10

Martin Dodge and Rob Kitchin

Charting Movement: Mapping Internet Infrastructures

Over the last decade or so there has been a phenomenal growth in the use and diversity of information and communications technologies (ICTs), with the rise of the Internet being of particular note. As of September 2002, there were 605.6 million people from around the world using the Internet for all manner of personal and business communications (Nua 2003). Along with this growth, there has been a multi-billion dollar investment in vast assemblages of powerful computer servers and the infrastructure necessary to support current and projected demand in information processing and exchange, including long-haul, fiber-optic backbone networks to link countries and metropolitan cores, high-speed routers and switches, and "last-mile" DSL and cable. Understanding the development and growth of ICTs, the myriad of their social, economic, and political consequences, as well as the practical tasks of planning infrastructure deployment, however, is no easy task. In this chapter, we argue that one useful strategy for analyzing and comprehending the Internet is the application of concepts and techniques from cartography and geographic visualization.

Maps and visualizations have long been used as a way of making the world more comprehensible. Mapping provides a means by which to classify, represent, and communicate information about areas that are too large and too complex to be seen directly. Well designed maps are relatively easy-to-interpret and constitute concentrated databases of information about the location, shape, and size of key features of a landscape and the connections between them. Moreover, the process of spatialization, where a spatial, map-like structure is applied to data where no inherent or obvious one exists, can provide an interpretable structure to large databases of abstract information (Couclelis 1998). In essence, maps and spatializations exploit the mind's ability to more readily see complex relationships in images, providing a clear understanding of a phenomenon, reducing search time, and revealing relationships that may otherwise not have been noticed. We illustrate the power of a mapping strategy by focusing on its usefulness in comprehending Internet infrastructure, although as we detail elsewhere, mapping and spatialization can be used to develop an understanding of a wide range of Internet uses and content (Dodge and Kitchin 2000a, 2001).

Internet infrastructure, and its use, is often taken for granted because, unlike roads or railways, it is largely invisible—buried underground, snaking across ocean floors, hidden inside wall conduits, or floating unseen in orbit above us. Indeed, given its invisibility, it is easy to assume that it is as ethereal and virtual as the information and communication that it supports. Consequently, there are a number of elements to Internet infrastructure that we presently have little systematic knowledge about, such as the form and function of backbone networks and their subsidiaries, network routing and traffic conditions, user demographics, marketing penetration and ownership, the physical location of computer servers (hosts) and Internet addresses, connectivity, and bandwidth. The mapping of these elements we believe serves a number of useful functions, providing important insights into who owns and controls infrastructure, who has access to the Internet, how the system can be surveyed, and how and from where the Internet is being used. This is vital information for the planning of new provision and the setting of policy and regulatory guidelines.

At a basic level, the maps provide a spatialized inventory and census of *where* Internet nodes and routes of connection are located, and in specific cases the traffic that flows through networks and their user profiles. Maps of network infrastructure can show clearly how computers are physically wired together to create complex networks that operate over several spatial scales, building into global scale systems. Depending on scale, these maps can be used by engineers to install and maintain the physical hardware of the networks, by system operators to manage networks more effectively, and by marketing and business development departments to demonstrate the size and penetration of networked services.

In addition, the maps have academic utility by showing significant trends and spatial patterns in the growth of network architecture, service provision, user profiles, and traffic flows across spatial scales, so for example, allowing comparison of neighborhoods, cities, and countries. As such, they reveal the growth of the "network society" and the "information economy." Maps also reveal the uneven and unequal distribution of infrastructure, showing those areas that have poor access to the Internet or are presently excluded altogether (Castells 2000, 2001). Moreover, they allow an analysis of the changes occurring in these patterns. As recent research highlights, although the Internet has expanded, diversified, and diffused greatly, basic infrastructure access and equity issues are still significant as can be seen by the proliferation of discussions about the "digital divide" that track the different ways and the different spatial scales at which access is fractured along lines of wealth, class, race, gender, and so on (Norris 2001, U.S. Department of Commerce 2000, Warf 2001).

Perhaps not unsurprisingly given the Internet's varied nature, maps of its infrastructure come in a variety of forms both in terms of what is mapped (e.g., network structure or traffic flows) and how it is mapped. The cartographic designs employed are various. Many examples use conventional approaches of shaded or symbol maps on a familiar geographic framework (these are often produced using standard GIS packages). However, other significant examples stretch the notion of a "map" using more diagrammatic approaches, for example showing the topology of network connections laid out in a non-geographic, abstract coordinate space. Some of the maps are interactive interfaces, using the medium of the map to allow users to access and query the data in novel ways. Some of the most potentially powerful and interesting "new breed" of infrastructure maps are dynamic in nature, constructed with live data gathered from the Internet every time the map is requested by a user.

In the remainder of the chapter, we provide a review of some different projects that have sought to map Internet infrastructure, dividing our discussion into four sections, themed by map purpose:

- maps for operational Internet management
- maps for Internet marketing
- maps for Internet policy and planning
- maps for academic Internet analysis.

Our selection of projects is limited by space, so we have chosen projects that have particular salience in relation to Internet infrastructure policy and planning, either for the public sector or commercial companies, and importantly are publicly available for wider analysis.* The maps are produced by many different people, ranging from interested individuals, to academic research groups, consultants and commercial analysts, to government regulators and network operators and marketing departments at Internet service providers (ISPs).

For a more comprehensive review, see Dodge and Kitchin 2000a, 2001.

Maps For Operational Internet Management

Managing large-scale and geographically distributed network infrastructure is a challenging and demanding task. Network managers need to insure the fast and uninterrupted flow of gigabytes of data traffic from multiple origin points to many destinations. It requires skill and attention to identify, correctly diagnose, and rectify faults in hardware and the complex software systems that control data traffic routing. This is made more challenging by the fact that (1) many ISPs have service agreements with customers that specify a minimum network performance and reliability at the 99.9 percent mark (or higher), which amounts to acceptable outages equivalent to just 4.4 hours per year, and (2) there are significant issues of cooperation between ISPs due to the decentralized and distributed nature of the global Internet. In relation to the latter point, it is often forgotten that the Internet is not a homogeneous single network, but rather a network composed of networks, each of which is owned and operated by separate (often competing) companies and organizations. This means that there is no central command or overall management of the Internet. Consequently, it is often the case that operational network problems, due to hardware failure or misconfiguration of software within one ISP, can have significant effects widely elsewhere in the Internet; a major event at a strategic location on the Internet can have widespread impacts across many networks and affect tens of thousands of users who may be many miles from the event itself.* These network problems can be caused by natural events such as hurricanes or earthquakes or can be man-made, like malicious distributed denial-of-service attacks and network viruses, or accidental "back-hoe" incidents that cut major backbone fiber-optic cables (Barrett 1999, Delio 2001). A third problem is that of handling unexpected surges in traffic in response to high-profile news events (Ewalt 2001, Manjoo 2000).

**Research is showing that the Internet is surprisingly vulnerable to disruption despite its decentralized nature. ("The Achilles' Heel of the Internet" 2000).*

In tackling these operational challenges, maps of network architecture and performance can be vital tools for managers and engineers. Maps can summarize and present complex, rapidly changing data on the operational state of a network in a single visual image, providing an easy-to-interpret overview of the system and, thereby, aiding in the diagnosis and solution of problems. For example, at network operation centers (NOCs) of large ISPs, just a handful of skilled operators using maps are responsible for keeping a complex and geographically distributed hardware infrastructure running smoothly (See Figure 1.) (Koutsofios et al. 1999, Wei et al. 2000). As a *New York Times* story noted on the huge stress on the U.S. telecommunications systems immediately following the attacks of September 11, 2001, "By watching computerized maps of the United States, [operators] can tell in an instant whether there are any jams in long-distance traffic." (Guernsey 2001).

However, the detailed network monitoring maps and tools used by operators in NOCs are not made public for reasons of security and commercial confidentiality. Also, most of these maps are not designed as general-purpose maps that can be read by the general-public. Instead, they are specialized management tools that require skilled interpretation. That said, some Internet networks, particularly those serving the research and education communities, do make summary network performance data publicly available using map interfaces. These interfaces are popularly referred to as "network weather maps." The maps are public-spirited information dissemination tools that provide network customers (usually universities and labs) with useful information (especially on how to identify trouble spots) and that can also have a marketing function (see next section).

Figure 1
View of AT&T's Large NOC with Large Wall Displays Showing Network Maps

Source: Wei et al. 2000:2

Below are two examples of network weather maps—the Abilene network in the United States (See Figure 2.) and NORDUnet serving Scandinavia (See Figure 3.). The maps are updated frequently (for example the Abilene map is updated every five minutes), allowing users a "peak inside" the network. Both maps provide a summary of overall network performance with links color coded by their traffic flows, but importantly they also provide an interactive, visual interface through which to browse more detailed performance statistics available as tables and statistical charts (accessed by clicking on links on the map).

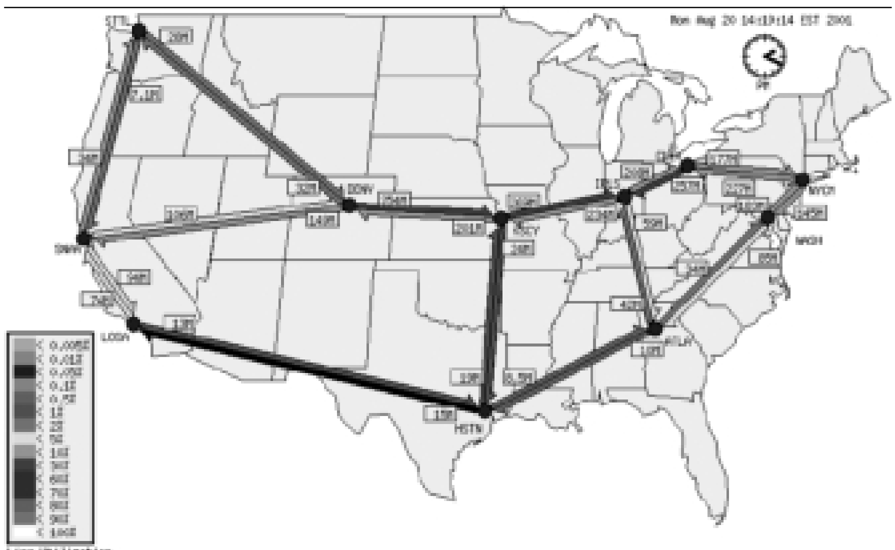

Figure 2
"Weather Map" of the Traffic Load on the Core Links of the Abilene Network

Source: Abilene Network Operations Center, Indiana University <http://hydra.uits.iu.edu/~abilene/traffic/>

