

The Matrix

*Computer Networks
and Conferencing Systems
Worldwide*

*John S. Quarterman
Texas Internet Consulting*

digital™

Digital Press

TK
5105.5
.Q37
1990

Copyright © 1990 by Digital Equipment Corporation

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the publisher.

9876543

Order number EY-C176E-DP.

Printed in the United States of America.

A list of trademarks and registered trademarks mentioned in this book appear as Credits and Trademarks at the end of this book.

The following figures are printed with permission from the authors referenced in the appropriate figure legends:

Figures: 7.1; 7.2; 10.1; 10.2; 10.3; 11.1; 11.2; 11.3; 11.4; 11.5; 11.6; 11.7; 11.8; 11.9; 11.10; 11.11; 11.12; 11.13; 11.14; 11.15; 11.16; 11.17; 11.18; 11.19; 11.20; 11.21; 12.1; 12.2; 12.3; 12.4; 12.5; 12.6; 12.7; 12.8; 12.9; 13.1; 13.2; 13.4; 14.1; 15.1; 15.2; 18.1

Copyediting: Barbara Jatkola
Index: Gordon Brumm
Typesetting: Jaap Akkerhuis
Production: Editorial Inc.

Library of Congress Cataloging-in-Publication Data

Quarterman, John S., 1954 -

The matrix: computer networks and conferencing systems worldwide
John S. Quarterman.

p. cm.

Includes index.

ISBN 1-55558-033-5

1. Computer networks. 2. Telecommunication. I. Title.

TK5105.5.Q37 1990

621.382-dc20

89-1392
CIP

A good example of a **NETSERV** is **KERMSRV** on **CUVMB.BITNET**. This is the primary distribution server for **Kermit** software and documentation. It replies to hundreds of requests a day, originating from all over the world. It is sufficiently sophisticated to be able to return a response in the format of a request—e.g., as an **NJE** interactive message, as electronic mail, or as a virtual card deck [da Cruz 1988].

Access

See the previous information about **NETSERV** and the sections on the constituent parts of the network later in this book: **BITNET** in the United States, **Net-North** in Canada, **EARN** in Europe, **ILAN** in Israel, and **GulfNet** in the Persian Gulf states.

10.2.2 USENET

USENET began in 1979 and is thus one of the oldest cooperative networks [Daniel et al. 1980; Emerson 1983; Tokuda et al. 1986]. It is also one of the largest networks by almost any measure, having about 265,000 users and 9,700 hosts on five continents, as shown in Figure 10.1, and ranging from individuals to people at the largest corporations and universities. In fact, one of its key distinguishing features is that essentially anyone and any organization can join. Its name is usually taken to mean "User's Network." This network is unusual in a number of ways, most remarkably because it supports only one basic service, *news*, a distributed conferencing service, and does not support electronic mail (the closely associated network **UUCP** supports mail but not news).

Administration

USENET is one of the most decentralized networks: all a new machine needs to join in most regions is another machine to communicate with [Todino and O'Reilly 1988a]. Although registration in the general maps is advisable, there is no central authority that determines access. There is also no central funding source or mechanism: each host pays for its own transmission costs and generally passes through traffic originating at other hosts. This has led to a general consensus that commercial use of the network is frowned upon. Product announcements are usually permitted, but advertisements are not. There is no way to enforce such policies directly. Instead it is done by public opinion and direct complaints: anyone posting an advertisement will get direct mail complaining about it and will see postings on the network showing that many people think the original posting was a bad idea. Since people and companies don't generally like bad publicity and castigation by their colleagues and customers, this usually works. In extreme cases, people have proposed putting pressure on the




Figure 10.1. LUSENET worldwide map (June 1989) [Raid 1989a]

administrators of machines neighboring the machine from which the offending message was posted in an attempt to remove the posting machine from the network. Apparently this has never actually been done.

The greatest degree of organization in the network is found in the *backbone*, shown in Figure 10.2 and Figure 10.3. This consists of the machines with the greatest amount of traffic and thus the highest costs. (The European backbone is shown in the *EUnet* section in Chapter 13.) Since these hosts bear much of the burden of the network, their administrators tend to take a strong interest in the state of the network. The exact influence of the backbone is unclear, but it is involved in the organization of both the distribution and the content of news. There is a backbone on each continent, with trees radiating from its nodes. There may be redundant links among the nodes of those trees as well.

Speed

Most backbone (and many other) links currently run at about 11000bps. The average speed of slower links is probably close to 2400bps. More than 90 percent of all messages reach 90 percent of all hosts for which they are intended within one day.

Articles and Newsgroups

The element of information is called an *article*. Its format, as given in RFC1036 [Horton and Adams 1987] is modeled after that of ARPANET mail, as given in RFC822 [Crocker 1982]. That is, articles are lines of 7 bit ASCII text with a header prepended. The *From:*, *Date:*, and *Subject:* header lines are used in common with RFC822, although RFC1036 makes more requirements on *From:* and *Date:* than does RFC822. There are additional header lines peculiar to news articles, some of which are mentioned later in this section.

Articles are grouped according to topics in *newsgroups*, of which there are about 350. Newsgroup names are intended to reflect the topics to be discussed in them.

<i>news.announce.important</i>	Used for important announcements that are intended for everyone on the network
<i>news.announce.newusers</i>	Used for new users of the network

These are both *moderated* newsgroups: that is, when a user posts an article, the news software arranges to mail it over a mail network (usually *UUCP*) to a *moderator*, who decides whether or not to post it on the basis of considerations such as relevance, interest, clarity, and nonredundancy.

Some other newsgroups play roles in the administration of the network.




Figure 10.2. USENET backbone topology (3 April 1988) [Spafford 1988a]

news.groups

comp.mail.maps

Newsgroups are created or deleted according to decisions made after discussion in this newsgroup. This is used to distribute maps of the UUCP network, with USENET hosts marked as annotations. These moderators form another group interested in USENET organization. Most newsgroups are not moderated, and postings to them are distributed directly.

