

Mapping Networks and Services

John S. Quarterman¹
 Smoot Carl-Mitchell²
 Gretchen Phillips³

Abstract

Computer networking is widespread, supported by the global Matrix of interconnected computer networks such as the Internet, BITNET, UUCP, and FidoNet. Perhaps 20,000,000 people use this global Matrix of computer networks, and several major global computer networks are growing exponentially, rapidly reaching new countries and new classes of users. This growth prompts network users and service providers to ask questions like:

- *Where are these users?*
- *What is the magnitude of the various network communities?*
- *What are the available services?*
- *How are services distributed?*
- *And, finally, how can this data be represented in a meaningful graphical way?*

We have chosen Latin America and the Caribbean (ALyC) as the geographic region for discussing these issues and demonstrating the mapping of currently available data. These are not traditional shaded geographic outlines nor node and link connectivity maps, but rather, maps that show locations and numbers of networked hosts and services. This representation of the network gives a presentation of where the networks go, rather than how they get there, and shows what areas are reachable by each network type. We do not show exactly where the users are, but we show where the hosts they use are. These maps can give network administrators and users alike a better understanding of the areas connected to the global Matrix, and the services available in each area.

1 John Quarterman is Senior Technical Partner in Texas Internet Consulting. He can be reached at jsq@tic.com

2 Smoot Carl-Mitchell is Managing Partner in Texas Internet Consulting. He can be reached at smoot@tic.com

3 Gretchen Phillips is the Manager of the UNIX Support Group at the University at Buffalo. She can be reached at phillips@acsu.buffalo.edu

I . Introduction

Four global networks carry most of the international network traffic; FidoNet, UUCP, BITNET, and the Internet. The first step in looking at these networks involves collecting appropriate data for each and performing some consistency checking. The data for these maps was initially derived from their node lists (BITNET and FidoNet) or host lists (UUCP). The Internet data came from a domain walk by Mark Lottor.

The data available from the node lists, public maps and the domain walk is not in a format that is readily usable for this type of mapping. Each map or other source of network information tends to include different details. We have correlated the data by using telephone country codes, area codes and local exchanges for determining locations. A more complete discussion of the methods is available in issues of *Matrix News*[5-7]

II . Classification

Classification of location by telephone country codes has been straightforward for most ALyC countries; for example, Brazil is telephone country code 55 and Mexico is 52; Costa Rica is 506 and Uruguay is 598. When the International Telecommunications Union (ITU) assigned these numbers, it made them all start with 5, using two digits for each populous country and three for smaller countries. There have been no recent political border changes to account for, unlike in Europe. Telephone country code 53, for Cuba [3], tends to be missing from many lists.

The Caribbean is interesting, since almost two dozen countries use telephone exchanges in area code 1-809, under the country code 1, that also covers Canada, the United States, and Guam. These countries may be hard to see in the maps as reproduced in black and white for the proceedings (the maps look better in color). Here is a list of what to look for: Puerto Rico (Internet, UUCP, FidoNet, BITNET), U.S. Virgin Islands (FidoNet, Internet, UUCP), Jamaica (UUCP), Dominican Republic (UUCP), Guadeloupe (UUCP and EARN), Grenada (UUCP), and Trinidad and

Tobago (UUCP).