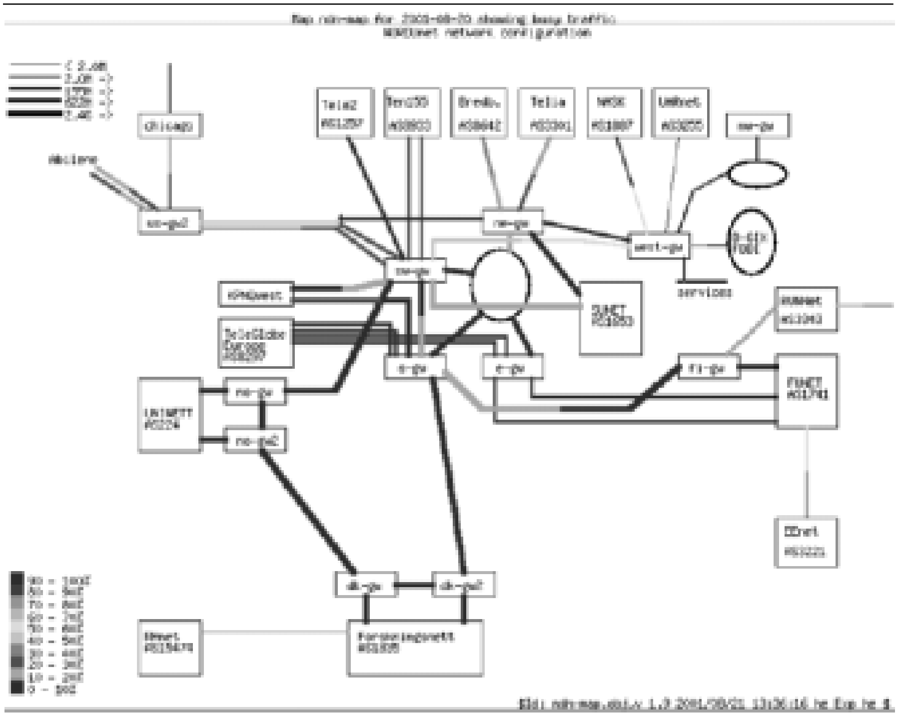

Source: NORDUnet <http://www.nordu.net>

Figure 3
"Weather Map" of Network Load for the NORDUnet Network

These two maps also illustrate the two major cartographic archetypes employed to represent computer networks—showing linkages and nodes either on a geographic base with a familiar template of cities and administrative boundaries or as a logical schematic. These maps can often be highly generalized, with, for example, the network architecture shown as straight lines, although they are topologically correct (as with conventional subway maps).

In addition to single network maps, there are also some attempts to provide dynamic "weather" maps of Internet-wide performance. For example, Matrix.Net's Internet Weather Report (IWR)* presents maps of network latency at many locations across the world using automated large-scale measurement of the Internet taken every four hours. Running continuously since 1993, IWR gives one of the few consistent, time-series measurement of global Internet performance (Quarterman et al. 1994). Figure 4 shows a frame from an animated IWR map at the global scale. Forecasts are made six times a day, every day of the year, for over 4,000 Internet sample points all around the world. These forecast measurements are turned into maps with graduated circle symbols representing latency (the larger the circle the longer the delay). In basic terms, small circles on the map show a healthy Internet, while large circles are indicative of poor performance and possible problems.

**<http://www.matrix.net/research/weather/>*

Another method for monitoring network performance are "traceroutes," Internet utilities that allow the active monitoring of real-time data routing and

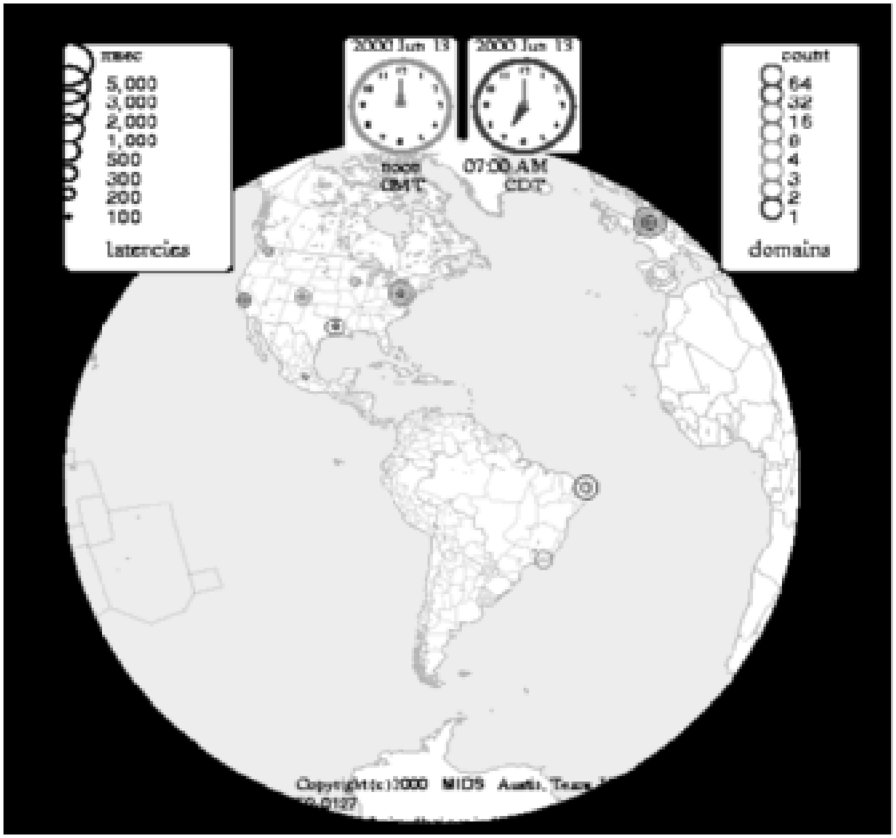

Figure 4
One Frame from the Animated Internet Weather Map

Source: Matrix.Net <http://www.matrix.net>

the “debugging” of connectivity problems. Traceroute works by reporting the routes that data packets travel through the Internet to reach a given destination and the time taken to travel between all the nodes along the route (Rickard 1996, Dodge 2000a). Traceroutes reveal the hidden complexity of data flows, traversing ten, twenty, or more nodes, seamlessly crossing oceans and national borders and moving through networks often owned and operated by competing companies, to reach a given destination. A typical output of the basic traceroute utility is shown in Figure 5. Each line in the output of traceroute represents a single “hop” the data takes through the Internet. In this case the data route took 23 hops to reach its destination. Each hop is generally a separate physical node consisting of a network switch or a router. The approximate locations of this routing hardware can also be plotted on a map to give a geographic traceroute, an example of which is given in Figure 6.

The physical infrastructure of the Internet is largely invisible to the casual observer because it is built into the fabric of buildings and under roads. Nevertheless, it has to be installed in the first place and subsequently maintained and upgraded. Highly detailed large-scale maps and plans of the physical infrastructure are routinely used for keeping track of network

```
Tracing route to walnut.may.ie [149.157.1.115]

over a maximum of 30 hops:

1  <10 ms  10 ms <10 ms  209-9-224-225.sdsl.cais.net[209.9.224.225]
2   30 ms  90 ms  50 ms  172.20.0.1
3  <10 ms  10 ms  10 ms  fe7-7.core1.mcl.cais.net [63.216.0.77]
4  <10 ms  20 ms  10 ms  pos3-2.core1.wdc.cais.net [63.216.0.69]
5   10 ms  10 ms  10 ms  pos3-0.core2.wdc.cais.net [63.216.1.14]
6   20 ms  30 ms  30 ms  pos5-3.core.pitt.cais.net [63.216.1.62]
7   40 ms 130 ms  30 ms  pos5-0.core1.pitt.cais.net [63.216.6.13]
8   50 ms  40 ms  60 ms  pos5-3.core.det.cais.net [63.216.7.58]
9   40 ms  40 ms  40 ms  pos5-0.core1.det.cais.net [63.216.8.13]
10  50 ms  70 ms  50 ms  pos5-2.core.chi.cais.net [63.216.8.58]
11  90 ms  81 ms  70 ms  uunet.a3-0.4.core2.chi.cais.net [63.216.9.65]
12  60 ms  70 ms  60 ms  0.so-5-1-0.XL1.CHI2.ALTER.NET [152.63.67.242]
13  50 ms  60 ms  80 ms  0.so-7-0-0.XR1.CHI2.ALTER.NET [152.63.67.130]
14 150 ms  60 ms 121 ms  0.so-3-0-0.TR1.CHI2.ALTER.NET [152.63.15.86]
15  80 ms 100 ms  70 ms  126.at-4-0-0.IR1.NYC9.ALTER.NET [152.63.1.121]
16  80 ms  70 ms  90 ms  so-1-0-0.IR1.NYC12.ALTER.NET [152.63.23.62]
17 131 ms 140 ms 190 ms  so-5-0-0.TR1.LND9.Alter.Net [146.188.15.49]
18 130 ms 141 ms 170 ms  pos0-1.cr2.dub2.gbb.uk.uu.net [158.43.253.58]
19 141 ms 120 ms 160 ms  ge0-0-0.gw4.dub2.gbb.uk.uu.net [158.43.152.6]
20 130 ms 151 ms 120 ms  158.43.111.102
21 161 ms 180 ms 140 ms  Oswald-f1-1.dublin.core.hea.net
                         [193.1.195.137]
22 151 ms 200 ms 170 ms  Uther-g1-0-0.dublin.core.hea.net
                         [193.1.195.242]
23 211 ms 180 ms 190 ms  nuim-kinnegad.atm.link.hea.net [193.1.194.22]
24 161 ms 200 ms 160 ms  walnut.may.ie [149.157.1.115]

Trace complete.
```

Figure 5
Traceroute Listing of Real-Time Internet Route Taken by Data between a PC in the Washington D.C. Area and a Web Server Located Just Outside Dublin, Ireland

architecture, for example schematics of the exact cable routes are needed by the engineers who actually drill the holes and dig up the roads. Here, CAD, AM/FM, and cable management systems that use spatial databases and map-layer representations are widely used (Fry 1999). However, these maps are generally not available to the public.

Maps For Internet Marketing

A large number of infrastructure maps of the different Internet networks have been produced primarily for the purposes of marketing. Indeed, a cursory examination of most any ISP Web site will reveal “high-gloss” marketing maps. This is, perhaps, not surprising as maps have long been created in the service of marketing and promotion (Tyner 1982, Monmonier 1991). Geographic maps can be seen in some senses as the natural visual representation of transportation and communications networks, able to effectively show potential customers how a particular network could expedite their travel needs. As a consequence, there is a long (dis)honorable tradition of promotional maps being used to highlight the advantages of the latest transportation

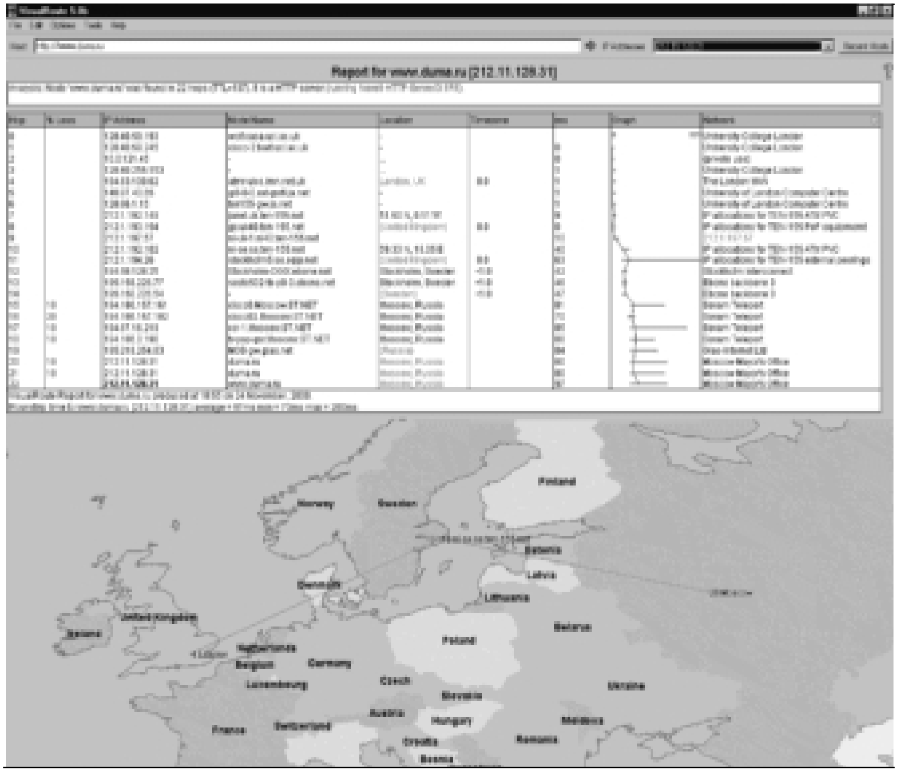

Figure 6
An Example of a Geographic Traceroute Using the VisualRoute Utility. The Internet Route is between London and the Russian Duma Web Site in Moscow.

network such as canals, oceanic shipping lines, railroads, highways, and, of course, airlines (Ackerman 1993, Fleming 1984).