The current top level newsgroups in USENET and their subjects are as follows [Spafford 1988b]:

comp	Computer science, software source, hardware and software systems
sci	Technical discussions about established sciences
misc	Topics that don't fit under another top level group
soc	Social issues and socializing
talk	Debates, lengthy discussions, few resolutions
news	The news network and the software that implements it
rec	Hobbies and recreational activities

Distributions

In Japan and Europe, the services of USENET are administered jointly with those of UUCP in the networks JUNET and EUnet, respectively, and there are large amounts of data that do not normally travel outside of Europe or Japan. Yet much news does travel the world, and much of the structure of

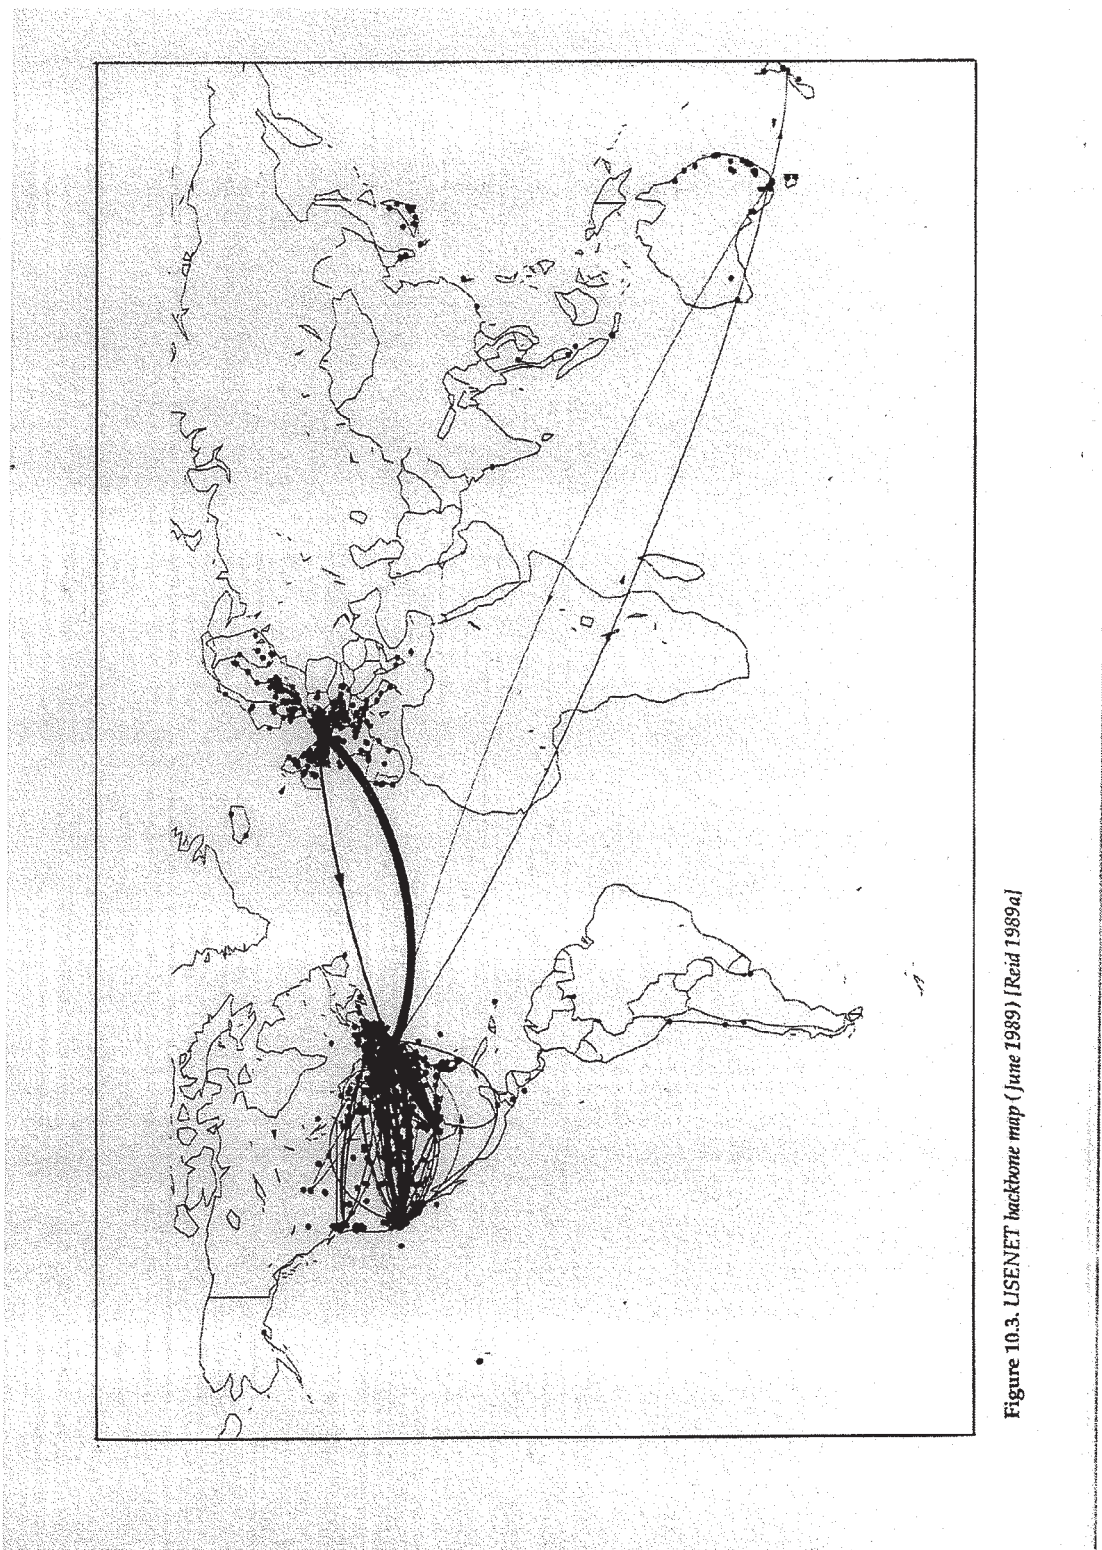


Figure 10.3. USNET backbone map (June 1989) [Reid 1989a]

Table 10.4. USENET distributions

Distribution	Geography
world	Everywhere
eunet	EUnet (Europe)
att	AT&T
can	Canada
na	North America
usa	United States
ca	California
ba	San Francisco Bay Area
ga	Georgia
ne	New England
tx	Texas
austin	Austin, Texas

USENET is reflected throughout the world. No explicit gatewaying is required to post an article that will travel to various parts of the world (although most users know that the main gateways to Europe, Japan, and Australia are *mcvax*, *kddlal*, and *munari*, respectively). Rather, *distributions* are used. Some current distributions are shown in Table 10.4. Many states, provinces, and cities have their own distributions. Many companies maintain internal distributions, and often there is a distribution just for the machine of posting; these most local distributions are not really part of USENET proper.