IIA . Internet Domain Walk

As mentioned, the Internet data mostly came from Mark Lottor's SRI NISC domain walk [2, 5]. The Internet uses the TCP/IP (and other) protocols to provide fast, sophisticated services, including not only electronic mail, mailing lists, and USENET news (group discussions), but also remote login (TELNET, rlogin, the X Window System, and others) and file transfer (FTP, rcp, NFS, AFS, and others). We will discuss some of these and some other Internet services later in this paper. The Internet is the hardest of these four networks to localize, because of the distribution of host naming authority with the Domain Name System (DNS). Unlike the other networks, there is no central node list for the Internet; instead, each organization controls its own name space, and may delegate part of it to suborganizations. Each country has a top level domain, such as BR for Brazil and MX for Mexico; CR for Costa Rica and UY for Uruguay. Second level domain registries, for domains like itesm.mx, are kept per country. This requires inquiries to approximately 30 country registries for complete localization of ALyC alone, not to mention the rest of the world. In addition, many organizations, particularly large universities, have third level subdomains in various locations, such as mty.itesm.mx, and these will usually not even be in the country registries. This distribution of DNS authority is good for the Internet, but not so good for mapping. This problem is evident in the maps shown here. The Internet in Brazil only appears in São Paulo and Rio de Janeiro. We don't know if there are Internet nodes elsewhere in Brazil.

Mexico is an interesting case of how much mapping can be done with a little information. The Instituto Tecnológico y de Estudios Superiores de Monterrey (ITESM) has numerous campuses across the country. Fortunately, they keep a list available by anonymous FTP of their own main subdomains, plus those of many other organizations. That list does not contain place names for all the domains it lists, but many of the domain abbreviations are clear, and some can be found in the BITEARN node list or the UUCP map. We have been a bit speculative in places. Is mor.itesm.mx Morelia, Michoacán or Cuernavaca, Morelos? We have guessed the latter. We would think gro.itesm.mx is for Quintana Roo, so the actual hosts are probably in Chetumal, but we

don't know for sure. Even with these uncertainties, a simple list of domains with some geographical annotations has permitted us to place 92 percent of the Mexican hosts from Lottor's domain walk with reasonable certainty, and all but about 1 percent of the rest with plausibility.

Of course, since domain walks show what hosts have IP (Internet Protocol) addresses advertised by DNS, not which hosts are actually reachable, we also do not know which hosts in Lottor's domain walk are really on the Internet, nor do we know how many hosts on the Internet are not in the domain walk. We do not know the corresponding information for the other networks either, because being listed in the UUCP map or the BITEARN or FidoNet nodelist is no guarantee that a host is reachable. The only way to be sure is to actually try to reach each host on each network with an appropriate protocol. We considered that beyond the scope of our paper,

IIB . FidoNet Hosts

Since the Fidonet nodelist is the definition of FidoNet, the primarily MS-DOS dial-up network that uses the Fido protocols to deliver mail and echomail (group discussion topics), we used only that list in determining the location and magnitudes of the FidoNet hosts [7].

IIC . UUCP Hosts

UUCP, the mostly UNIX mostly dialup network, uses the UUCP (UNIX to UNIX CoPy) protocol to supply electronic mail (and often also USENET news) to even the most remote parts of the world. Most of the UUCP data came from the UUCP maps, that are posted to USENET newsgroup comp.mail.maps, and are available by anonymous FTP and UUCP from many hosts. The UUCP data was augmented from several sources.

The folks at ORSTOM in Montpellier, France were kind to send the node list for the RIO network (you can get it by anonymous FTP from orstom.orstom.fr), which uses mostly UUCP and sometimes X.25 and TCP/IP to connect interesting parts of the world. Most RIO nodes are in Africa, the Indian Ocean, or the South Pacific. However, RIO also connects Guadeloupe in the Caribbean, and French Guiana in South America. These are both politically departments of France, and their main language is French, not Spanish or Portuguese. Nonetheless, they are physically located in the southern part of the western hemisphere (RIO lists them as being in the

Antilles), so we map them.

IID . BITNET Hosts

For BITNET we used only the BITEARN nodelist. We are using BITNET as the name of the worldwide collection of interconnected NJE (Network Job Entry) networks. In ALyC, there are actually many such networks, including SCARNET in Argentina, Chile, Peru, and Uruguay, ANSP in Brazil, RUNCOL in Colombia, and ECUANET in Ecuador [6]. Guadeloupe has a connection to EARN, the European Academic and Research Network. Puerto Rico is connected to BITNET, the NJE network run by the Corporation for Research and Education Networking (CREN) in the United States, as are the U.S. states visible on the maps. We map all the hosts within the borders of our maps that these networks have listed in the BITEARN nodelist.