Given that the provision of Internet network services is a highly competitive business, dominated by large corporations many of whom operate globally, effective marketing is a vitally important activity. Here, maps are employed to provide a selective and positive view of a network, emphasizing its extent (e.g., demonstrating the geographic reach of the network, emphasizing all the distant places that are linked together) and capabilities (e.g., illustrating the tremendous capacity of the “pipes” of the network to cope with huge user demands) in order to attract and compete for custom. In many respects, Internet network provision is such an intangible commodity that the map is powerful in making it seem more “real.” The maps generally show a generalized and simplified view of the network, usually in a bright, colorful, and visually effective manner. Most often the maps are drawn on a template of real-world geography and have many design commonalities with the airline route maps in the back of in-flight magazines.

While these maps do provide a selective picture, a reflection of what the company wants to emphasize, they also allow academic researchers and others to chart the range and make-up of each company’s network, to document different kinds of provisions at a range of scales, and, importantly, to note how this has changed over time. For example, Gorman and Malecki

(2000), Moss and Townsend (2000), and Wheeler and O'Kelly (1999) have undertaken useful analyses of the geography of Internet network topology based on data gathered, in part, from ISP marketing maps. This can be illustrated in reference to an analysis of UUNet's (part of Worldcom) infrastructure. Growing at over a self-reported rate of 1000 percent per year* a longitudinal study of their maps at a variety of scales allows us to see the company's strategy for delivering infrastructure services and to project the likely consequences of this strategy on issues such as the digital divide, urban-regional restructuring, local and regional economic development, and so on. (See Figure 7.) What is clear from these maps is that UUNet is a global supplier of network services, and that the network is confined to the three main pan-regional trading zones (North America, Europe, and Asia), and to the principal cities (hubs) in these regions that are most likely to hold potential customers. Lower-level cities have lower capacity linkages, and other potentially less profitable areas and cities are bypassed altogether (e.g., most of middle America).

**The 1000 percent figure might well be apocryphal and has been disputed, see for example Odlyzko (2000).*

Maps For Strategic Planning And Policy

There is a long history of using maps as instruments of planning and policy. Maps have been key strategic devices used in planning and implementing urban and regional development, plotting military strategy and the conquest of new lands, and legally contesting land ownership and use. Unsurprisingly, then, they are also being used in the short- and long-term strategic planning of Internet development by commercial enterprises, governmental, quasi-governmental, and other interested bodies (e.g., the Internet Society). That said, the extensiveness and impact of their use is difficult to gauge quantitatively. While we give several examples where maps have been used, we suspect that their full potential is not yet being realized (this is based on the fact that we could locate relatively few examples of where maps had been used as key analytical resources). This under-usage is, we suspect, because there is a perception that the Internet is somehow non-material in substance, due to its mode of interaction, and the relative invisibility of infrastructure. In addition, data to create useful maps is often closely guarded by service providers and its use restricted from the public domain, and other forms of data generation are costly and technically difficult. In order to structure our analysis, we have divided our discussion into two related themes. The first concerns the planning and development of infrastructure, the second, regional development, the attraction of inward investment, and the monitoring and addressing of inequalities.

At one level, maps have been used in the planning, development, and expansion of network infrastructure at a variety of scales from individual buildings to global networks. Planning the optimum topology for a communications network to efficiently interconnect geographically dispersed loca-

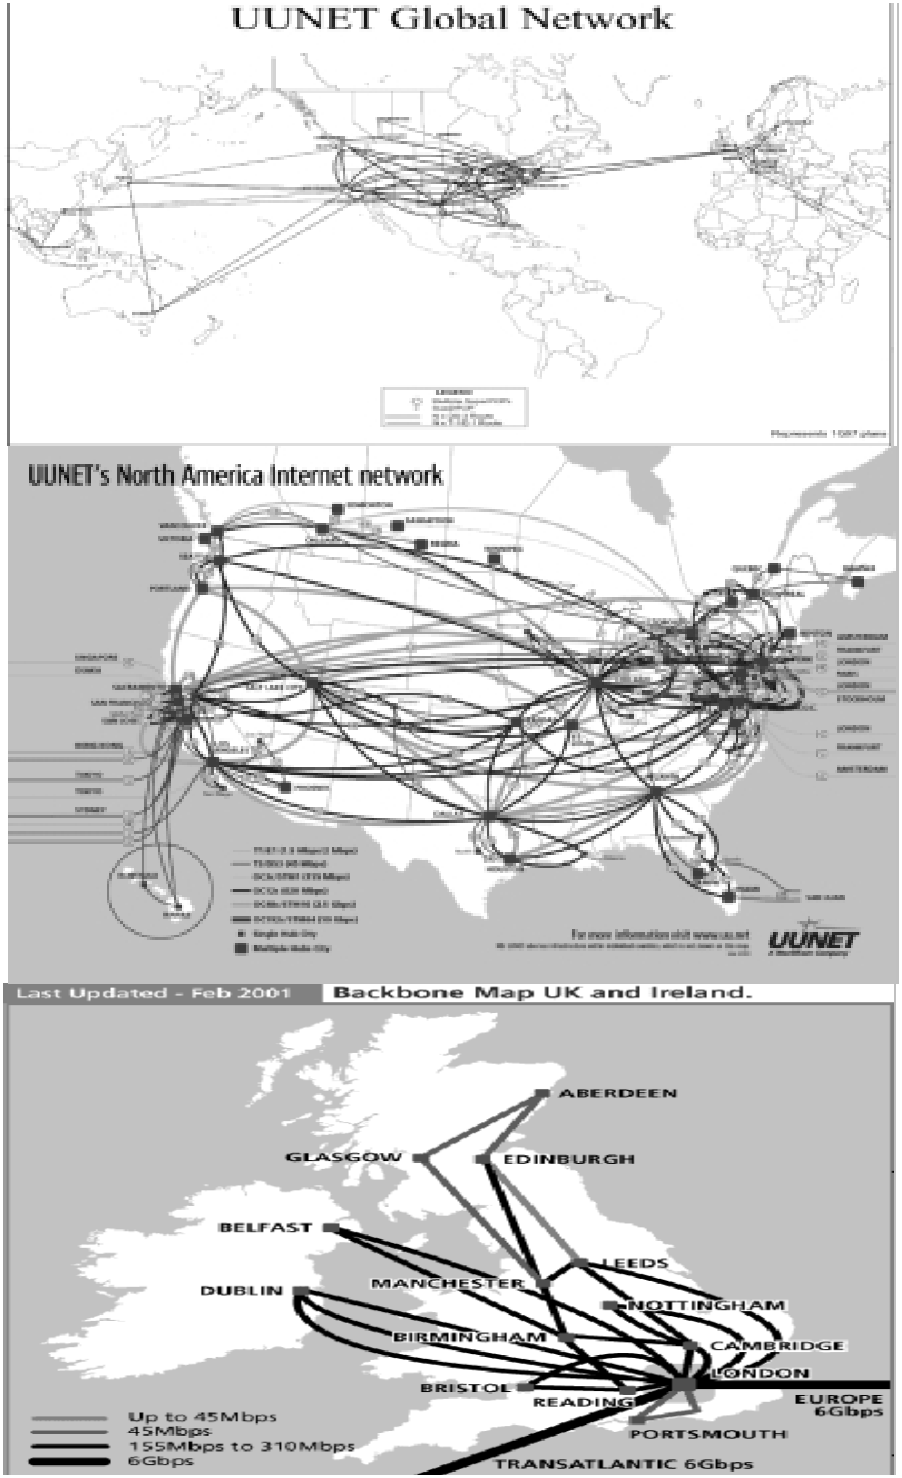

Source: UUNET <http://www.uu.net/>

Figure 7
Sample Marketing Maps Showing the Internet Network of UUNET, One of the Largest Providers: (a) UUNET Global Network as of First Quarter 1997, (b) U.S. Network from June 2000, and (c) the United Kingdom and Ireland, February 2001.

tions is an exacting task. Maps help visualize complex network topologies and how new configurations will look and operate. Figure 8 is a “back of the envelope” hand-drawn sketch map from the early planning of ARPANET,* plotted by the project manager Larry Robert in the late 1960s. It shows the

**ARPANET pioneered wide-area packet-switching networking and laid much of the foundation of the Internet as we know it today, developing both the technical and social infrastructure of internetworking (Abbate 1999, Hafner and Lyon 1996). It was initially conceived as a method to link several incompatible computer systems located at various points across the USA so that resources could be shared and was funded by the US military, through the ARPA agency.*