Distributions are specified by the header line *Distribution:* in the outgoing article. For example, in

Distribution: world

the *world* distribution is used (this is also the default when no distribution is specified), and the article will go everywhere that USENET does. Accidental use of this distribution by naive users can cause problems, as in the canonical case of car for sale advertisements posted in New Jersey that aren't of any interest to people in Australia or Europe, particularly considering that people on those other continents have to pay for articles received from North America. Thus, most news posting software asks the user posting to such a distribution whether broadcast to the whole world is really intended. Also, not all newsgroups are accepted everywhere.

Protocols and Management

Articles are transferred by a flooding algorithm: each host sends each article to each other host with which it communicates. There are two mechanisms for preventing loops, each involving a header line in each article as well as an expiration mechanism:

Path: longway!ames!mailrus!mcgill-vision!gate—This *Path:* header is updated at each host to reflect the host from which the article was just received, and thus reflects the entire path the article has traveled from the originating host to the receiving host. The receiving host will not send the article back out to any host that appears in this header.

Message-ID: <1064@mcgill-vision.UUCP>—This header line is set when the article is posted and is never changed; it is required not only for news articles, but also by RFC822. It is supposed to be unique throughout the network. This is usually accomplished by having it include the host name and a sequence number for articles posted from that host, or the time of day, as in the above example. This header line is used to construct a history file on each host of all articles received on that host. If an article with the same *Message-ID:* as one in the history file is seen, it is considered to be a duplicate and is thrown away. This is not quite as good a loop-prevention mechanism as the *Path:* header because someone has paid to transfer a duplicate article, but it protects against corrupted *Path:* lines and broken software on other hosts.

Expires: 24 Apr 88 23:44:56 GMT—Since hosts do not have unlimited storage, news articles are usually *expired* after some rather short period, by default 14 days. Expired articles are also removed from the history file. It is possible to set the expiration date of an article with the *Expires:* header line, but this should be done only for important articles.

There are also some fine-tuning mechanisms:

- | | |
|----------|---|
| Outgoing | Each host has lists of the hosts it communicates with and of the newsgroups and distributions to send to each of those hosts. If an outgoing article's newsgroup and distribution are not included in the lists for a given destination host, the article is not sent to that host. |
| Incoming | Each host also has a list of the newsgroups and distributions it will accept. Thus, a host may decide not to accept an article it has received. Once again, someone has to pay for transferring the article anyway, so this situation is to be avoided by tuning the specifications on the source host. It is also possible to reject articles that were posted more than a specified time ago, on the theory that they must be retransmissions from hosts that have been down. The posting rejection time is usually the same as that for article expiration because expired articles are also |

removed from the history file, and such retransmissions would thus not be rejected as duplicates.

The transport layer is UUCP for most links, although many others are used, including Ethernets and long-haul packet switched networks; sometimes UUCP is run on top of other transport layers, and sometimes UUCP is not used at all.

Naming, Addressing, and Routing

Every *USENET* host has a name. That host name and the name of the poster are used to identify the source of an article. Those hosts that are on both the *UUCP* mail and *USENET* news networks usually have the same name on both networks, but mail addresses have no meaning on *USENET*. Mail related to *USENET* articles is usually sent via *UUCP* mail; by definition, it cannot be sent over *USENET*. The two networks have always been closely related, but there are far more hosts on *UUCP* than on *USENET* and far more connections between hosts. In Australia, the two networks do not even intersect except at one host.

Interconnections

Since *USENET* is not a mail network, there are no mail gateways in the usual sense. However, *USENET* newsgroups are commonly interconnected with mailing lists on other networks because users on those networks want to participate in the discussions in the newsgroups, but their networks don't support the news service. Meanwhile, implementations of news for non-UNIX machines are becoming increasingly prevalent. For example, one of the largest machines on *USENET*, with about 30,000 users, is *CUNYVM.CUNY.EDU* [Reber 1988], which is best known as the root of the *BITNET* tree. It runs news and is connected as part of *USENET* as well. Here is how you would reach some networks from *USENET*:

- | | |
|-----------------|--|
| <i>Internet</i> | The machine that does the most gatewaying between <i>Internet</i> mailing lists and <i>USENET</i> newsgroups is known as <i>ucbvax</i> on <i>USENET</i> and as <i>ucbvax.berkeley.edu</i> on the <i>Internet</i> . It bursts digested mailing lists into individual articles. |
| <i>BITNET</i> | Although many newsgroups are being gatewayed with <i>BITNET</i> mailing lists, there appears to be no single machine that does most of the gatewaying. |
| <i>EASNet</i> | The main gateway for <i>EASNet</i> is <i>decurl</i> or <i>RHEA</i> , also known as <i>decurl.dec.com</i> . Of course, the user does not have to know the names of these gateways, or even that they exist. Specialized software at the gateways handles the conversions between news and mail automatically. |

Uses

USENET users pursue both technical and social ends, many of them organized by the newsgroup structure. But others are more general than that. At least one newsgroup has produced a series of conventions of its participants, not to mention t-shirts. There are even feuds between newsgroups, as in the long-running attempts of some *soc.women* participants to ban all male posters from that newsgroup; this has included the tactic of creating *soc.men*. More recently, *comp.society.women* has been created as a moderated newsgroup.

No description of *USENET* would be complete without mention of some of the notorious April Fool's Day jokes. These are derived from the ease of forging articles (such forgeries are usually readily detectable by a skilled eye). The most notorious example happened in 1984, when one "Konstantin Chernenko" announced that *kremvax*, *moskvax*, and he, *mcvax!moskvax!kremvax!chernenko*, were joining the network. Many users believed it, and the resulting follow-up articles and mail replies were numerous and interesting [Neumann 1984]. A more recent example, in 1988, was a purported warning by a well-known backbone host administrator about such April Fool's Day jokes: it was itself a forgery.

Effects

Relationships, marriages, conventions, lost friendships, software packages, papers, and books have resulted at least partly because of the use of *USENET*; many of them would not have happened without the network.

History

Jim Ellis and Tom Truscott were graduate students at Duke University (Duke) when they originated the idea of *USENET* in 1979 [Spafford 1988c]. According to an early description [Daniel et al. 1980], "A goal of *USENET* has been to give every UNIX system the opportunity to join and benefit from a computer network (a poor man's *ARPANET*, if you will)." The transport mechanism was UUCP, as distributed with the newly released Seventh Edition UNIX from AT&T Bell Laboratories. Access was limited to UNIX source licensees, as they were the only ones who could legally have the UUCP software (binary UNIX distributions had not been thought of yet).