III . Correction of Data

Larry Landweber also researches the extents of networks [1], producing a list every few months indicating 0, 1-5, or 6 or more hosts on each network in each country. Since we use different techniques for discovering the data, hosts are counted here for some countries that he does not list. Specifically, these are: BITNET hosts in Peru; Internet hosts in Costa Rica and the U.S. Virgin Islands; and UUCP hosts in Honduras and the U.S. Virgin Islands. From his list, we have added UUCP data for Bolivia, Cuba, Ecuador, Grenada, Panama, Suriname, Trinidad and Tobago, and Venezuela. We have used his list to adjust upwards the UUCP counts for Brazil, Mexico, Puerto Rico, and Uruguay. Apparently some countries are not listing their nodes in the UUCP map, or at least are not listing them all. We have also received information from contacts in ALyC.

IIIA . Mapping Difficulties

As we began working on the display of this data, we tackled one network at a time, displaying a single network on a map. The first problem was to determine the longitude and latitude of the hosts and then apply from the data the correct number of hosts at each location. Once each network was handled individually, we began to build a composite map of all the networks and their magnitudes on a single map. Display of the data clearly called for an exponential scale, particularly when we began displaying more than one network on a map. Brazil and Mexico each have more than

1,000 Internet hosts, and some sites have only one host. Both large and small sites need to be visible across three or more orders of magnitude.

The various networks must be visible separately on the composite map, so we used both color coding and different shapes to make the networks visually distinct. As we began building the composite map, it was confirmed that networks like UUCP and FidoNet provide connectivity to remote locations, while the Internet brings services into a country primarily at more populous locations. The locations of Internet hosts indicate where services such as TELNET and FTP are available. UUCP and FidoNet provide mostly mail and news, and some file transfer, as does BITNET. But these networks do not provide services such as TELNET, archie, gopher, and WAIS that the Internet provides. However, an Internet connection to the capital or largest city may serve to funnel mail in and out of the country for many non-Internet hosts in the rest of the country, as is the case in Argentina [4].

IV . Internet Services

The Internet provides services that are interesting to consider separately, including anonymous FTP, archie, WAIS (Wide Area Information Server), and gopher. An anonymous FTP permits any of the approximately 8 million Internet users in more than 40 countries to connect from any of more than 1 million Internet hosts and retrieve files without a personal login name. This service is commonly used for publishing documents, software, and other information on the Internet. There are probably more than 2,000 anonymous FTP servers on the Internet. As anonymous FTP servers became numerous and diverse, a tool for discovering materials was necessary. Archie, the query tool for indexes of anonymous FTP, helps the community to help itself, by polling more than 1,000 anonymous FTP servers, and providing searches of their listings. Distribution of archie servers, as with the distribution of anonymous FTP sites, is driven both by need and interest. WAIS servers (and clients) were developed as a tool for indexing, searching, and retrieving a variety of documents and other information currently available in electronic form. There are more than 300 WAIS databases worldwide. Gopher has been called "duct tape for the Internet," because it glues Internet services together, providing an easy menu-oriented user interface to files, TELNET

services, WAIS databases, and other services. There are more than 300 gopher servers in the main gopher list.

These services are all accessible from anywhere on the Internet, including ALyC hosts, assuming the local host has appropriate client software. Some of these services are starting to be provided on servers located in Latin America. This means some ALyC countries or organizations are starting to become producers, as well as consumers, of these sophisticated Internet services. We have indication of two anonymous FTP servers and a WAIS server in Monterrey, Mexico. Gopher servers are available from sites in Santiago, Chile (3); Quito, Ecuador (2); and São Paulo, Brazil (1). Additionally, there are undoubtedly other gophers that are not in the main list, including one at ITESM in Monterrey. (Choosing a set of gopher servers to map is an interesting exercise in itself, since there are several main lists, and you can find other servers by looking in the servers you already know.) We do not have any record ofarchie servers in Latin America, however, as the Internet expands throughout Latin America, we do expect to see increases in the number of sites providing these services.