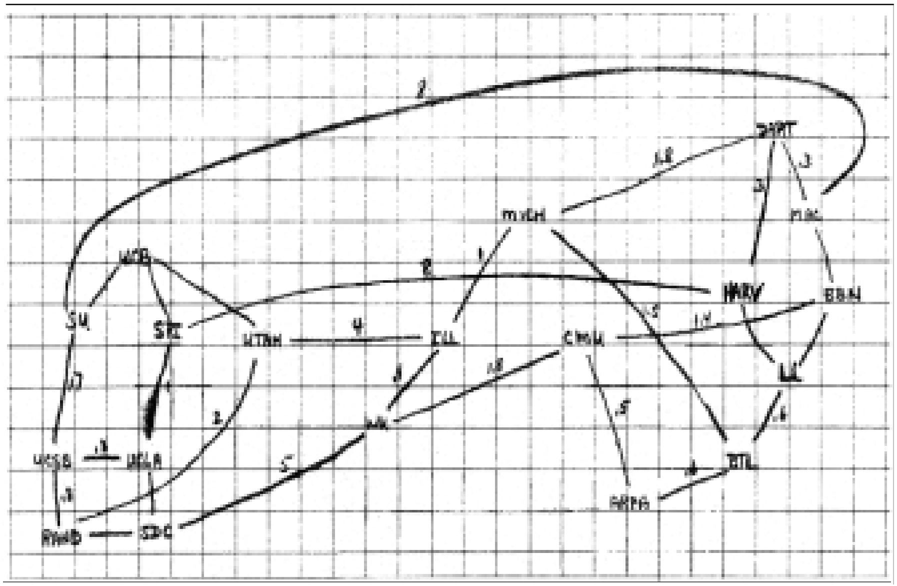

Source: Hafner and Lyon 1996:50

Figure 8
A “Back of the Envelope” Style Sketch Map for Network Topology Planning

projected topological routing of the fledgling Internet between nodes. Figure 9 shows the fiber-optic cable routing in downtown Philadelphia, a city home to 270 technology firms in 2001, 60 per cent of which were located in the center city, requiring high-speed Internet connections. Many of these companies are members of ePhiladelphia Technology Alliance, an organization dedicated to creating and fostering a vibrant technology community within the city. By mapping companies in relation to cable-routing, the city can adequately provide network connections and plan extensions that will hopefully attract new customers. At a larger-scale, countries are crisscrossed by many interconnected networks. An important function for ISPs is to easily and efficiently interconnect and exchange local traffic at neutral peering points. Figure 10 shows two examples of national-level maps tracking the Internet infrastructure in the Republic of Korea produced by Korean Network Information Center (KRNIC), based in Seoul (http://stat.nic.or.kr/english/netowrk.html). Analysts at KRNIC have produced a whole series of maps over the past five years using topological graphic representations. The two maps clearly reveal the tremendous growth in the number of ISPs, their interconnections, and the capacity of links within and external to Korea. The maps are valuable policy and research resources creating a census of the growing complexity of the links between ISPs and their capacity.

Figure 9
Fiber-Optic Routes in Central Philadelphia

Source: Central Philadelphia Development Corporation <http://www.centercityphila.org/it.html>

At a second level, maps have been employed in the strategic planning and implementation of regional development and in monitoring and addressing inequalities, the so-called digital divide, between places. Again, the data relates to several scales from intra-urban to global. As widely documented, cities are increasingly becoming competitive enterprises, vying to attract investment of the high-tech sector (Graham and Marvin 2001). Here, the "where" of infrastructure is important, with decisions about structural investment tied to a city's economic future. Here public-private partnerships between city government, commercial ICT infrastructure companies, a range of economic and public policy consultancies, and local development and community groups seek to maximize their connectivity within optimal constraints (e.g., profit). Maps are a potentially important tool for illustrating high-capacity Internet infrastructure to potential inward investors and encouraging economic development. Examples include the "Bandwidth Bay Fiber Network Mapping" (http://www.bandwidthbay.org/main.htm) by the City of San Diego (Abouna 2001; Figure 11) and the "Georgia High-Speed Telecommunications Atlas" (http://maps.gis.gatech.edu/telecomweb/index.html) in the state of Georgia, USA (French and Jia 2001; Figure 12).

In addition to being used as ways to gain competitive advantage (exploiting the differences between cities), paradoxically these data can also be used to close the digital divide within cities. Indeed, it is a policy of most Western governments at this point to try to ensure widespread access to the

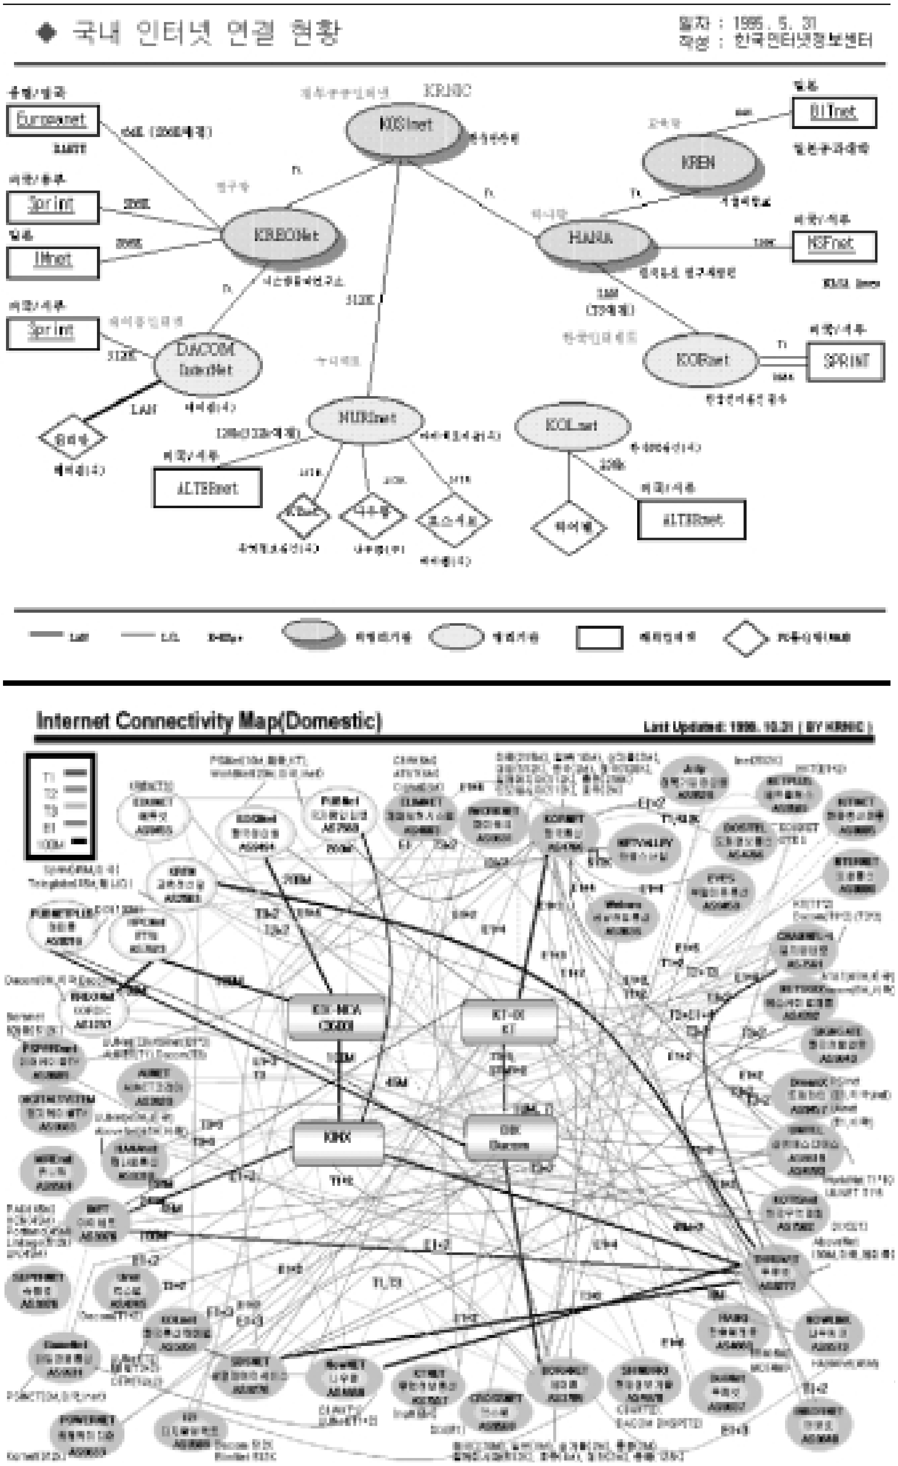

Source: Korean Network Information Center < http://www.nic.or.kr/>)

Figure 10
Topology Maps of ISP Interconnections in the Republic of Korea from (a) May 1995 and (b) October 1999

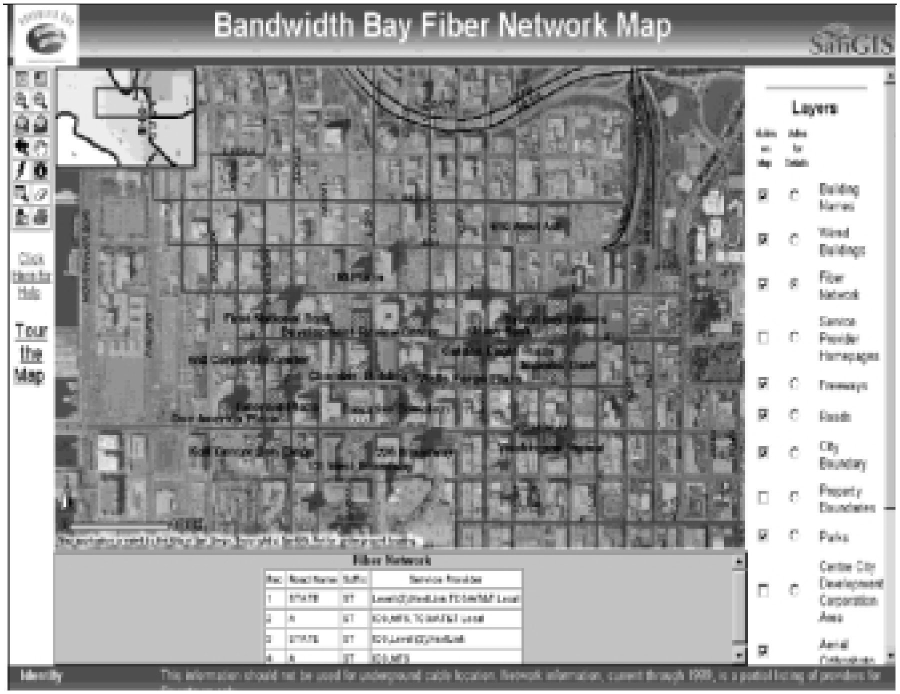

Figures 11
Map of the Internet Fiber-Optic Networks and Wired Buildings in Downtown San Diego from the Bandwidth Bay System

Source: San Diego Geographical Information Source (SanGIS) <http://www.bandwidthbay.org/main.htm>

Internet so that communities, at all scales—local, regional, national—are not left too far behind. For example, two federal U.S. schemes designed to