The news software was not part of the Seventh Edition distribution but rather was written by Steve Bellovin, then a graduate student at the University of North Carolina (UNC) at Chapel Hill. The first two hosts were called *unc* and *duke*. A third host, *phs*, was added by 1980; it was administered by Dennis Rockwell, who also participated in the late 1979 planning sessions. Steve Bellovin rewrote the original shell script

implementation into the C programming language, but this was never released beyond the original universities. Stephen Daniel did another C implementation, which, after modifications by Tom Truscott, became *A news* and was later distributed widely, although it was never intended for more than a hundred hosts and a few articles per day [Spafford 1988c].

The originators soon invited others to join, partly by distributing a flyer [Ellis et al. 1980] at the next USENIX conference, 29 January through 1 February 1980 in Boulder, Colorado. The news software was distributed on the June 1980 USENIX Delaware Conference distribution tape [Daniel et al. 1980]. The USENIX Association is the oldest and largest group of users of the UNIX operating system. The name *USENET* was coined by James Ellis as a variant of "Usenix Network." The name USENIX itself is derived from UNIX, which is a pun on *Multics* [Organick 1975], the name of a major predecessor operating system. (The pun indicates that in areas where *Multics* tries to do many things, *UNIX* tries to do one thing well.) Alternative names for the network were considered, such as *CHAOSNET*, *ARACHNET*, and *Spidernet*, but most of them were already taken. The users would have liked a name including the word UNIX but that was not considered appropriate because of AT&T's trademark. Later, Mark Horton or Matt Glickman claimed *USENET* meant "User Network." Since the users had gained quite a bit of control from the founders (the existence and popularity of the newsgroup *net.jokes* attested to this), the new meaning was appropriate [Truscott 1988]. Even that meaning has changed slightly to "User's Network," which is the one most commonly accepted today.

According to Daniel et al., "New systems will probably find that the rapid access newsletter is the initially most significant service" [Daniel et al. 1980]. Newsgroups were originally thought of as "newsletters" partly because of the analogy with *:login:*, the USENIX paper newsletter, which was in a brief hiatus when *USENET* was started. Newsgroups were seen as a logical extension of mailing lists when those grew to have too many participants. Initially, there were only *NET.general* *NET.v7bugs*, with *NET.test* following soon afterward. There was no newsgroup hierarchy originally: that was added by Stephen Daniel. For example, you could subscribe to "ALL,!ALL.play," meaning everything but newsgroups with *play* as the second component.

Early traffic was low: "The articles from the Usenet Epoch (lets call it 00:00 Jan 1 1980) to 1982 would fit comfortably on a single 800bpi tape, not to mention our 200 megabyte /usr/spool filesystem. Once we might have retrieved them, but they are surely gone now" [Truscott 1988]. All the articles from 1982 to the end of 1987 fill 26 reels of 6250bpi tape, as recorded by Tom Truscott. Enough individual articles have survived from the early years to show that even then the network had a mix of technical and social

discussions, although the latter have outgrown the former in recent years. Growth, as measured in machines, users, newsgroups, articles, and data transferred and stored, has been the biggest historical problem with this network [Durlak et al. 1987].

The first of several news software rewrites to cope with growth of the network was done in 1981 at the University of California at Berkeley (UCB) by graduate student Mark Horton and high school student Matt Glickman. This was **B news**, first released to the public as **B news 2.1** in 1982. Mark Horton remained the principal maintainer through **B news 2.10.1** in 1984 [Spafford 1988c].

One of the effects of growth is articles that duplicate information in other articles, as when one article asks a question and many people post similar answers. This happens both because of the inherent delay of network propagation and because of new users who do not realize the likelihood that their articles will be duplicates in such a case. More users also lead to more unfocused or vituperative discussions, as well as other problems such as possible posting of licensed or copyrighted material. Moderation was introduced in 1984 as a solution to this problem after a great deal of debate. The idea was derived from the moderated mailing lists that had existed on the ARPANET for many years. The first moderated newsgroup was *mod.announce*, which was an alternative to *net.general*, a newsgroup formerly intended for general discussions throughout the network, but that had become useless because of the above problems.

Initial code for moderation was in **B news 2.10.1**, the last version released by Mark Horton, and depended on the first element of a moderated newsgroup's name being *mod*. **B news 2.10.2** was released by Rick Adams of the Center for Seismic Studies (CSS) in Arlington, Virginia, and completed on 18 September 1984. Adams added code that handled moderated newsgroups without depending on the *mod* keyword on 15 September 1986 and released **B news 2.11** in late 1986 [Horton and Adams 1987]. The distinction between distributions and newsgroups was not in early news software; it was also added in. The purpose is to allow users to limit individual postings to geographical areas appropriate to them, while still keeping such articles in newsgroups appropriate to their general topics.

Actual use of this software involved changing the names of newsgroups, and the opportunity was taken to change the names of *all* newsgroups in order to give them a more rational structure. This was the only time this has ever been done. There were originally only two main top level newsgroup elements, *net* and *mod*, and some others that indicated distributions. The name changes were done in two phases, each affecting about half of all newsgroups. The unmoderated newsgroups were changed in September 1986, and the moderated newsgroups were changed in April 1987.

Basically the same structure is still used, but some top level names and many more newsgroups have since been added [Durlak et al. 1987].

The most controversial decision involved in the newsgroup renaming was the creation of the *talk* top level name and the moving of many newsgroups to names under it. This was desirable because the number of non-technical newsgroups and their volume had become large enough that many backbone hosts did not want to carry them: a separate top level category for them would make that easy. Many users complained that this would cause the death of those newsgroups [Durlak et al. 1987]. Instead, a kind of alternative backbone, called the *jawbone*, was developed to carry them. The current *alt* and *bionet* top level newsgroups [Gilmore and Spafford 1988] are also results of this reorganization, and are also generally not carried by the backbone. Interestingly, some newsgroups that the backbone refused to carry, such as *net.rec.drugs*, have been created under new names as *alt* groups. There was an attempt before the reorganization to form a *funny bone* to carry such newsgroups, but that disappeared in the reorganization [Durlak et al. 1987].

The decisions to rename the newsgroups and to divide them into classes the backbone would carry and ones it would not led to a great deal of antipathy toward the reputed *backbone cabal* that made the decisions [Durlak et al. 1987]. This group actually comprised the administrators of the machines carrying most of the traffic, and thus paying most of the bills, and they considered the changes necessary to keep the network working. This was probably the high point of the backbone's power, for three main reasons: (1) the group deliberately encouraged the creation of an alternative backbone; (2) technological advances, such as much faster modems and *NSFNET*, have made the backbone less necessary; and (3) organizational developments such as the *inet* distribution (see below) and *UUNET* have also reduced the need for a backbone.