V . Maps of the Matrix in ALyC

The point of this paper is the two attached maps (pages 7 and 8). One is of South America, in Lambert Azimuthal projection; the other an Albers Conic Equal-Area projection of northern ALyC, from Mexico through the Gulf of Mexico to the Caribbean, with some of South America repeated. Both maps use the same scale, of 1:40,000,000, and together cover all of Latin America and the Caribbean. The projections were done with the GMT software from the University of Hawaii. The software for computing the locations of network hosts and combining networks, with map legends, was written by the authors and all the data is from late January or early February 1993.

Latin America and the Caribbean show almost every networking approach used anywhere in the world, from the UUCP quick and inexpensive path of Peru to the initial Internet connectivity of Venezuela to the banking backing of IP connectivity in Ecuador to the big city Internet connectivity of Brazil and Argentina feeding mail to the UUCP and FidoNet connectivity of the smaller cities and towns. This last pattern is similar to that of India.

VA . South America

Peru, on the west coast, has a bit of BITNET connectivity and a lot of UUCP nodes, but no FidoNet and no Internet connection. UUCP was inexpensive and easy to install to connect universities and research institutions. External mail transport has been provided by a dialup link to Oregon, in the United States. Mail to Peruvian hosts uses DNS domain names and looks like Internet mail, though no Peruvian hosts are actually on the Internet. Peru is now seeking an Internet connection.

No Peruvian hosts are registered in the UUCP maps, either. We obtained the UUCP count from Peruvian networkers. We know about 15% of the machines indicated are not in Lima, but we don't know exactly where they are, so we lumped them in with the Lima hosts.

This illustrates an important aspect of these maps: hosts are not omitted if we don't know where they are. We just aggregate them with the nearest (sometimes in subtle senses) neighbor hosts in their country until we can separate them out.

Venezuela, on the north coast, has much more Internet connectivity than anything else. They started relatively late, and they have an international Internet connection that is partly supported by the U.S. National Science Foundation (NSF), as does Ecuador and perhaps some other countries.

Ecuador is perhaps unique in the world in having one of the largest banks in the country backing installing IP connectivity throughout the country. Expect rapid growth in Ecuadorian networking.

Chile has long been connected to many networks, partly because of the presence of astronomical observatories, which made it an attractive place for the U.S. National Aeronautics and Space Administration (NASA) to encourage connectivity. Whatever the early history, Chile is increasingly well-connected, and is at the moment the Latin American country advertising the most gopher servers.

Argentina was one of the earlier places in Latin America to start networking, and shows a pattern that is visible other places, such as India. There is Internet connectivity to the capital, Buenos Aires, on the Atlantic, partly from some funding from the United Nations Development Programme (UNDP). Mail flows from Buenos Aires to the rest of the country over the UUCP

network (actually organized as several networks). Most of these UUCP nodes are actually supported by UUPC, an MS-DOS implementation, running on Commodores. Internet connectivity will spread from Buenos Aires to the rest of the country. For more on Argentina, see *Matrix News*, November 1991. [4].

Brazil shows Internet hosts in São Paulo (the large green circle) and Rio de Janeiro (the smaller circle). Some other cities are probably also on the Internet, but we need access to the Brazilian national domain registry to determine that.

FidoNet is for some reason more popular in Brazil than Argentina, and UUCP the reverse. The reasons are probably historical; probably one network started earlier in each country.

VB . Mexico to Caribbean

Much network activity in Mexico starts in Monterrey, in the far northeast. The Internet has already spread to Mexico City, Puebla, and many other cities throughout the country.

