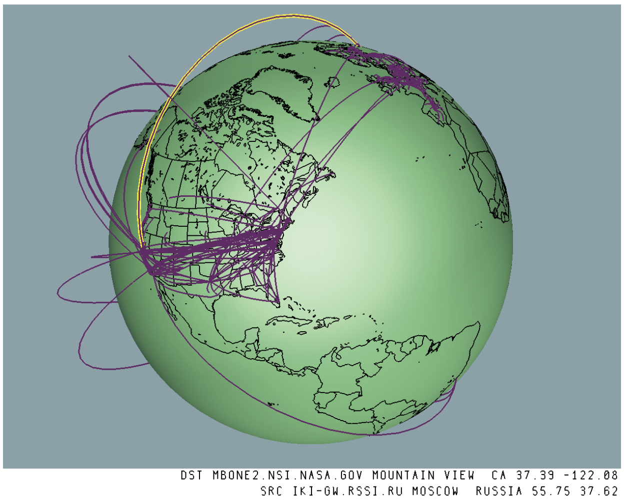

Figure 1: 3D arcs on a globe representation of the Internet MBone network.
(Courtesy of Tamara Munzner)

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, and are conceptually very similar to airline route maps found in the back of in-flight magazines. The geographic visualisation techniques employed can vary substantially, including interactive 3D globe representations and dynamic maps that update in real-time.

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 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.

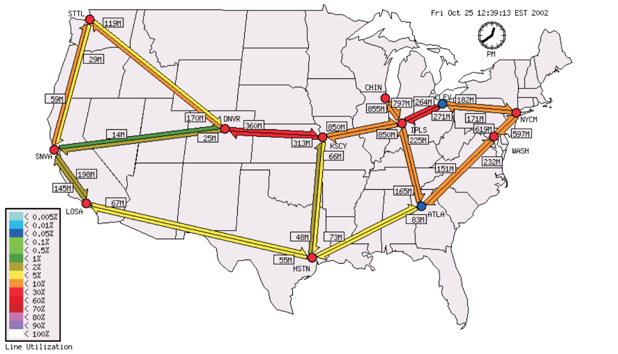

Figure 2: 'Weather map' of the traffic load on the core links of the Abilene network as of Friday 25t[h] October 2002
(Courtesy of the Abilene Network Operations Center, Indiana University, <http://loadrunner.uits.iu.edu/weather maps/abilene/)

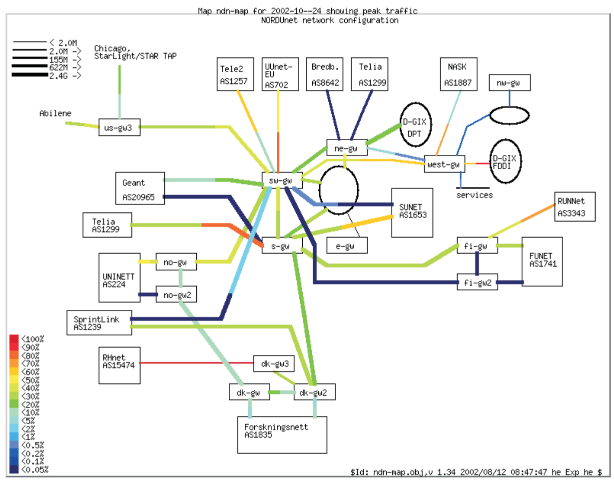

Figure 3: 'Weather map' summarising network load for the links in the NORDUnet network for the 25th of October 2002.
(Courtesy of NORDUnet, <http://www.nordu.net>)

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 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 below shows the topology of the large campus IP network of University College London (UCL). Each cluster in the spatialisation is an individual university building or large department and the distance metric is based on speed of network links. The long grey lines are generally network links between buildings. In total, some 4,600 nodes of the network are shown.