B news 2.11 included enhanced mechanisms for batching together groups of articles for transmission and for compressing them before transmission, as well as other features. This is the most common version of news software in use, but the distributed nature of the network means that there is no way to compel hosts to upgrade. There are still versions of **B news 2.9**, and maybe even **A news**, in use. Fortunately, **B news 2.11** is relatively clever at dealing with problem articles produced by such old software.

Geoff Collyer and Henry Spencer of the University of Toronto (Toronto) produced an alternative version of the basic news software, called **C news**, which they made available in 1987. Its main purpose was speed and reliability [Spencer and Collyer 1987]. This is not the only alternative version of the software in current use. The **notesfiles** software was

developed at the University of Illinois at Urbana-Champaign (UIUC) by two graduate students, Ray Essick and Rob Kolstad, starting in 1979 [Kolstad 1988]. They were inspired by the UIUC CDC *PLATO NOTES* system. The *notesfiles* software was first released to the public at the January 1982 USENIX conference in Santa Monica. This software originally spread separately from the news software. But even though different interchange formats and internal organizations were used, the *notes* and *USENET* networks were connected within six months of the *notesfiles* announcement [Kolstad 1988] and are now considered to be a single network, carrying most of the same newsgroups (or *notesfiles*) [Spafford 1988c].

Another new top level newsgroup is *inet*. This is used to distribute news over networks such as the *Internet*, producing marked savings by allowing messages to exist as one *USENET* news article per machine instead of many mail messages in individual users' mailboxes. UUCP can be used over TCP/IP, and there is even a way to do it so that its functions that duplicate those of TCP are not used. However, the usual method of transporting news over TCP/IP is the Network News Transfer Protocol (NNTP) [Kantor and Lapsley 1986].

The many problems with *USENET* (e.g., reader overload, old software, slow propagation speed, and high and unevenly carried costs of transmission) have raised the possibility of using the experience gained in *USENET* to design a new network to replace it. These problems have been considered so bad that the network has been declared almost dead several times. But it is not dead yet, due largely to technology having kept up with growth. For example, Spencer Thomas wrote a program that used Lempel-Ziv [Welch 1984] compression, and Rick Adams applied it to news transmission in *B news 2.10.2* [Adams 1988]. As shown in Table 10.5, effective transfer speeds in 1988 were an order of magnitude greater than they were in 1983. Of course, not every host connects to every other one at 12000bps yet, but most of the backbone hosts do, and they have always been the biggest bottleneck. Table 10.5 does not take into account cost-cutting factors such as more efficient news software, *PC Pursuit*, *UUNET*, or NNTP. It is not clear what the next advance will be, but DS1 and DS3 (1.544Mbps and 45Mbps) speeds are becoming increasingly plausible. In fact, faster speeds are already in use on some links: some backbone hosts exchange news over *NSFNET*, with effective speeds of 40Kbps or more.

Most of the *USENET* problems are hard to measure, but numbers of hosts and users, as well as volume of total articles, have been estimated since 1986 by a project headed by Brian Reid [Reid 1988b]. These are shown in Table 10.6, with the earlier figures coming from various sources [Frey and Adams 1987]. Details from that project on growth in recent years are shown in Table 10.7. The readership numbers are based on actual readers,

Table 10.5. *USENET technological improvements*

Date	Advance	Effective bps	Factor	Comment
1979	—	300	1×	news and notesfiles invented
1980	—	300	1×	A news distributed
1981	1200bps	1200	4×	
1982	—	1200	1×	B news 2.1, notesfiles, EUnet, SDN
1983	—	1200	1×	4.2BSD, System V, vnews
1984	Moderation	1700	1.4×	B news 2.10.1, m, JUNET, compression, revised moderation and distribution, B news 2.10.2
1985	2400bps	3400	2×	
1986	NNTP	3400	1×	4.3BSD, B news 2.11, jawbone
1987	—	3400	1×	System V.3, UUNET, newsgroup reorganization completed
1988	12000bps	12000	3.5×	Telebit Trailblazers

Sources: [Spafford 1988d; Truscott 1988; Kolstad 1988]

Table 10.6. *USENET growth*

Date	Newsgroups	Articles/day	MBytes/day	Hosts	Readers
1979	3	2	—	3	10?
1980	37	10	—	15	20?
1981	—	20	0.05	150	—
1982	—	35	—	400	—
1983	—	120	—	600	10,000?
1984	—	225	—	900	—
1985	—	375	>1	1,300	—
1986	241	946	2.0	2,200	53,000
1987	259	957	2.1	5,200	100,000
1988	381	1,933	4.4	7,800	141,000

Sources: [Spafford 1988d; Reid 1988a]

not just users, in a continuing survey of about 7 percent of all *USENET* hosts.

Comparing these figures with those of Table 10.5, we see that the hardest years were about 1983–1986, when growth had exceeded the capability of the original technology and further developments had not yet alleviated the problem. Disk space has also been a problem because all those data have to be stored on the machine for users to read. Cost per megabyte has diminished rapidly, however, and with developments such as 800Mbyte half-height Winchester drives for a few thousand dollars expected by 1989, an IBM PC could easily have enough space for the current traffic.

Table 10.7. Recent USENET growth

Month	News- groups	Mes- sages-	Mbytes	HC12	HC3	HC1	Sam- ple	Reader est.	User est.
Apr 1986	241	28,388	59,512	2,193	2,193	2,315	249	53,000	380,000
May 1986	244	27,476	59,155	2,618	2,618	2,352	296	59,000	420,000
Jun 1986	243	20,706	49,024	2,951	2,951	1,724	293	75,000	500,000
Jul 1986	236	19,521	46,218	3,052	2,828	1,692	320	95,000	640,000
Aug 1986	241	19,057	44,846	3,354	2,311	1,641	351	78,000	630,000
Sep 1986	262	19,652	45,593	3,607	2,506	1,765	320	82,000	660,000
Oct 1986	298	20,579	49,490	3,878	2,502	1,816	328	88,000	640,000
Nov 1986	314	18,259	43,233	4,104	2,550	1,684	354	88,000	590,000
Dec 1986	242	17,396	44,914	4,358	2,585	1,656	356	84,000	620,000
Jan 1987	252	26,809	62,321	4,683	2,653	2,056	347	79,000	570,000
Feb 1987	250	29,410	62,735	5,053	2,873	2,223	362	90,000	590,000
Mar 1987	259	28,223	62,036	5,376	3,149	2,238	398	90,000	550,000
Apr 1987	—	—	—	5,200	2,765	2,238	412	86,000	570,000
May 1987	259	29,689	61,272	5,052	2,945	2,249	435	105,000	550,000
Jun 1987	269	32,349	71,316	5,323	3,453	2,359	452	120,000	530,000
Jul 1987	262	33,493	77,531	5,519	3,514	2,585	476	118,000	640,000
Aug 1987	265	26,689	57,344	5,675	3,552	2,302	497	125,000	720,000
Sep 1987	312	32,665	74,395	5,929	3,711	2,632	514	134,000	810,000
Oct 1987	319	33,161	69,593	6,195	3,871	2,837	534	131,000	770,000
Nov 1987	321	40,843	82,451	6,368	4,064	2,797	589	140,000	730,000
Dec 1987	328	34,839	77,438	6,615	4,130	2,854	596	129,000	770,000
Jan 1988	330	50,593	110,838	6,925	4,363	3,337	645	137,000	840,000
Feb 1988	—	—	—	7,362	4,838	3,498	632	143,000	850,000
Mar 1988	347	53,144	121,561	7,362	4,398	3,498	641	142,000	910,000
Apr 1988	381	57,979	130,945	7,810	4,803	4,152	809	141,000	880,000