Costa Rica has long had at least one UUCP node, and has recently joined the Internet. Panama has FidoNet and BITNET connectivity. Nicaragua and Honduras have minimal connectivity. Guatemala has no direct connectivity, although many Guatemalans log in by X.25 onto a host in Costa Rica.

Havana, Cuba, is connected somewhat circuitously to the global mail network by UUCP links, and the DNS domain CU is registered and working. There is apparently also a FidoNet node at the U.S. naval base in Guantanamo Bay.

Jaimaca, Grenada, Trinidad, the U.S. Virgin Islands, and especially Puerto Rico are connected to various networks. Santo Domingo in the Dominican Republic only shows one UUCP node, but it is managed to provide access to hundreds of researchers across the country.

VC . Corrections or Further Information

The additions to the networks are surprising each time we obtain new data. The Internet is beginning to have wide exposure in ALyC and is reaching new cities each month. Internet services like gopher and WAIS will continue to grow as the Internet continues to expand. FidoNet and UUCP are reaching even the remotest areas of the world, providing basic electronic communication to an increasing number of people. Our graphical

representations of the networks and services are intended to help the electronically networked community have a better understanding of where the network goes, who is served and where services are available. If you notice hosts missing from the maps, please contact us with corrections or further information.

VI . Acknowledgements

We are most grateful to those individuals who have helped in the collection of the data. In particular, we would like to thank Mark Lottor for the Internet domain walk, ITESM for the Mexican Internet node list, Randy Bush for access to the FidoNet nodelist, EARN for access to the BITEARN nodelist, Monique Michaux for the RIO nodelist, Jose Soriano for Peruvian information, and Cristian Matias for 1-809 information.