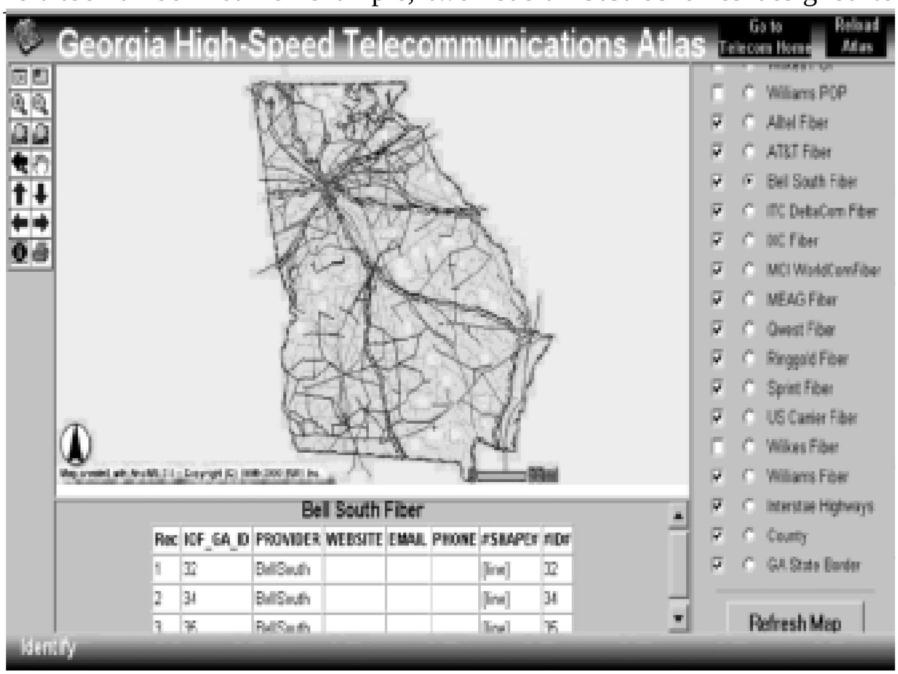

Figures 12
Map of the Commercial Networks Infrastructure in Georgia, USA

Source: Center for Geographic Information Systems <http://maps.gis.gatech.edu/telecomweb/index.html>

facilitate connecting disadvantaged communities to the Internet are the Community Technology Center (CTC) programs (Office of Adult and Vocational Education, Department of Education <http://www.ed.gov/offices/OVAE/CTC>) and the Technology Opportunity Program (Department of Commerce <http://www.ntia.doc.gov/otiahome/top/>). These are supplemented by a wide range of other programs at the state and city levels. For example, the Federal Communications Commission (FCC) (the U.S. telecom regulator) is concerned with issues of access and equity for different communities. Public policy makers obviously recognize that wiring areas requires considerable infrastructure investment on the part of commercial providers, with the incentive to concentrate on those areas most likely to return an operating profit. Consequently, regulators are concerned that planned high-speed Internet delivery systems are available, at affordable costs, to all members of a community, in particular, low-income communities or those in more sparsely populated rural areas. Clearly, here, the geography of access is crucial and one strategy open to regulators to make visible inequalities "on the ground" is to make use of maps that show spatial patterns of broadband Internet availability. Figure 13 provides two examples, at different scales, from a recent Federal Communications Commission (FCC) report on broadband Internet access. The first map shows the number of broadband providers for ZIP-code areas across the whole of the United States, while the second map focuses just on the local geography of DSL coverage in Los Angeles county, California.

These maps were part of a large report on the FCC regulatory monitoring of providers to insure that they met the provision of the 1996 Telecommunications Act that encouraged the deployment of advanced telecommunications capability to all Americans in a reasonable and timely fashion. The general conclusion of the report, supported by tables and maps, was that commercial providers were generally meeting targets with 59 percent of the U.S. ZIP codes (which represent 91 percent of the resident population) showing evidence of high-speed Internet access. However, they also issued one crucial caveat:

> …the data support the troubling conclusion that market forces alone may not guarantee that some categories of Americans will receive timely access to advanced telecommunications capability. We identify certain categories of Americans who are particularly vulnerable to not having access to advanced services. These include low-income consumers, those living in sparsely populated areas, minority consumers, Indians, persons with disabilities, and those living in the U.S. territories. (U.S. Federal Communications Commission 2000).

In addition, in the United States, the Census Bureau, the Department of Commerce's National Telecommunications and Information Administration,

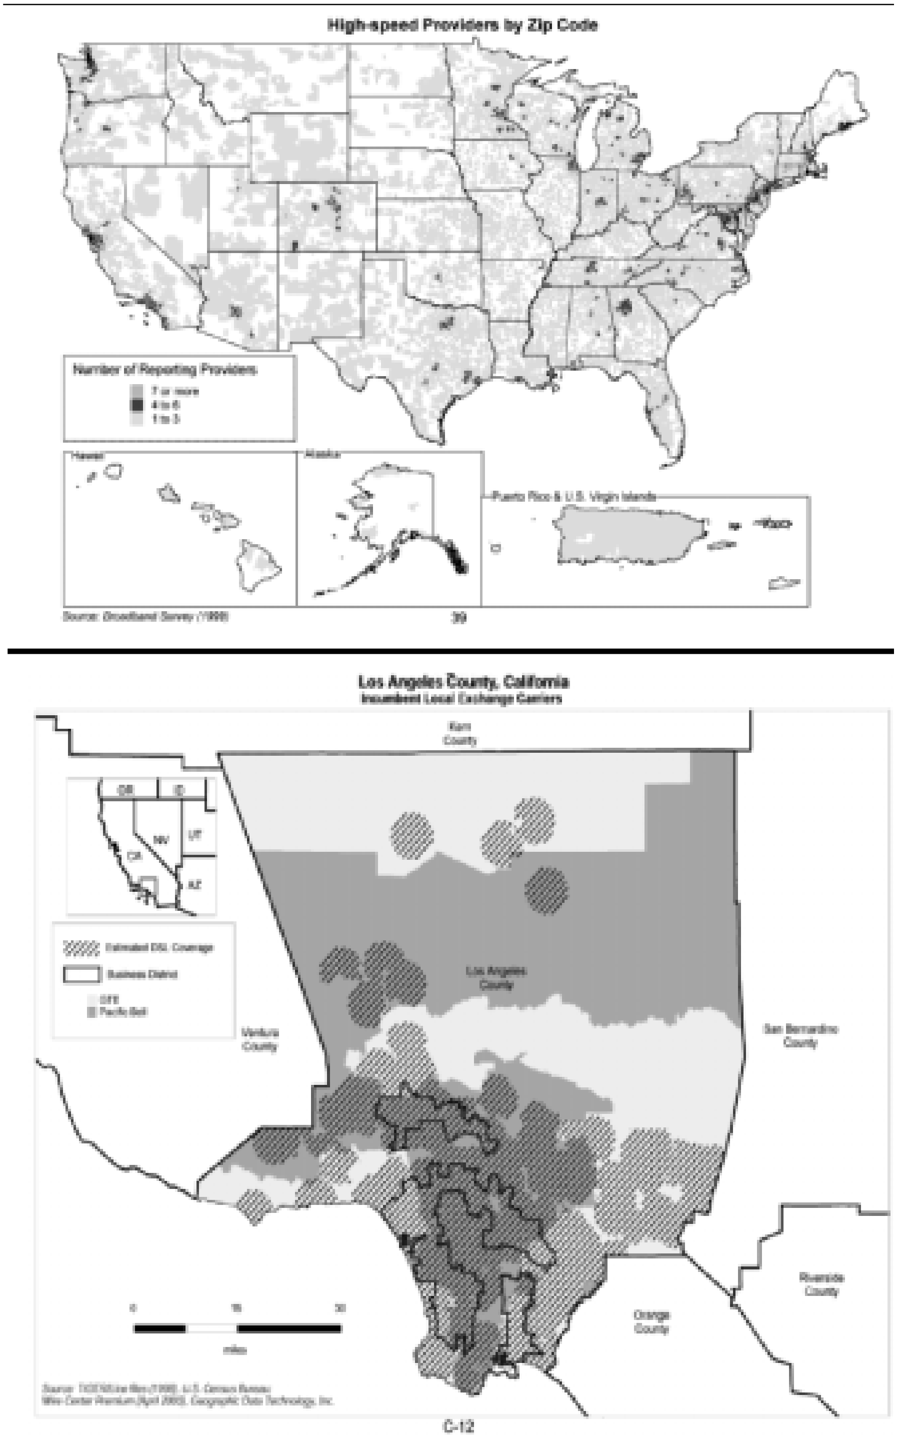

Figure 13
Maps of Broadband Provision in (a) the United States and (b) Central Los Angeles

Source: Federal Communications Commission <http://www.fcc.gov/broadband/>

and the Economics and Statistics Administration generate official statistics on Internet and telecommunications access at national and regional scales that are analyzed for their economic policy potential by a range of groups, including local and state governments and commercial companies. For example, the Progressive Policy Institute uses a range of these data in formulating their “New Economy Index” reports (Atkinson 2002, Atkinson and Gottlieb 2001, Atkinson et al. 1999). The self-stated aims of these reports are to “…offer policy makers a framework for economic development strategies aimed at promoting fast and widely shared economic growth and prosperity.” Maps are used prominently throughout the report and Figure 14 shows an example mapping the online population, from *The Metropolitan New Economy Index* (Atkinson and Gottlieb 2001) for the top 50 metropolitan regions in the United States. These are grouped into four percentile groups. Other maps in the report rank the regions according to 16 indicators that are used to create an overall index of economic competitiveness in the information economy.

Source: Progressive Policy Institute

Figure 14
Map of the Top 50 U.S. Metropolitan Areas in Terms of Online Population

Likewise, Mark Krymalowski has been analyzing data at the country level, plotting the geographical distribution of *.de* domain registrations in Germany.* Figure 15, drawn from his research, shows the relative number of domains per capita in 2000 for German counties. Krymalowski’s analysis and maps of domains were subsequently used in analyzing high-tech, economic, and regional development (Sternberg 2001). This analysis concluded that the city of Munich and its wider region scored much more heavily in domain name registrations than would be expected simply based on population. This, he hypothesized, is because this region is the leading zone of IT and multimedia production in Germany. Importantly, Sternberg concludes that, “the Internet does not create new regions but it replicates, at least in Germany, the well-known ranking of regions in terms of high-tech” (Sternberg 2001:3). In other

**<http://www.denic.de/doc/DENIC/presse/stats2000.en.html>*

words, the information economy is likely to grow most quickly around existing IT hubs, rather than invest in new, potentially cheaper, locations. This clearly has implications to regional development designed to address regional inequalities and attract inward investment given the widespread shift towards an IT-centered service economy.