Source: [Reid 1988a]

Note: All data were measured at *decvri*.*Newsgroups* is the number of newsgroups that had nonzero traffic during the measurement interval.*Messages* is observed messages per month.*Mbytes* is megabytes of data per month.*HC12*, *HC3*, and *HC1* are three different ways of computing the "host count."*HC12* is the total number of distinct hosts mentioned in *Path:* lines in the past 12 months.*HC3* is the total number of distinct hosts mentioned in *Path:* lines in the past 3 months.*HC1* is the total number of distinct hosts in *Path:* lines in the past 30 days.*Sample* is the number of sites submitting arbitron reports.*Reader est.* is an estimate of the total number of readers.*User est.* is an estimate of the total number of users.

Regardless of how fast machines can transfer or store news, there is still the problem of readers being able to keep up with masses of material [Durlak et al. 1987]. The original user interface was called *readnews* and was line oriented. A screen-oriented interface called *vnews* was written by Kenneth Almquist and released in 1983. It is similar to the interface to the *notesfiles* software, and it allows moving among several menus, can list

contents of newsgroups, and is quite fast. The *rn* interface was written by Larry Wall of JPL and NASA and was originally released in 1984 [Spafford 1988c]. It provides ways to read, discard, or otherwise process articles based on patterns (regular expressions and logical expressions) specified by the user. One of its most popular commands is the *K* command, which allows skipping all articles with the subject of the current one, even in later reading sessions. All news reading interfaces allow the user to subscribe or unsubscribe to individual newsgroups at will, without having to go through a system administrator (assuming the host system itself subscribes to the desired newsgroups). All of them also allow a user who is reading a particular article to post a follow-up news article, to reply by mail (usually over the *UUCP* network), to save the article to a file, to mark it as unread so it will be seen again while the user is reading news, or to mark it as read so that it will not be seen again.

Relations between *USENET* and *USENIX* have been obscure since 1980, when Thomas Truscott said, "It is hoped that *USENIX* will take an active (indeed central) role in the network" [Truscott 1988]. Michael Tilson of Human Computing Resources (HCR) in Toronto proposed about 1982 that *USENIX* buy a Digital PDP-11 to use as a central node for *USENET*, but the *USENIX* board of directors has been wary of assuming any direct responsibility or control for a network as amorphous as *USENET*. Nonetheless, *USENIX* is willing to attempt technical improvements to the network and did sponsor the *Stargate* experiment, to use part of the vertical blanking interval on a satellite television transmission to carry news [Weinstein 1985; Weinstein 1986]. There were political problems with this service, such as the unpopularity of its decision not to permit forwarding of articles received from it. *USENIX* declared the technological experiment a success in February 1987 and left the operational service to survive as it might [Durlak et al. 1987]. Tilson's idea was later revived in an updated form as *UUNET*, which was sponsored by *USENIX* and is discussed elsewhere in this book.

Access

The *UUCP* map includes *USENET* map information as annotations. A list of legitimate netwide newsgroups is posted to several newsgroups monthly. Volunteers keep statistics on the use of the various newsgroups and on frequency of posting by persons and hosts. These are posted to *news.lists* once a month, as is the list of newsgroups. Important announcements are posted to two moderated newsgroups, *news.announce* and *news.announce.newusers*, which are intended to reach all users (the current moderators are Mark Horton and Gene Spafford, respectively).

For information on the network, contact:

uunet!usenet-request usenet-request@uunet.uu.net

CERN appear in the *HEPnet* section in this chapter and the CERN section in Chapter 13.

10.5 Bibliographic Notes

References for worldwide networks are few, but some of interest are Jennings 1987 and Landweber et al. 1986. There is some information on actual networks among the technical information in Tanenbaum 1988 and Partridge 1988.

10.6 References

- Adams 1988. Adams, Rick, Personal communications, 16 May 1988.
- Bach 1987. Bach, M. J., *The Design of the UNIX Operating System*, Prentice-Hall, Englewood Cliffs, NJ, 1987.
- Blokzijl 1988. Blokzijl, Rob, Personal communications, October–November 1988.
- Bush 1986. Bush, Randy, "A Basic FidoNet Technical Standard," IFNA, St. Louis, MO, 10 September 1986.
- Carpenter et al. 1987. Carpenter, B. E., Fluckiger, F., Gerard, J. M., Lord, D., and Segal, B., "Two Years of Real Progress in European HEP Networking: A CERN Perspective," *Computer Physics Communications*, vol. 45, pp. 83–92, August 1987.
- Chartrand and Demar 1985. Chartrand, Greg, and Demar, Philip, "DECnet and the High Energy Physics Network," *DECUS US Chapter SIGs Newsletters*, vol. 1, no. 1, pp. NTW-17–NTW-19, FERMILAB, P.O. Box 500, Batavia, IL 60510, September 1985.
- Crocker 1982. Crocker, David H., "Standard for the Format of ARPA Internet Text Messages; RFC822," *ARPANET Working Group Requests for Comments*, 13 August 1982.
- da Cruz 1987. da Cruz, Frank, "Kermit in Antarctica," *Kermit News*, vol. 2, no. 1, p. 6, November 1987.
- da Cruz 1988. da Cruz, Frank, Personal communications, October–November 1988.
- Daniel et al. 1980. Daniel, Stephen, Ellis, James, and Truscott, Tom, "USENET—A General Access UNIX Network," Duke University, Durham, NC, Summer 1980. This is a slightly modified copy of the original flyer of January 1980.
- Demco 1988. Demco, John, Personal communications, August 1988.
- Durlak et al. 1987. Durlak, Jerome, Yigit, Ozan, and O'Brien, Rory, "USENET: An Examination of the Social and Political Processes of a Cooperative Computer/Communications Network Under the Stress of Rapid Growth," *Proceedings of the Canadian Communications Association (Montreal, 20 May 1987)*, UQAM, Montreal, 1987.
- Eden 1988. Eden, Paul, Personal communications, November 1988.
- Ellis et al. 1980. Ellis, James, Truscott, Tom, and Daniel, Stephen, "Invitation to Join a General Access UNIX Network," USENIX Association, Boulder, CO, January 1980. About 40 copies of this flyer were distributed; none still exist.
- Emerson 1983. Emerson, Sandra L., "USENET: A Bulletin Board for UNIX Users," *BYTE*, vol. 8, no. 10, pp. 219–236, October 1983.