VII . References

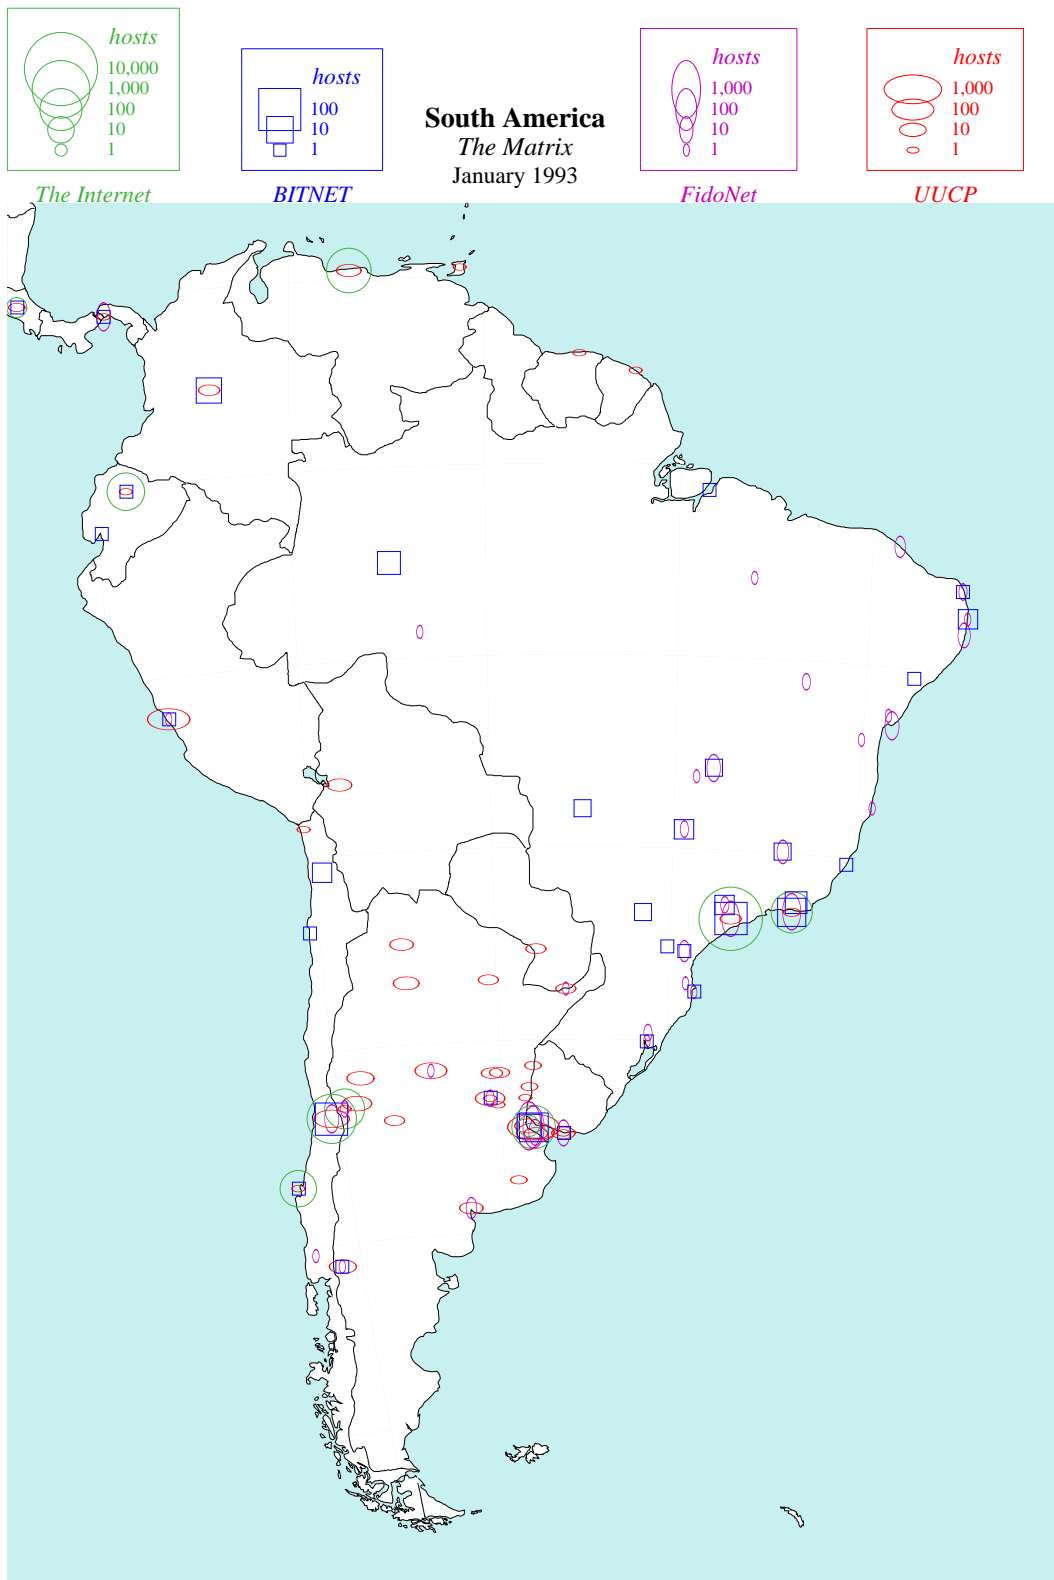
1. Landweber, Larry, "International Connectivity," *Internet Society News* 1(4)(Winter 1993).
2. Lottor, Mark, "Internet Growth (1981-1991); RFC1296," *Network Working Group Request for Comments*, Network Information Systems Center, SRI International, (January 1992).
3. Press, Larry and Snyder, Joel, "A Look at Cuban Networks," *Matrix News* 2(6) pp. 1-2 Matrix Information and Directory Services, Inc. (MIDS), (June 1992). lpress@isi.edu.
4. Quarterman, John S., "Networks in Argentina," *Matrix News* 1(8) pp. 1-6 MIDS, (November 1991). jsq@tic.com.
5. Quarterman, John S., "Where is the Internet?," *Matrix News* 2(8) pp. 9-15 MIDS, (August 1992). jsq@tic.com.
6. Quarterman, John S., "Where is BITNET?," *Matrix News* 2(10) pp. 3-7 MIDS, (October 1992). jsq@tic.com.
7. Quarterman, John S., "Where is FidoNet?," *Matrix News* 2(6) pp. 7-9 MIDS, (June 1992). jsq@tic.com.