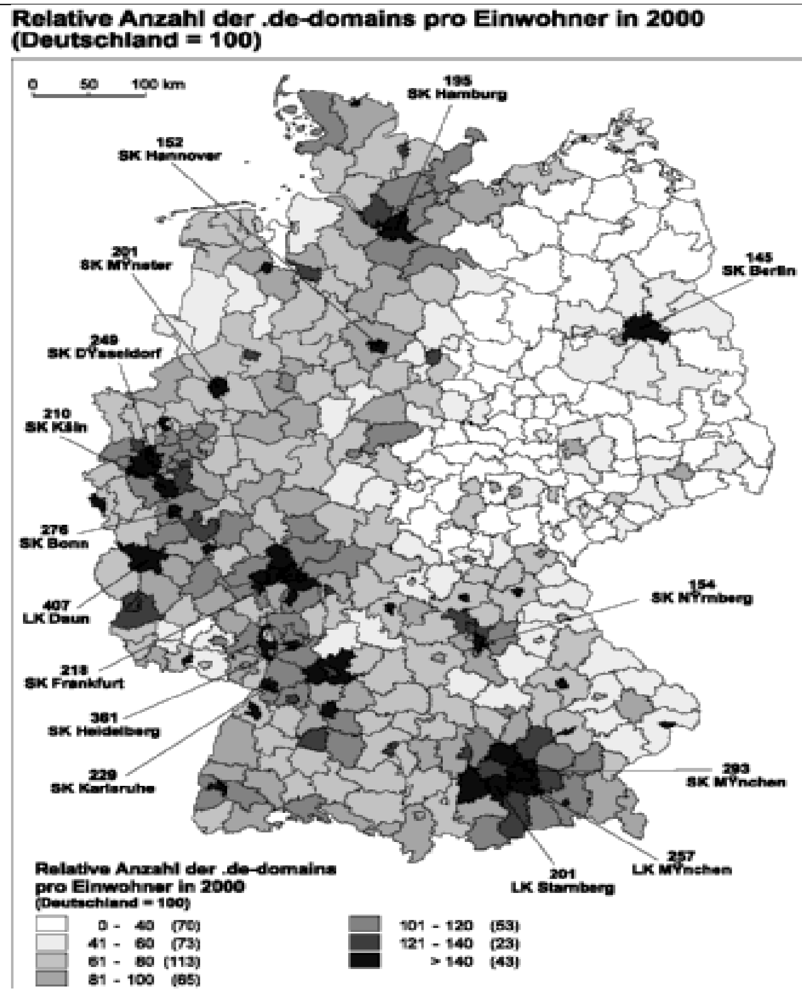

Figure 15
Per Capita Measure of .de *Domains in German Counties*

Source: Mark Krymalowski and DENIC <http://www.denic.de>

These kinds of maps, when put together in a timeline, form a powerful means for tracking development and for predicting future change. One project that illustrates this is that by Larry Landweber, and several organizations that have taken his lead to produce longitudinal maps at different scales (e.g., TeleGeography and Matrix.Net), used by both governments and commercial enterprises to formulate strategies of investment. During the 1990s, the Internet spread across the globe so that by the end of the decade virtually all nations were connected (although the number and capacity of connections still varies greatly). This global diffusion of the Internet was tracked by Landweber and charted in a series of maps (See Figure 16.), providing a useful baseline census for policy of the spread of international network connectivity (Dodge 2000b). Countries are shaded according to a four-fold classification of network connectivity, with permanent Internet the darkest shade.

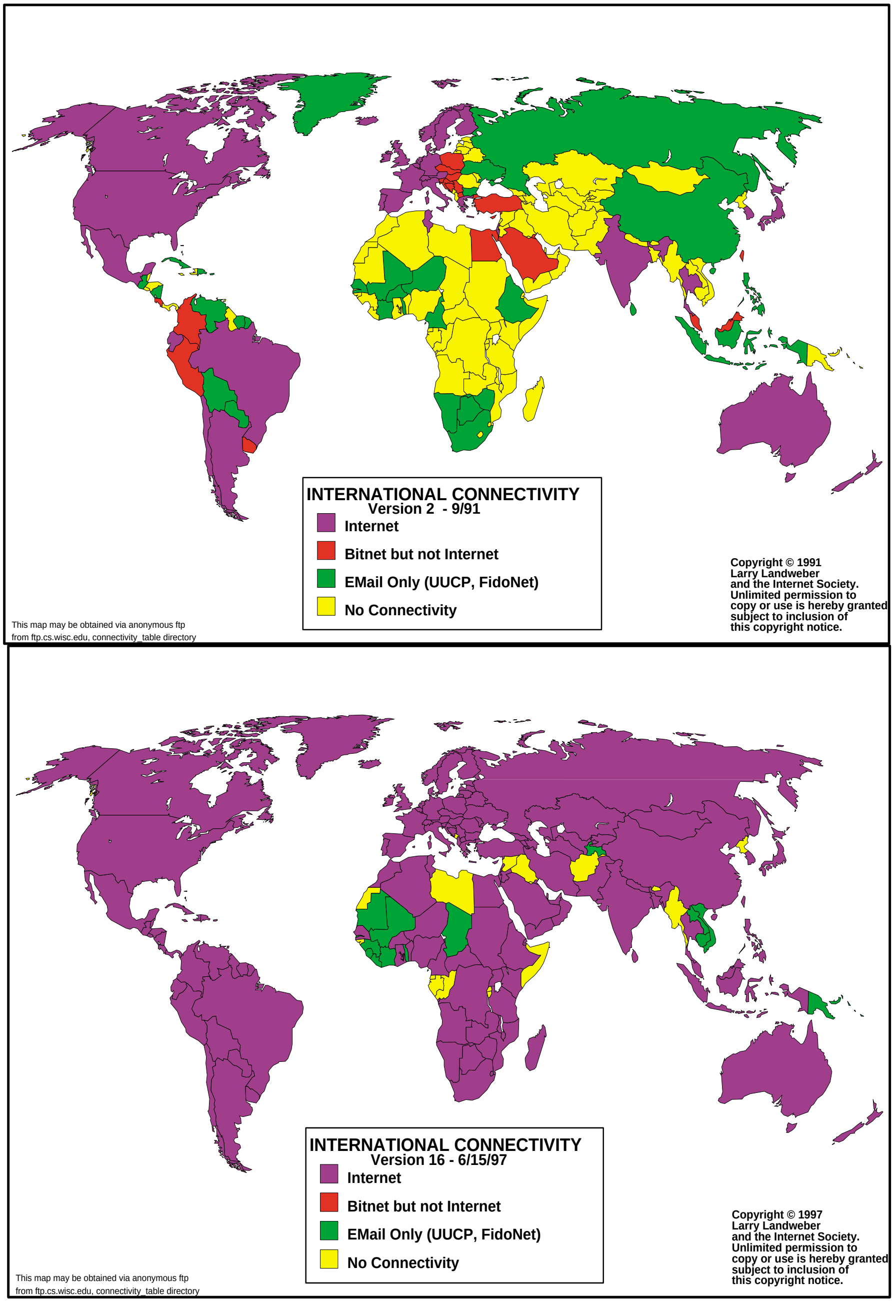

Source: Larry Landweber and the Internet Society <http://www.cs.wisc.edu/~lhl/maps/>

Figure 16
Maps of the Global Diffusion of Internet Connectivity at the National Level by Network Infrastructure from (a) 1991 and (b) 1997

These maps provide a partial, but useful, picture of global Internet diffusion through the 1990s. The first map, from 1991, shows that a large number of countries, particularly in the Americas and in Northern Europe, had full Internet connectivity. However, an equally large measure of the world's nations are shaded light gray, indicating that they had no international Internet connectivity. In fact, this category included well over half the nations of the world, though these were clearly concentrated in the less developed regions of Africa and central Asia. By 1997, the majority of the nations of the world were the darkest shade. The Internet, as measured by Landweber's survey, was so widespread that the exceptions really stand out. (It was at this point that tracking diffusion at this scale using Landweber's criteria became redundant and, hence, this is the last map in the series). The light-shaded exceptions were nations suffering from extreme poverty, war, and civil conflicts (such as Afghanistan and Somalia) or from external geopolitical isolation (e.g., Libya, North Korea, Burma, Iran, and Iraq).

Maps For Academic Internet Analysis

It has been widely argued by academics that the ICTs are transformative technologies that are having significant impacts on social, economic, and political life, engendering widespread changes (e.g., Castells 2000, 2001, Graham and Marvin 2001, Kitchin 1998). The process of mapping has been used to comprehend three different sorts of projects aimed at furthering our understanding of these changes in relation to infrastructure: urban-regional restructuring, the digital divide, and measuring the Net.

As noted above, maps reveal visually the nature and extent of the "digital divide" in society. They have, therefore, been used by a number of academics such as Holderness (1998), Moss and Townsend (1997, 2000), Sternberg (2001; see above), and Zook (2000, 2001) to explore social and economic divides in access to Internet infrastructure at a variety of scales. For example, Matthew Zook has analyzed the spatiality of the Internet content production industries in the United States through the detailed mapping of the geographic location of domain name registrations at different scales. (See Figure 17.) Just as postal addresses in geographic space identify a unique location, domain names perform the same function for the Internet, allowing users to visit the site. Importantly, the geographic location of the owner of these domains can be determined from registration databases, which have a billing postal address, containing ZIP codes that can easily be mapped to street-level locations using off-the-shelf GIS software and map data. Figure 17 displays maps for downtown San Francisco using proportional map symbols, with background road and town data to add context. The densest concentration of ZIP codes are located in the financial district and "South of Market" area (famed as the Multimedia Gulch). This mapping led Zook to conclude that the

"Internet industry exhibits a remarkable degree of clustering despite its reported spacelessness" (Zook 1998:18). This approach provides a valuable quantitative measurement for policy analysis on Internet economic activity and reveals where the connections are and where they are not (Zook 2000, 2001).

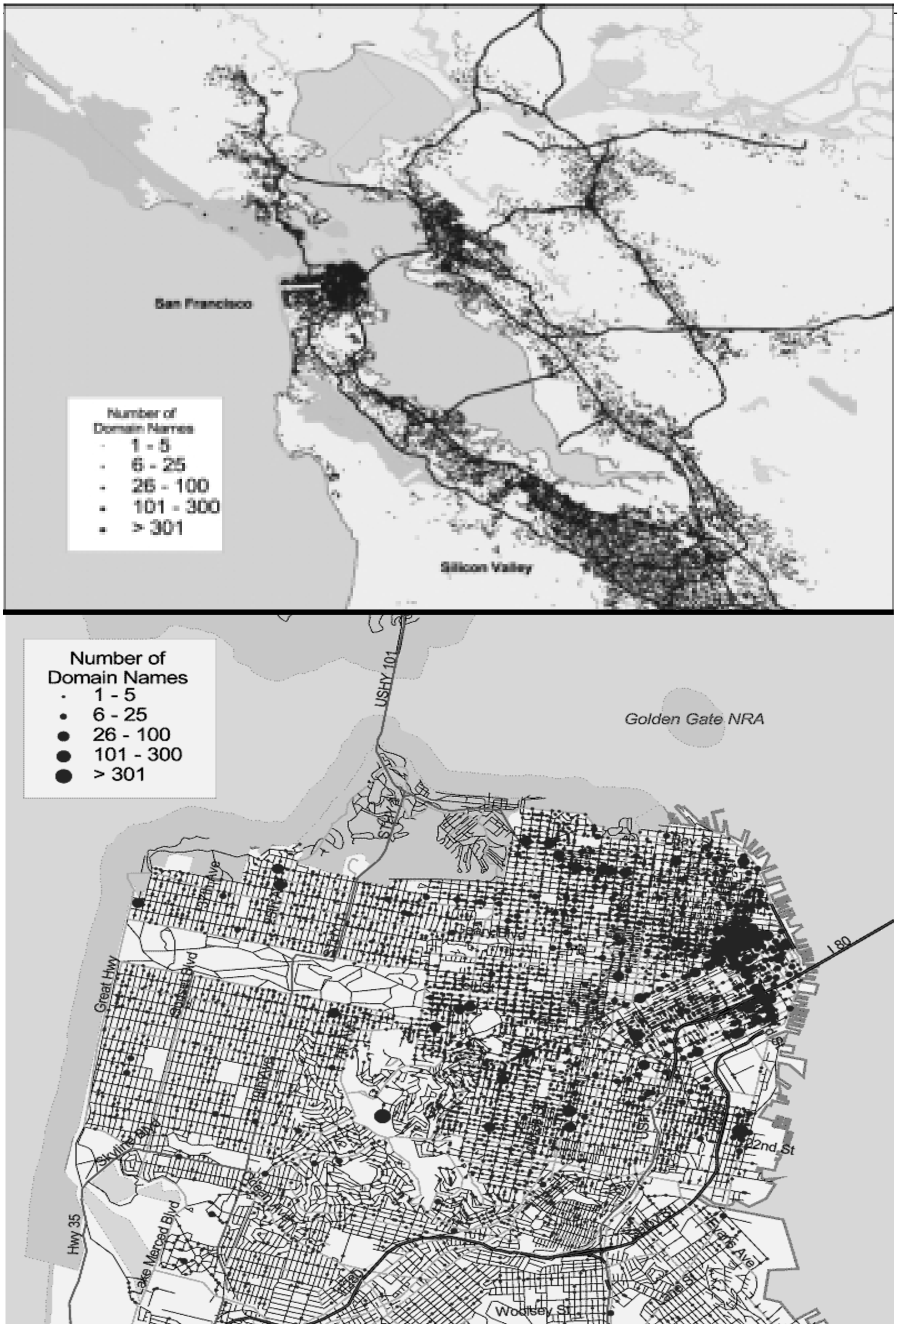

Source: Matthew Zook <http://www.zooknic.com/>

Figure 17
The Number of Domain Names in (a) the Bay/Silicon Valley Area of California, (b) San Francisco, 1999

The final way that maps have been used by academics and commercial research teams is as a means by which to display measurements that quantify the extent and use of Internet infrastructure so as to gain a better understanding of its distribution, diffusion, and utilization. Maps have particular appeal because they reveal discernable patterns in very large data sets and so provide panoramic overviews of where changes are occurring. To date, a number of such mapping projects have been instigated (see Dodge and Kitchin 2000a, 2001) and here we discuss three in brief.