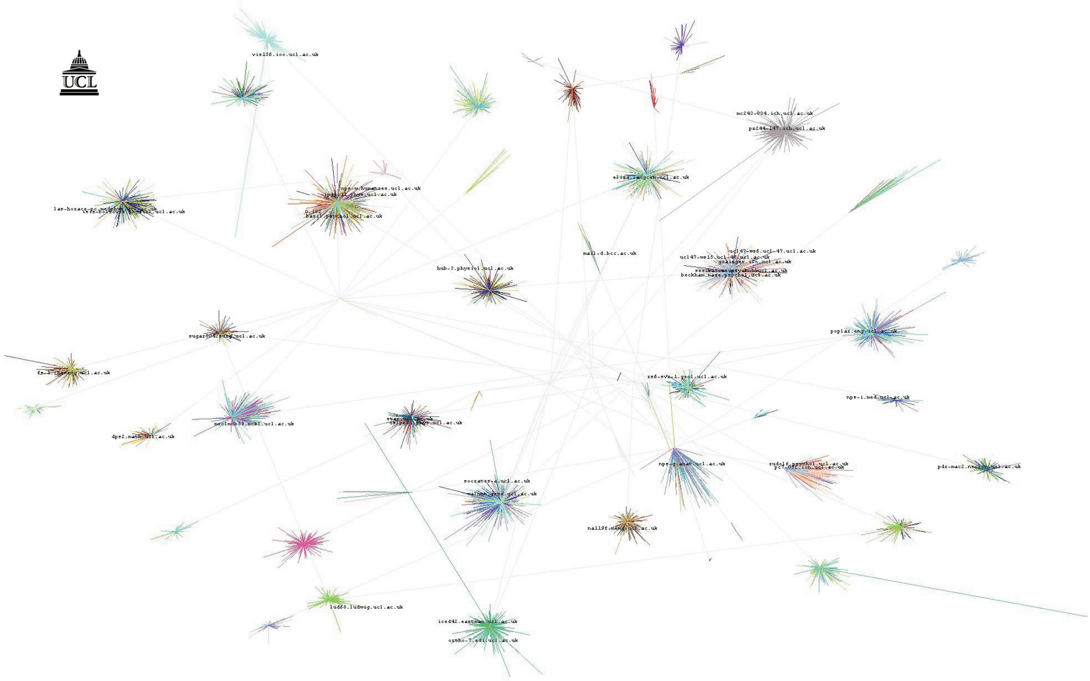

Figure 4: The topology of the UCL campus network by Steve Coast.
(Courtesy of Stephen Coast, Department of Physics, UCL, < http://www.fractalus.com/steve/stuff/ipmap>)