Author Information

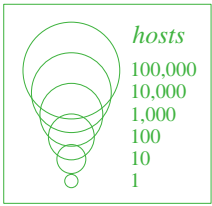
Smoot Carl-Mitchell and John S. Quarterman are partners in Texas Internet Consulting (TIC), which consults in networks and open system with particular emphasis on TCP/IP networks, UNIX systems and standards. They also

write articles and books. The books include *Practical Internetworking with TCP/IP and UNIX*, 1993, and *How to Connect to the Internet*, forthcoming in 1993, both from Addison-Wesley. Quarterman wrote *The Matrix: Computer Networks and Conferencing Systems Worldwide* (Digital Press, 1990), and co-authored *The Design and Implementation of the 4.3BSD UNIX Operating System* (Addison-Wesley, 1989), and *UNIX, POSIX, and Open Systems: The Open Standards Puzzle* (Addison-Wesley, 1993). Carl-Mitchell and Quarterman, as Matrix Information and Directory Services, Inc., publish *Matrix News*, a monthly newsletter, on paper and online, about contextual issues crossing network, geographic, and political boundaries.

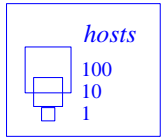
Gretchen Phillips is currently the manager of the UNIX Support Group at the University at Buffalo. As such, she is responsible for directing the group in all areas of UNIX support, including system configuration and installation, networking, and daily operations, as well as documentation and end user support. This group provides support for several hundred UNIX workstations, primarily Sun and SGI, networked across two campuses and centralized UNIX timesharing for over 3000 faculty, researchers and students. She has been involved in UNIX system administration since 1983.



+1-512-451-7602 Copyright © 1993 **Matrix, Inc.** Austin, Texas, USA mids@tic.com
1:40,000,000 Lambert Azimuthal projection 93.03.29_18:28