Figure 18 displays an "arc map" of Internet traffic flows between fifty nations, from February 1993. The color, thickness and height of the arcs are used to encode the traffic statistics for particular inter-country routes (Becker et al. 1995, Cox et al. 1996). The arcs are also partially translucent so as not to completely obscure lines at the back of the map, while their height above the base map is in relation to total volume of traffic flowing over a link. This has the effect of making the most important (high traffic) links, the highest and therefore most visually prominent on the map. In the SeeNet3D application in which the image was generated, the user had considerable interactive control and was able, for example, to vary the arc height, scaling, and translucency. The map could also be rotated and scaled, so that the user could view it from any angle. The map shows that there was significant traffic, in the early 1990s, between three areas of the world, North America and Europe, Europe and Australiasia, and Australiasia and North America, with most traffic crossing the Atlantic. The map does not show all traffic, however, because it is limited to just fifty countries. As such, it portrays a selected image, one that is dominated by developed countries that were the principal nations connected to the Internet in 1993.

Figure 18
Interactive Visualization of Internet Traffic in the See Net3D Network Analysis Application

Source: Stephen Eick, Visual Insights

Figure 19 is a 3-D, interactive geographic visualization of the Internet MBone network (Munzner et al. 1996). The MBone comprises a special set of routes, known as “tunnels” in technical jargon, which run on top of the ordinary Internet and are used to deliver multicast data. Multicasting is an Internet protocol designed for delivering efficiently a single copy of a chunk of data to many different people. It is especially useful for distributing real-time audio and video. Munzner and her colleagues map these tunnels as arcs on a 3-D model of the globe, which the user can manipulate to rotate and view from any angle. Line color 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. Before their mapping, it was very difficult to determine the extent of MBone infrastructures because they were created by several different organizations and their characteristics were documented using text listings (some seventy-five pages in length in June 1996) from which it was very difficult to determine the topology.

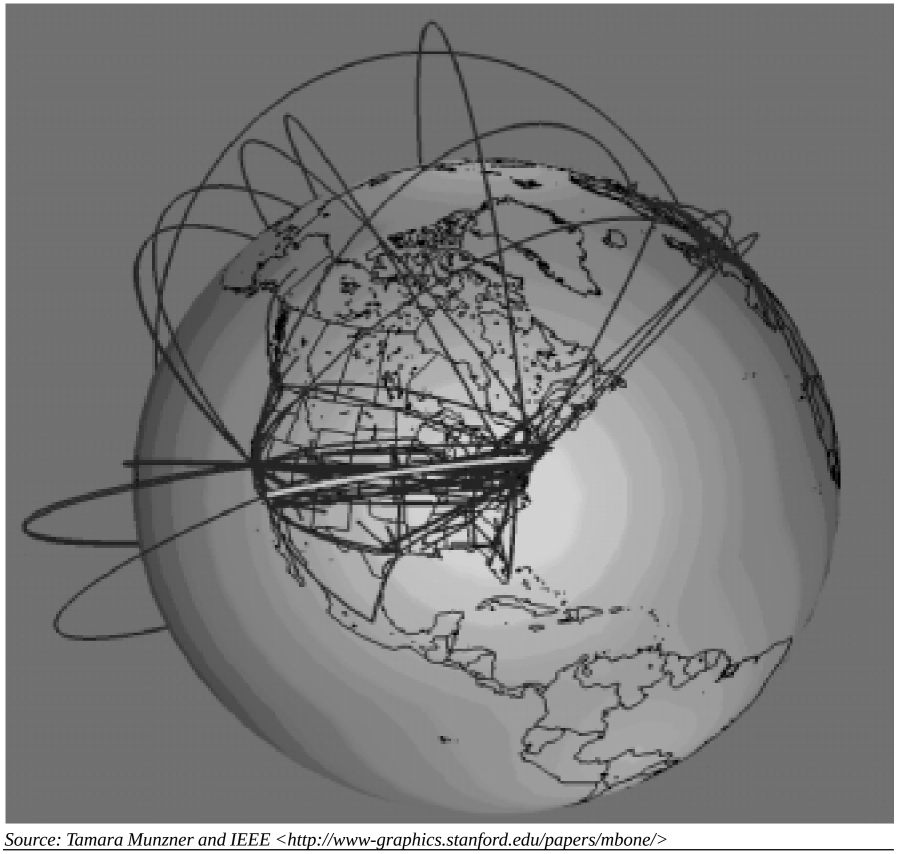

Source: Tamara Munzner and IEEE <http://www-graphics.stanford.edu/papers/mbone/>

Figure 19
3D Arcs on a Globe Representation of the Internet MBone Network

The final example is the Internet Mapping Project being undertaken by Hal Burch and Bill Cheswick at Lumeta Corporation (formerly at Bell Labs)

(Branigan *et al.* 2001) (www.cs.bell-labs.com/~ches/map/index.html). 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 (i.e., using a process of spatialization), thus disregarding the actual location of nodes in physical space. 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, United States. The resulting spatialization maps show how hundreds of networks connect to form the core of the Internet. Figure 20 shows the structure of the Internet from December 2000, representing nearly 100,000 network nodes. This highly complex spatialization 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 color-codings of the links according to different criteria (such as network ownership, country, etc.). In the example shown, links have been color-coded (depicted here in shades of gray) 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 for the use of other researchers. Over time, it is hoped that the data will be useful for monitoring growth and changes in the structure of the Internet. The experience gained in mapping the Internet is also being applied commercially, using network scanning and visualization techniques to chart the structure of corporate intranets for the purpose of identifying security weaknesses and unauthorized nodes.

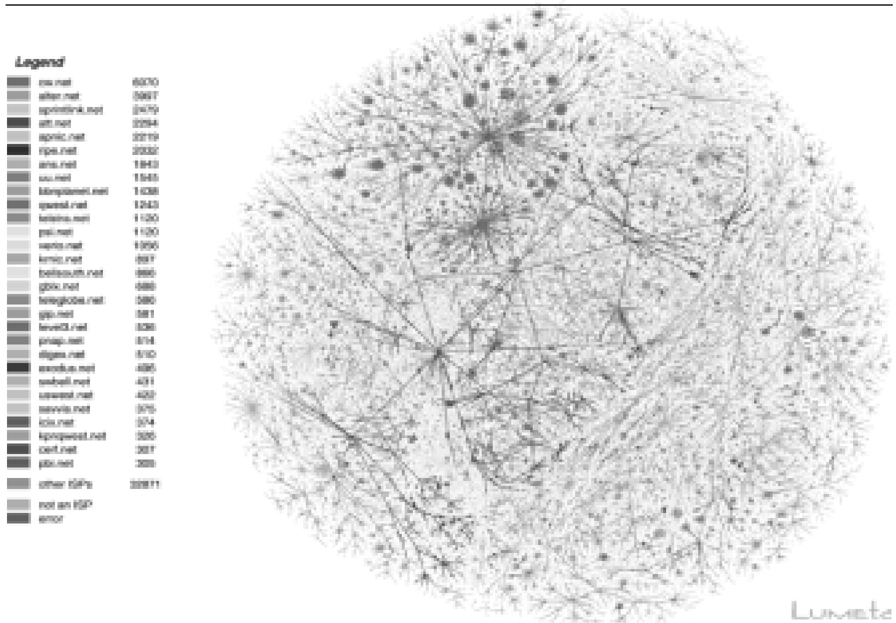

Figure 20
Map of the Internet Topology by Hal Burch and Bill Cheswick

Source: Lumeta Corporation <http://www.lumeta.com>

Conclusions

We have argued in this chapter that mapping can be used as a significant tool of analysis for managing Internet infrastructure, developing and implementing policy, and understanding the information economy. Maps can be used to reveal the range, extent, and density of Internet infrastructure in relation to real-world geography at a variety of scales.

We finish on a note of caution, however. While mapping is a useful strategy, with many of the maps visually striking and persuasive, there are four main reasons they need to be created, used, and interpreted with care. First, maps are only as accurate as the data used in their construction. While it is generally recognized that all spatial data are of limited accuracy because of the inherent error in data generation (e.g., surveying) or source materials, there are particularly acute problems in relation to data concerning the Internet. This is because what sources of data there are, are limited and fragmented, with few attempts to systematically measure the various components of Internet infrastructure. The problem is exacerbated by the Internet's fast growing and dynamic nature that makes keeping up with changes almost impossible. Consequently, maps are out-of-date before they are created as the component data they are constructed from has altered. In addition, the provision of both infrastructure and content services has become an intensely competitive and profitable business. As such, corporations are wary of giving away details that may aid competitors or threaten security, hence they police data relating to their own infrastructure (e.g., in relation to traffic flows). A further problem is there are no data standards for what data is produced. Hence, different agencies produce different kinds of data measured using varied techniques. This makes comparison of data from different sources difficult. Consequently, most maps, while fascinating, are often limited in scope, coverage, and currency because they are based on limited data.

Second, good maps require skilled construction. Maps necessarily depict a selective distortion of that which they seek to portray because they employ processes of generalization and classification. Weak cartographic technique—and poor judgment on how best to generalize and classify—can lead to poorly constructed maps that have low communicability. At present, many of the maps of Internet infrastructure are not being created by trained cartographers. This means that many have poor cartographic design standards, using inappropriate styles or poorly chosen categorization. Consequently, many maps are lacking in legibility and some may be misleading.