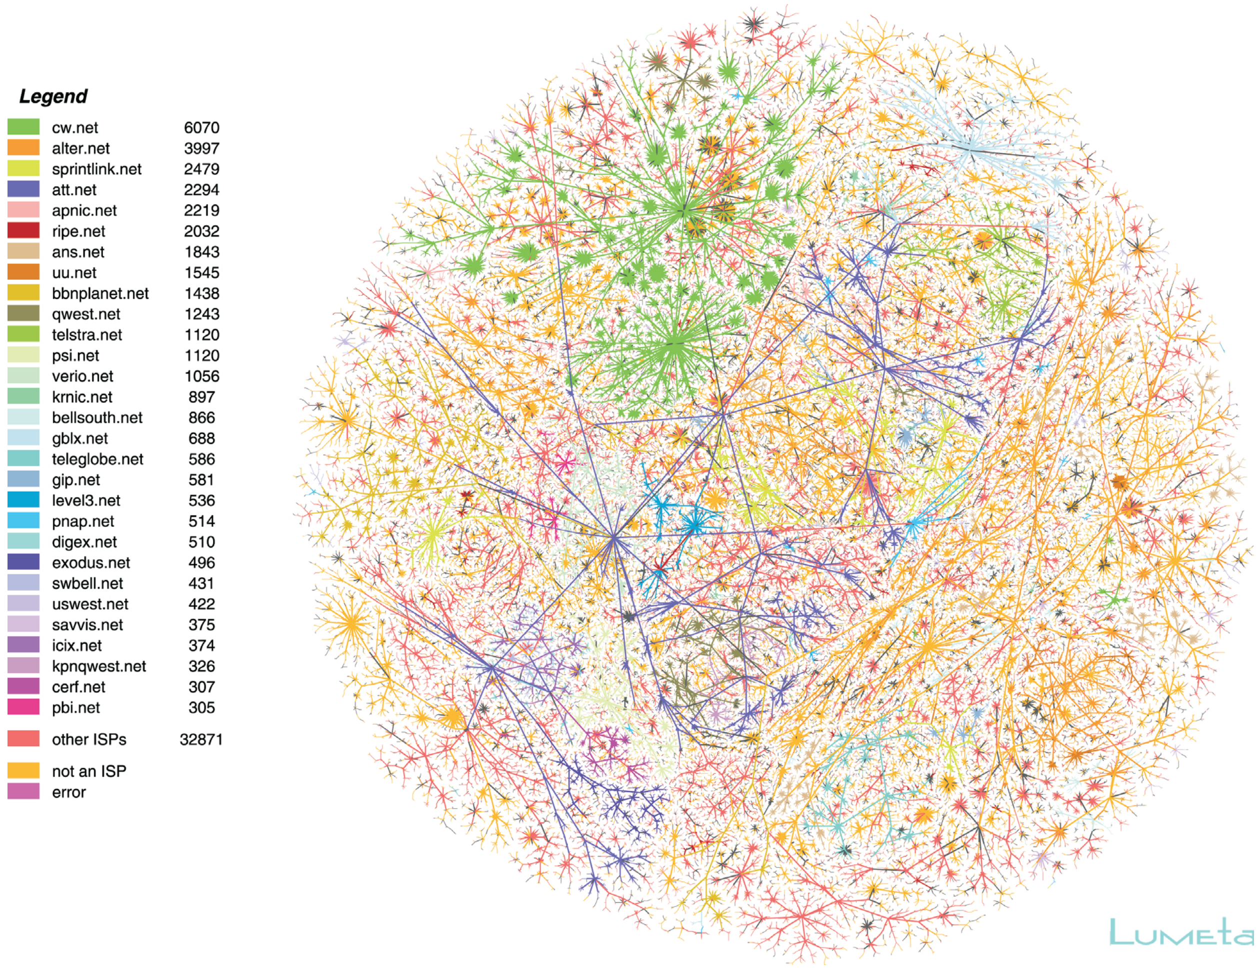

Figure 5: Map of the Internet topology by Hal Burch and Bill Cheswick.
(Courtesy of Peacock Maps Inc., <http://www.peacockmaps.com>)

Another example, at a much larger scale is the Internet Mapping Project being undertaken by Bill Cheswick and Hal Burch, at Lumeta Corporation (Branigan *et al.* 2001). Their work began in 1998 as a research project at Bell Labs, to map the topology of thousands of interconnected Internet networks and it now provides one the best available large-scale overviews of the core of the Internet in a single visual snapshot.
(You can buy versions of their maps as wall posters from Peacock Maps, <http://www.peacockmaps.com>).

They map the Internet in an abstract space, disregarding the actual location of nodes in physical space. As Cheswick has says, *"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 spatialisation maps how the hundreds of networks and many thousands of nodes connect together to form the core of Internet. The striking example (Figure 5) shows the structure of the Internet from December 2000, representing nearly 100,000 network nodes. This highly complex spatialisation 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 Cheswich 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. Indeed, there is increasing interest from the military and police for maps of cyberspace as concerns about threats of 'cyberterrorism' and 'info-warfare' grow with societies increasing dependence on data networks and information space for daily life.

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 be too aware. Research in information visualisation is using many map-like representations to display the structure of abstract information and provide interactive interfaces for navigation (for a review see Card et al. 1999).

One experimental solution emerging from information visualisation research has been to spatialise large information collections – that is, to summarize and characterize 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. Example of such spatialisations include Map.net, which maps a large directory of websites and Map of the Market, which shows the changing stock market.

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-leveled 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 categories on cooking. There are a range of more detailed categories for different types of cooking (e.g. 'Soups and Stews', 'World Cuisine') available shown by the different sized and coloured polygon, arranged in alphabetical order across the map extent.