- Frey and Adams 1987. Frey, Donnalyn, and Adams, Rick, "USENET: Death by Success?," *UNIX REVIEW*, vol. 5, no. 8, pp. 55-60, August 1987.
- Fuchs 1983. Fuchs, Ira H., "BITNET: because it's time," *Perspectives in Computing*, vol. 3, no. 1, pp. 16-27, IBM Corporation, March 1983.
- Gilmore and Spafford 1988. Gilmore, John, and Spafford, Gene, "Alternative Newsgroup Hierarchies (Updated: 11 March 1988)," *news.lists, news.groups, news.announce.newusers*, USENET, 3 April 1988.
- Henderson 1986. Henderson, Thom, SEAdog Electronic Mail System Version 3, Thom Henderson, April 1986.
- Hendricks and Hartman 1979. Hendricks, E. C., and Hartman, T. C., "Evolution of a virtual machine subsystem," *IBM Systems Journal*, vol. 18, no. 1, p. 111, 1979.
- HEPnet 1988. HEPnet, "High Energy Physics Computer Networking, Report of the HEPNET Review Committee," Department of Energy, Washington, DC, June 1988.
- Honeyman and Bellovin 1986. Honeyman, Peter, and Bellovin, Steven M., "Pathalias or the Care and Feeding of Relative Addresses," *Proceedings of the 1986 Summer USENIX Conference (Atlanta, 9-13 June 1986)*, pp. 126-141, USENIX Association, Berkeley, CA, 1986.
- Horton 1984. Horton, Mark, "What is a Domain?," *Proceedings of the Summer 1983 USENIX Conference (Salt Lake City, Utah, 12-15 June 1984)*, pp. 368-372, USENIX Association, Berkeley, CA, 1984.
- Horton 1986. Horton, Mark R., "UUCP Mail Interchange Format Standard; RFC976," *ARPANET Working Group Requests for Comments*, February 1986.
- Horton and Adams 1987. Horton, Mark R., and Adams, Rick, "Standard for Interchange of USENET Messages; RFC1036," *ARPANET Working Group Requests for Comments*, December 1987.
- Horton et al. 1984. Horton, Mark, Summers-Horton, Karen, and Kercheval, Berry, "Proposal for a UUCP/Usenet Registry Host," *Proceedings of the Summer 1983 USENIX Conference (Salt Lake City, Utah, 12-15 June 1984)*, p. 373, USENIX Association, Berkeley, CA, 1984.
- Jennings 1987. Jennings, Dennis M., "Computing the best for Europe," *Nature*, vol. 329, no. 29, pp. 775-778, October 1987.
- Jennings 1983. Jennings, Tom, "Extending XMODEM/MODEM File Transfer Protocol to support DOS," IFNA, St. Louis, MO, 20 September 1983.
- Jennings 1985a. Jennings, Tom, "Fido's Internal Structures," IFNA, St. Louis, MO, 13 September 1985.
- Jennings 1985b. Jennings, Tom, "FidoNet Electronic Mail Protocol," IFNA, St. Louis, MO, 8 February 1985.
- Kantor and Lapsley 1986. Kantor, Brian, and Lapsley, Phil, "Network News Transfer Protocol; RFC977," *ARPANET Working Group Requests for Comments*, February 1986.
- Kiessig 1986. Kiessig, Richard, *UUCP Network Directory*, vol. 2, no. 2, Spring 1986. (P.O. Box 50174, Palo Alto, CA 94303)
- Kille 1986. Kille, S. E., "Mapping between X.400 and RFC822; RFC987," *ARPANET Working Group Requests for Comments*, June 1986.
- Kille 1987. Kille, S. E., "Addendum to RFC 987 (Mapping between X.400 and RFC-822); RFC1026," *ARPANET Working Group Requests for Comments*, September 1987.
- Kincl 1989. Kincl, Norman, Personal communications, 10 March 1989.
- Kincl and Michaels 1988. Kincl, Norman, and Michaels, Robert, "Inside Autonomous System 71 — The HP Internet," *ConneXions — The Interoperability Report*, vol. 2, no. 9, pp. 2-7, Advanced Computing Environments, Mountain View, CA, September 1988.