Third, due to a combination of the first two issues, many maps can propagate severe interpretation problems centered around issues of ecological fallacy. In regards to maps of infrastructure, ecological fallacy relates to the aggregation of data within spatial units—otherwise known as the Modifiable Areal Unit Problem (Openshaw 1984). The presentation of aggregated data can give the impression that all phenomenon within an area are similar, when in fact there could be significant variation. This can lead to inappropriate conclusions about that area. This is perhaps best revealed when the same data is mapped onto differing sets of spatial units (e.g., wards, districts, counties, states), as this can produce significantly different patterns across scales. Ecological fallacies are quite common (see Landweber example above), particularly when using secondary, "off-the-shelf" data such as that published by the World Bank, OECD, and International Telecommunications Union for example, because the data often relates to a particular scale (e.g., nations) but has no sub-scale variability. Consequently, there is little choice but to map it at the scale collected (see Dodge and Kitchin 2000b for a fuller discussion).

Lastly, all the maps we have discussed in this chapter have been created by people with a wide variety of motivations and agendas. As a consequence, all the maps are selective and subjective presentations of their underlying data, telling the "story" their creators have designed them to tell—even if created in a so-called scientific fashion, decisions have to be made over scale, symbols, layout, category classes, and what to map and what to omit. In many cases, this "story" will be benign; in others, it will be carefully constructed. For example, maps used for marketing purposes are essentially pieces of corporate propaganda designed to highlight the range and scope of services on offer, communicating to a potential customer that they offer the "right" network for the customer's needs. As such, it is necessary to think about who the map was made for, by whom, why it was produced, and what are the implications of its message and use.

Given the diversity of map purpose, the variety of mapping techniques adopted, the problems with data capture and availability, and the subjective decisions made in their creation, it should be noted that there is no one single map or technique that can capture all the complexities of the Internet's infrastructure, and no such map can be created. Instead, there are a multiplicity of different Internet maps that focus on different components of the infrastructure. Perhaps, even, our knowledge is diminishing as the scale and complexity of infrastructure grows and information about it becomes less open to scrutiny. That said, we believe based on our review of the projects in this chapter that mapping can provide a highly useful tool in understanding and managing Internet infrastructure.

References

Abbate, J. (1999) *Inventing the Internet*, Cambridge, Mass: MIT Press.

Abouna, A.G. (2001) “A Tour of the Bandwidth Bay Fiber Network Map,” *Proceedings 2001 ESRI International User Conference* (San Diego, July 9-13) http://www.esri.com/library/userconf/proc01/professional/abstracts/a963.html>

“The Archilles’ Heel of the Internet” (2002) *Nature* (July 27) <www.nature.com/nature/fow/00727.html> **(AUTHOR??) (PAGE NUMBER??)**

Ackerman, J. (1993) “Blazing a Well Worn Path: Cartographic Commercialism, Highway Promotion, and Automobile Tourism in the United States, 1880-1930, *Cartographica* 30(1):10-20.

Atkinson, R. (2002) *The 2002 State New Economy Index*, report, Washington, D.C.: Progressive Policy Institute (June) <www.neweconomy.org>

Atkinson, R. and R. Court (1998) *The New Economy Index*, report, Washington, D.C.: <www.newecomonyindex.org>

Atkinson, R. and P.D. Gottlieb (2001) *The Metropolitan New Economy Index*, report, Washington, D.C.: Progressive Policy Institute (April) <www.neweconomyindex.org>

Atkinson, R., R. Court, and J.M. Ward (1999) *The State New Economy Index*, report, Washington, D.C.: Progressive Policy Institute (April) <www.neweconomyindex.org>

Barrett, R. (1999) “Fiber Cuts Still Plague ISPs,” *Interactive Week* (June 4) <http://www.zdnet.com/zdnn/stories/news/0,4586,406174,00.html>

Becker, R.A., S.G. Eick, and A.R. Wilks **(**1995) “Visualizing Network Data,” *IEEE Transactions on Visualization and Computer Graphics* 1(1):16-28.

Branigan, S., H. Burch, B. Cheswick, and F. Wojcik (2001) “What Can You Do with Traceroute?” *Internet Computing* 5(5):96.

Castells, M. (2000) *The Rise of the Network Society*, second edition, Oxford: Blackwell.

Castells, M. (2001) *Internet Galaxy: Reflections on the Internet, Business and Society,* Oxford: Oxford University Press.

Cox, K.C., S.G. Eick, and T. He (1996) “3D Geographic Network Display,” *ACM Sigmod Record* 25(4):50-54.

Couclelis, H. (1998) “Worlds of Information: The Geographic Metaphor in the Visualization of Complex Information,” *Cartography and Geographic Information Systems* 25:209-220.

Delio, M. (2001) “From Code Red to Code Dread,” *Wired News* (August 1) <http://www.wired.com/news/lycos/0,1306,45760,00.html>

Dodge, M. (2000a) “Mapping Where the Data Flows,” *OnTheInternet* (April) <http://www.isoc.org/oti/dodge.html>

Dodge, M. (2000b) “Mapping the Global Spread of the Net,” *Mappa Mundi Magazine* (April 4) <http://mappa.mundi.net/maps/maps_011/>

Dodge, M. (2001) *Atlas of Cyberspace* Web site, Centre for Advanced Spatial Analysis, University College London. <http://www.cybergeography.org/atlas/>

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

Dodge, M. and R. Kitchin (2000b) “Exposing the ‘Second Text’ of Maps of the Net,” *Journal of Computer-Mediated Communication* 5(4) **PAGE NUMBERS??**.

Dodge, M. and R. Kitchin (2001) *Atlas of Cyberspace*, London: Addison-Wesley.

Ewalt, D.M. (2001) “Attacks Lead To Internet Slowdowns,” *InformationWeek.com* (September 12) <http://www.informationweek.com/story/IWK20010912S0013>

Fleming, D.K. (1984) “Cartographic Strategies for Airline Advertising,” *Geographical Review* 74:76-93.

French, S.P. and X. Jia (2001) “Georgia High-Speed Telecommunications Atlas: An ArcIMS Implementation,” *Proceedings 2001 ESRI International User Conference* (San Diego, July 9-13) <http://www.esri.com/library/userconf/proc01/professional/abstracts/a259.html>

Fry, C. (1999) “GIS in Telecommunications,” in P. Longley, M. Goodchild, D. Maguire, and D. Rhind, eds., *Geographical Information Systems Volume 2: Management,* New York: John Wiley and Sons.

Gorman, S.P. and E.J. Malecki (2000) “The Networks of the Internet: An Analysis of Provider Networks in the USA,” *Telecommunications Policy* 24(2):113-134.

Graham and Marvin 2001 **(COMPLETE CITATION??? See Page 13)**

Guernsey, L. (2001) “An Unimaginable Emergency Put Communications to the Test,” *The New York Times* (September 20) **PAGE NUMBERS??** <http://www.nytimes.com/2001/09/20/technology/circuits/20INFR.html>

Hafner K. and M. Lyons (1996) *Where Wizards Stay Up Late: The Origins of the Internet,* New York: Simon and Schuster.

Holderness, M. (1998) “Who Are the World’s Information-Poor?” in B.D. Loader, ed., *The Cyberspace Divide,* London: Routledge.

Kitchin, R. (1998) *Cyberspace*, Chichester: John Wiley.

Koutsofios, E.E. et al. (1999) “Visualizing Large-Scale Telecommunications Networks and Services,” *Proceedings of IEEE Visualization’99* (October 1999).

Manjoo, F. (2000) “Net Traffic at All-Time High,” *Wired News* (November 8) <http://www.wired.com/news/business/0,1367,40043,00.html>

Monmonier, M. (1991) *How to Lie with Maps*, Chicago: University of Chicago Press.

Moss, M.L. and A. Townsend (1997) “Tracking the Net: Using Domain Names to Measure the Growth of the Internet in U.S. Cities,” *Journal of Urban Technology* 4(3):47-59.

Moss, M. and A. Townsend (2000) “The Internet Backbone and the American Metropolis,” *The Information Society* 16:35-47.

Munzner T., E. Hoffman, K. Claffy, and B. Fenner (1996) “Visualizing the Global Topology of the Mbone,” *Proceedings of the 1996 IEEE Symposium on Information Visualization* (San Francisco, October 28-29) <http://graphics.stanford.edu/papers/mbone/>

Norris P. (2001) *Digital Divide: Civic Engagement, Information Poverty, and the Internet Worldwide*, Cambridge, U.K.: Cambridge University Press.

Nua (2003) *How Many Online?*
< http://www.nua.com/surveys/how_many_online/index.html>

OECD (2001) *Communications Outlook 2001*, Paris: Organization for Economic Co-operation and Development.

Openshaw, S. (1984) *The Modifiable Areal Unit Problem: Concepts and Techniques in Modern Geography 38*, Norwich: Geo Books.

Odlyzko, A.M. (2000) “Internet Growth: Myth and Reality, Use and Abuse,” *Information Impacts Magazine* (November 2000) <http://www.cisp.org/imp/november_2000/odlyzko/11_00odlyzko.htm>

Petrazzini, B. and M. Kibati (1999) “The Internet in Developing Countries,” *Communications of the ACM* 42(6):31-36.

Quarterman, J.S., S. Carl-Mitchell, and G. Phillips (1994) “Internet Interaction Pinged and Mapped,” *Proceedings of INET’94 Conference* (Prague, June) <http://www.isoc.org/inet94/papers/522.ps.gz>

Rickard, J. (1996) “Mapping the Internet with Traceroute,” *Boardwatch Magazine* (December) **PAGE NUMBERS??.**

Salus, P.H. (1995) *Casting the Net: From Arpanet to Internet and Beyond*, New York: Addison Wesley.

Sternberg, R. (2001) “Internet Domains and the Innovativeness of Cities/Regions: Evidence from Germany and Munich,” paper presented at *ESRC CITIES Competitiveness and Cohesion Programme Workshop* (Oxford, March 28-30) <http://cwis.livjm.ac.uk/cities/JSworkshop/Sternberg.pdf>

Tyner, J.A. (1982) “Persuasive Cartography,” *Journal of Geography* 81:140-144.

U.S. Department of Commerce, National Telecommunications and Information Administration (2000) *Falling Through the Net: Toward Digital Inclusion*, Washington, D.C.: Government Printing Office <www.ntia.doc.gov/ntiahome/digital divide/>

U.S. Federal Communications Commission (2000) *Availability of High-Speed and Advanced Telecommunications Services*, order number 00-290, Washington, D.C.: Government Printing Office <http://www.fcc.gov/Bureaus/Common_Carrier/Orders/2000/fcc00290.pdf>

Warf, B. (2001) "Segueways into Cyberspace: Multiple Geographies of the Digital Divide," *Environment and Planning B: Planning and Design* 28(1):3-19.

Wei, B. et al. (2000) "Visualization Research with Large Displays," *IEEE Computer Graphics and Applications* (July/August) **(PAGE NUMBERS??)**.

Wheeler, D. and M. O'Kelly (1999) "Network Topology and City Accessibility of the Commercial Internet," *Professional Geographer* 51:327-339.

Zook, M.A. (2000) "The Web of Production: The Economic Geography of Commercial Internet Content Production in the United States," *Environment and Planning A* 32:411-426.

Zook, M.A. (2001) "Old Hierarchies or New Networks of Centrality? The Global Geography of the Internet Content Market," *American Behavioral Scientist* 44(10):1679-1696.

Zook 1998 p21 **(WHAT IS THIS CITATION??--SEE PAGE 22)**