Another good example of the potential power and usefulness of spatialization 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

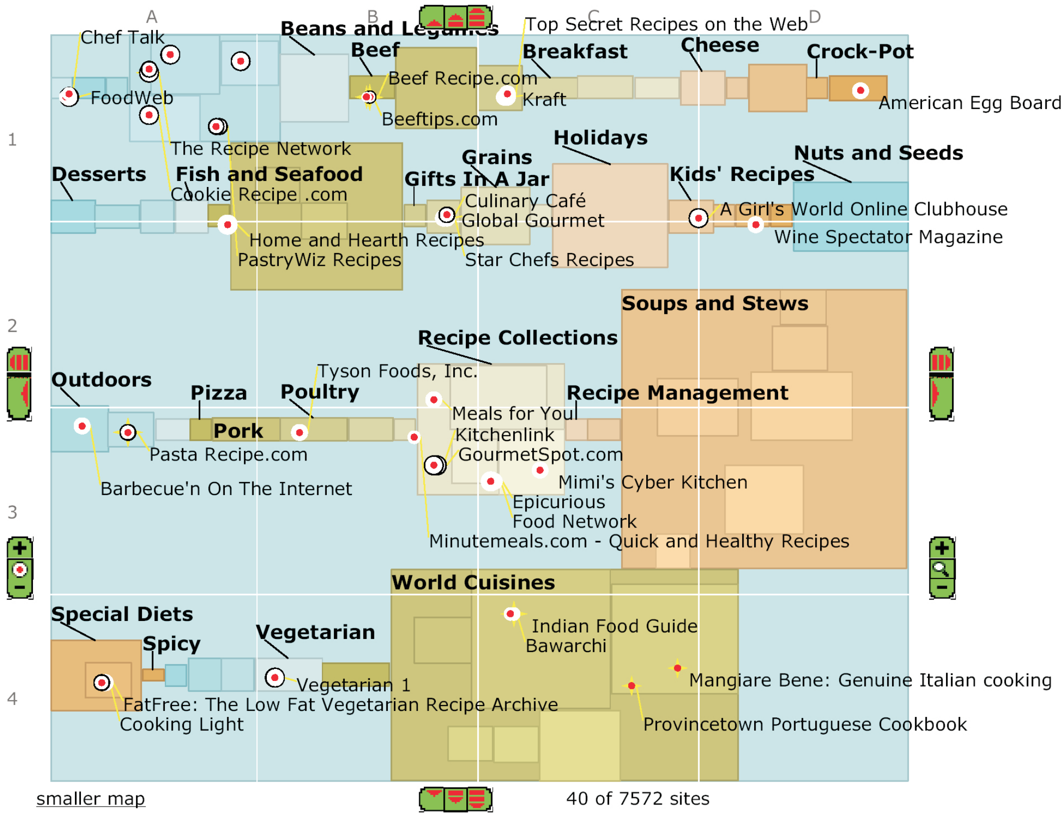

Figure 6: The Map.net map of a small section of Open Directory showing categories of cooking websites.
(Courtesy of Antarcti.ca Systems Inc., <http://www.map.net>)