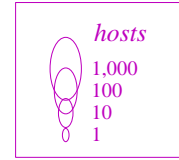


The Internet

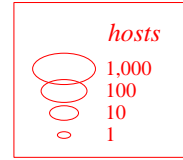


BITNET

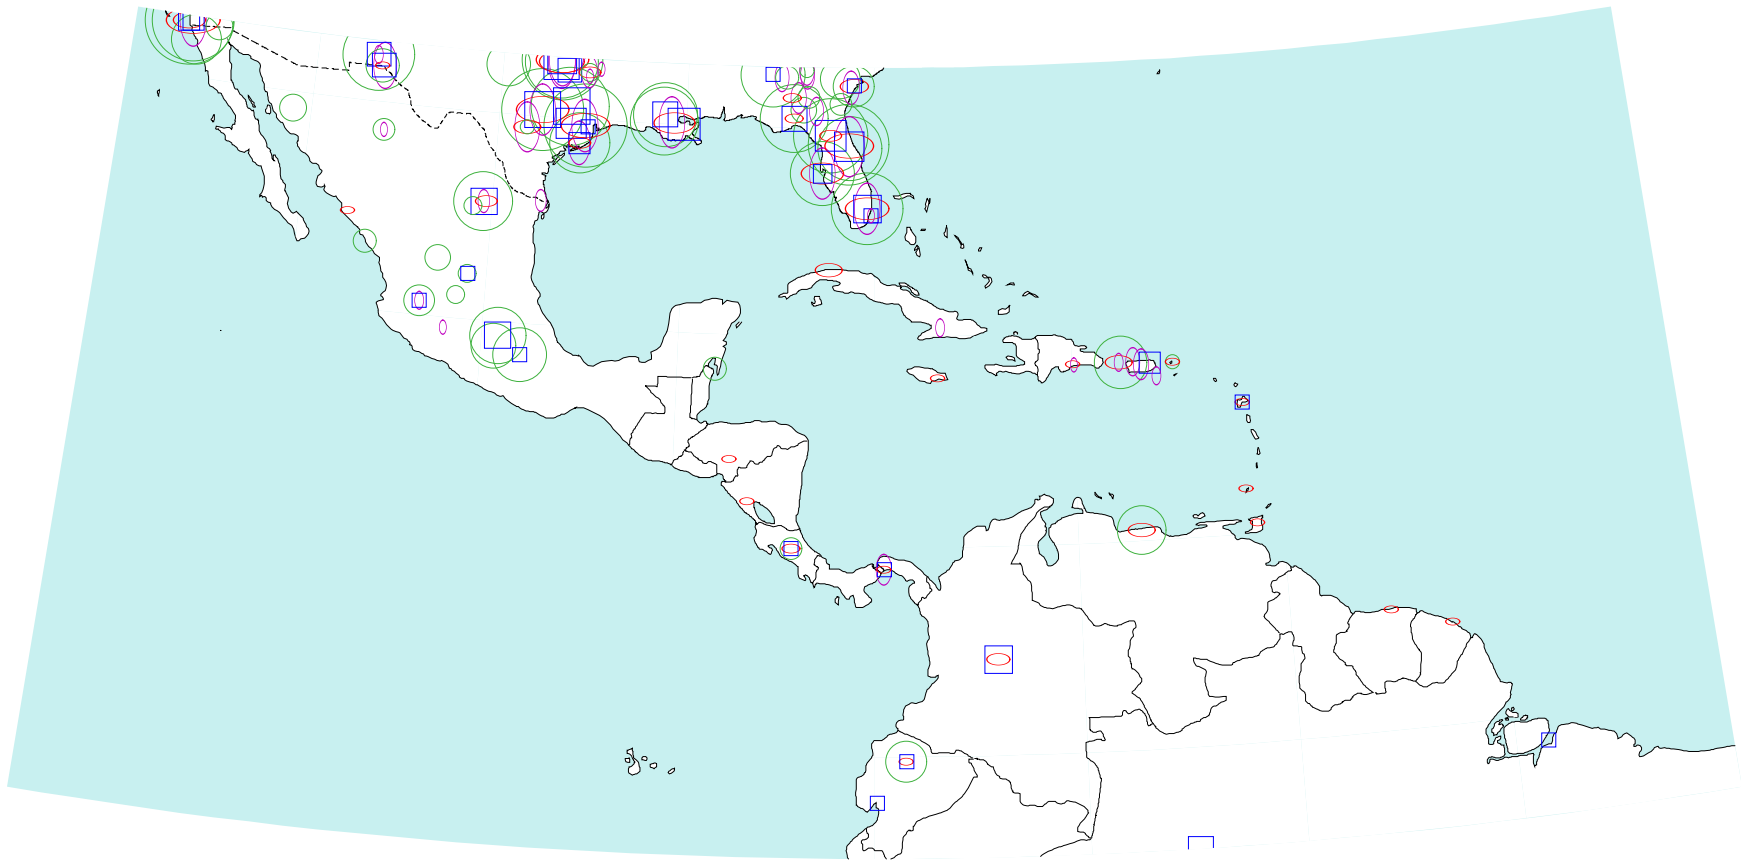
Mexico to the Caribbean
The Matrix
 January 1993



FidoNet



UUCP



+1-512-451-7602 Copyright © 1993 **Matrix, Inc.** Austin, Texas, USA mids@tic.com
 1:40,000,000 Albers Conic Equal-Area projection 93.03.29_18:47