- Klein 1988. Klein, Ben, "Networking at CUNY: an Overview," *CUNY University Computer Center Communications*, vol. 14, no. 3/4, pp. 15-21, CUNY, New York, March/April 1988.
- Kolstad 1988. Kolstad, Rob, Personal communications, 1 December 1988.
- Landweber et al. 1986. Landweber, Lawrence H., Jennings, Dennis M., and Fuchs, Ira, "Research Computer Networks and Their Interconnection," *IEEE Communications*, vol. 24, no. 6, pp. 5-17, June 1986.
- LaPedis and DeBra 1987. LaPedis, Ron, and DeBra, Corinne, "Network Noted," *Communications of the ACM*, vol. 30, no. 3, p. 195, March 1987.
- LaQuey 1988. LaQuey, Tracy Lynn, *Users' Directory of Computer Networks*, Office of Telecommunication Services, University of Texas System, Austin, July 1988.
- Leffler et al. 1989. Leffler, Samuel J., McKusick, Marshall Kirk, Karels, Michael J., and Quarterman, John S., *The Design and Implementation of the 4.3BSD UNIX Operating System*, Addison-Wesley, Reading, MA, 1989.
- Madsen and Foley 1985a. Madsen, Kent, and Foley, David, "When the shoe's on the other foot: A look at Tandem's corporate network," *Data Communications*, vol. 14, no. 9, pp. 183-193, August 1985.
- Madsen and Foley 1985b. Madsen, Kent, and Foley, David, "How multi-processor nodes can become more sociable," *Data Communications*, vol. 14, no. 10, pp. 279-294, September 1985.
- Neumann 1984. Neumann, Peter G., "Spring Fever—Ape Rills are Flowing (or Fooling?)," *Software Engineering Notes*, vol. 9, no. 4, pp. 6-8, ACM SIGSOFT, July 1984.
- Nowitz 1979. Nowitz, D. A., "Uucp Implementation Description," in *UNIX Programmer's Manual, Seventh Edition*, vol. 2, pp. 577-591, Holt, Rinehart and Winston, New York, 1983.
- Nowitz and Lesk 1979. Nowitz, D. A., and Lesk, M. E., "A Dial-Up Network of UNIX Systems," in *UNIX Programmer's Manual, Seventh Edition*, vol. 2, pp. 569-576, Holt, Rinehart and Winston, New York, 1983.
- Nussbacher 1987. Nussbacher, Henry, "Conversion Plans from RSCS to Tcp/Ip for Israel Academic Network," *TCP-IP@SRI-NIC.ARPA*, Israeli University Telecommunications Subcommittee, Tel Aviv, 7 July 1987.
- Nussbacher 1988. Nussbacher, Henry, "BITNET GATES Version 88.09," *BITNET/EARN NETSERV*, 25 August 1988.
- Organick 1975. Organick, Elliot I., *The Multics System: An Examination of Its Structure*, MIT Press, Cambridge, MA, 1975.
- Partridge 1988. Partridge, Craig, *Innovations in Internetworking*, ARTECH House, Norwood, MA, October 1988.
- Price 1988. Price, Larry, Personal communications, 7 December 1988.
- Quarterman et al. 1985. Quarterman, John S., Silberschatz, Abraham, and Peterson, James L., "4.2BSD and 4.3BSD as Examples of the UNIX System," *ACM Computing Surveys*, vol. 17, no. 4, pp. 379-418, December 1985.
- Reber 1988. Reber, Pat, "NETNEWS: A Mainframe Bulletin Board," *CUNY University Computer Center Communications*, vol. 14, no. 5-6, pp. 29-31, CUNY, New York, May-June 1988.
- Reid 1988a. Reid, Brian, "Network Maps (DECWRL netmap 1.5)," DEC Western Research Lab, Palo Alto, 19 June 1988.
- Reid 1988b. Reid, Brian, "USENET Readership Summary Report for Mar 88," *news groups, news lists*, USENET Measurement Project, DEC Western Research Laboratory, Palo Alto, CA, 3 April 1988.
- Ritchie and Thompson 1974. Ritchie, D. M., and Thompson, K., "The UNIX Time-Sharing System," *Communications of the ACM*, vol. 7, no. 7, pp. 365-375, July 1974. This is the original version, which describes Sixth Edition UNIX.

- Ritchie and Thompson 1978. Ritchie, D.M., and Thompson, K., "The UNIX Time-Sharing System," *Bell System Technical Journal*, vol. 57, no. 6, part 2, pp. 1905-1929, July/August 1978. This version describes Seventh Edition UNIX.
- Schroeder et al. 1984. Schroeder, M.D., Birrell, A.D., and Needham, R.M., "Experience with Grapevine: The Growth of a Distributed System," *ACM Transactions on Computer Systems*, vol. 2, no. 1, pp. 3-23, February 1984. Also in Partridge, *Innovations in Internetworking*, 1988.
- Spafford 1988a. Spafford, Gene, "The USENET Backbone (Updated: 3 April 1988)," *news.misc, news.config*, USENET, 4 April 1988.
- Spafford 1988b. Spafford, Gene, "List of Active Newsgroups (Updated: 2 April 1988)," *news.lists, news.groups, news.announce.newusers*, USENET, 3 April 1988.
- Spafford 1988c. Spafford, Gene, "USENET Software: History and Sources (Updated: 2 April 1988)," *news.admin, news.announce.newusers*, USENET, 3 April 1988.
- Spafford 1988d. Spafford, Gene, "USENET," *IETF Proceedings*, IETF, Washington, DC, October 1988.
- Spafford 1988e. Spafford, Gene, "Publicly Accessible Mailing Lists (Updated: 14 April 1988)," *news.lists, news.announce.newusers*, USENET, 17 April 1988.
- Spencer and Collyer 1987. Spencer, Henry, and Collyer, Geoff, "News Need Not Be Slow," *Proceedings of the Winter 1987 USENIX Conference (Washington, D.C., 21-23 January 1987)*, USENIX Association, Berkeley, CA, 1987.
- Summers-Horton and Horton 1985. Summers-Horton, Karen, and Horton, Mark, "Status of the USENIX UUCP Project," *Proceedings of the Winter 1985 USENIX Conference (Dallas, Texas, 23-25 January 1985)*, p. 183, USENIX Association, Berkeley, CA, 1985.
- Tanenbaum 1988. Tanenbaum, Andrew S., *Computer Networks*, 2d ed., Prentice-Hall, Englewood Cliffs, NJ, 1988.
- Todino and O'Reilly 1988a. Todino, Grace, and O'Reilly, Tim, *Using UUCP and Usenet*, O'Reilly & Associates, Inc., Newton, MA, 1988.
- Todino and O'Reilly 1988b. Todino, Grace, and O'Reilly, Tim, *Managing UUCP and Usenet*, O'Reilly and Associates, Inc., Newton, MA, 1988.
- Tokuda et al. 1986. Tokuda, T., Landweber, L.H., Faber, D.J., Horton, M.R., Tokuda, H., Adrion, W.R., and Horton, K.S., "Toward the network community for R&D (kenkyuukaihatsuyou no network community wo mezashite)," *BIT Magazine*, vol. 18, no. 2, p. 122, 1986. Transcript of a panel discussion, in Kanji.
- Truscott 1988. Truscott, Thomas, Personal communications, 16 February 1988.
- Weinstein 1985. Weinstein, Lauren, "Project Stargate," *Proceedings of the Summer 1985 USENIX Conference (Portland, Oregon, 11-14 June 1985)*, pp. 79-80, USENIX Association, Berkeley, CA, 1985.
- Weinstein 1986. Weinstein, Lauren, "Project Stargate," *EUUG Spring '86 Conference Proceedings (Florence, Italy, 21-24 April 1986)*, European UNIX systems User Group, Florence, Italy, 1986.
- Welch 1984. Welch, Terry A., "A Technique for High Performance Data Compression," *IEEE Computer*, vol. 17, no. 6, pp. 8-19, June 1984.
- Wells 1987. Wells, William C., "PHYSnet/CITnet/HEPnet/NRAO/SPAN — Physics Community DECnet Mail," Data Communication and Network Services, University of California at Berkeley, 14 September 1987.
- Wells 1988a. Wells, William C., "HEPNET, PHYSNET and SPAN," *info-nets@think.com*, BITNET, Internet, Berkeley, CA, 6 June 1988.
- Wells 1988b. Wells, William C., Personal communications, 12 September 1988.