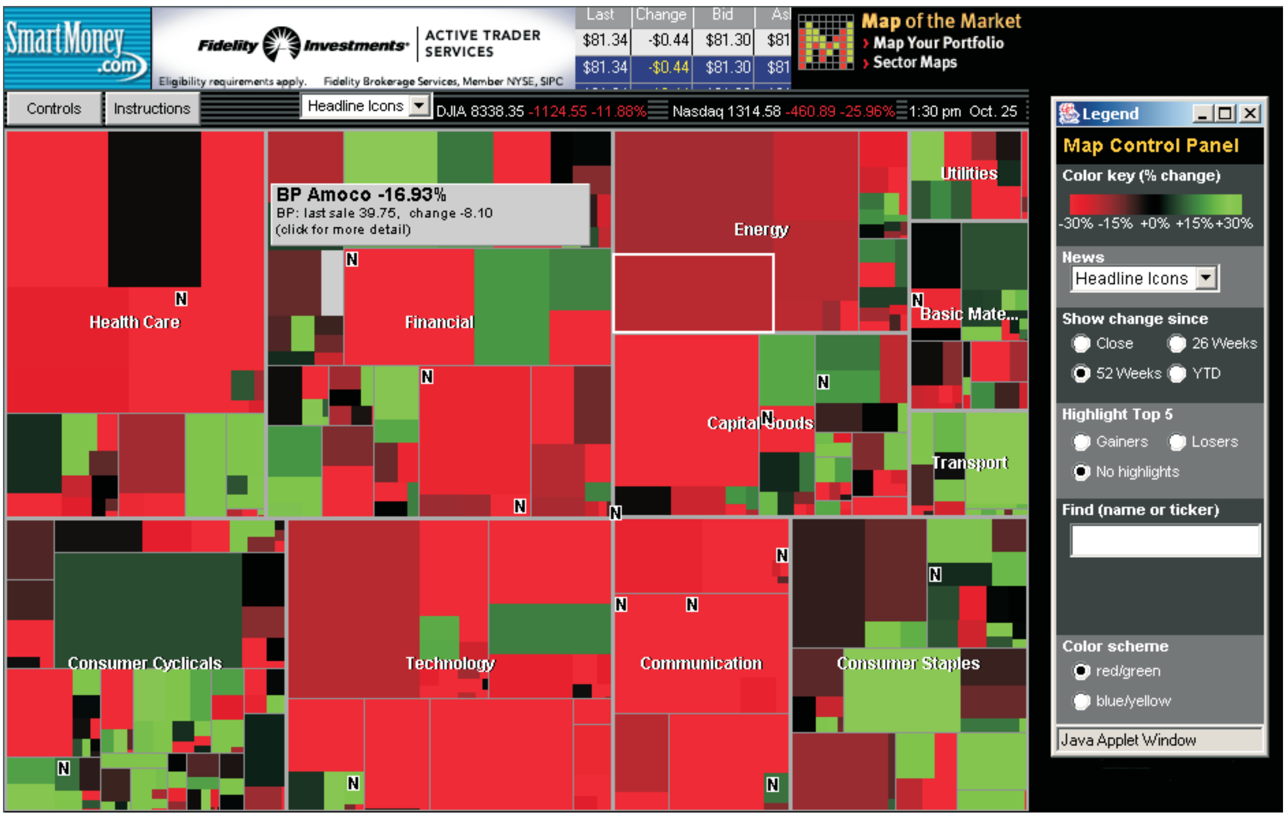

Figure 7: The *Map of the Market*.
(Courtesy of SmartMoney.com, <http://www.smartmoney.com/>)

the user to access a great deal of information, statistics and news on the companies by clicking on their tiles.

Map of the Market is a compact and elegant spatialization 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 25th of October 2002. Each tile represents a single company, with the size of the tile being proportional to its market capitalization. So, the larger the tile, the greater 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 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 which 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.

References:

Branigan, S., Burch, H., Cheswick, B., and Wojcik, F. (2001) "What Can You Do with Traceroute?". *Internet Computing*, September/October 2001, Vol. 5, No. 5, page 96. <http://computer.org/internet/v5n5/index.htm>, last accessed 25th October 2002. See also <http://www.lumeta.com/mapping.html>.

Card S.K., MacKinlay J., and Shneiderman B. (1999) *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufman Publishers, San Francisco.

Dodge, M. (2002) *Atlas of Cyberspace website*. Centre for Advanced Spatial Analysis, University College London. <http://www.cybergeography.org/atlas>, last accessed 26th October 2002.

Dodge, M. and Kitchin, R. (2000) *Mapping Cyberspace*. Routledge, London.

Kellerman, A. (2002) *The Internet on Earth: A Geography of Information*. John Wiley & Sons Ltd, London.

Munzner, T., Hoffman, E., Claffy, K. and Fenner, B., (1996) "Visualizing the Global Topology of the Mbone". *Proceedings of the 1996 IEEE Symposium on Information Visualization*, 28-29th October 1996, San Francisco, CA, pp. 85-92. <http://www-graphics.stanford.edu/papers/mbone/>, last accessed 25th October 2002

Staple, G.C. (1995), "Notes on Mapping the Net: From Tribal Space to Corporate Space". In Staple G.C. (ed.) *TeleGeography 1995: Global Telecommuncations Traffic Statistics and Commentary*, page 66-73.

About authors

Martin Dodge works as a researcher in the Centre for Advanced Spatial Analysis and a lecturer in the Department of Geography 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 Director of the National Institute for Regional and Spatial Analysis and 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 six other books and is Managing Editor of the journal *Social and Cultural Geography.*