

ORNL/TM-12680

Engineering Physics and Mathematics Division
Mathematical Sciences Section

NETLIB SERVICES AND RESOURCES

Shirley V. Browne †
Jack J. Dongarra ‡‡
Stan C. Green †
Keith Moore †
Thomas H. Rowan †
Reed C. Wade †

† Department of Computer Science
University of Tennessee
107 Ayres Hall
Knoxville, TN 37996-1301

‡ Mathematical Sciences Section
Oak Ridge National Laboratory
P.O. Box 2008, Bldg. 6012
Oak Ridge, TN 37831-6367

Date Published: April 1994

Research was supported by the National Science Foundation under Grant. No. ASC-9103853, by the Advanced Research Projects Agency under contract DAAL03-92-G-0284 administered by the Army Research Office, and by the Applied Mathematical Sciences Research Program of the Office of Energy Research, U.S. Department of Energy.

Prepared by the
Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831
managed by
Martin Marietta Energy Systems, Inc.
for the
U.S. DEPARTMENT OF ENERGY
under Contract No. DE-AC05-84OR21400

MASTER
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED *870*

Contents

1	Introduction	1
2	Repository Contents	3
2.1	Software and Documents	3
2.1.1	Software and Document Submissions	3
2.2	Performance Database	3
2.3	Conferences Database	4
2.4	Numerical Analysis Net (NA-NET)	4
3	User Interfaces	5
3.1	Email	5
3.1.1	Support Addresses	5
3.1.2	Email Interface to the Netlib Repository	5
3.1.3	Email Interface to NA-NET	6
3.1.4	Email Interface to the NA-NET Whitepages	8
3.2	X Windows Interface – the Xnetlib Client	10
3.2.1	Acquiring the Xnetlib Software	11
3.2.2	System Requirements	11
3.2.3	Building Xnetlib	12
3.2.4	Xnetlib Man Page and Quick Reference Card	12
3.2.5	Xnetlib Client Program Operation	13
3.2.6	Command Line Options	16
3.2.7	Application Defaults File	16
3.2.8	Support Address	16
3.3	Netlibget, a Command-line TCP/IP Client	17
3.4	Anonymous Access	18
3.4.1	Instructions for Anonymous FTP	18
3.4.2	Instructions for Anonymous RCP and RSH	18
3.5	Access via Gopher	18
4	Future Plans	19
5	Netlib Repository Setup and Maintenance	20
5.1	Netlib Index File Format	20
5.2	Repository Replication in Netlib	22
6	Email Netlib Server Installation	23
6.1	Acquiring and Installing the Netlib Software	23
7	NA-NET Database Setup and Maintenance	25
7.1	Acquiring the NA-NET Software	25
7.2	The NA-NET Program	25
7.3	NA-NET Files	25
7.4	File Formats	27
7.5	NA-NET Source Files	28
7.6	Database Changes, Backups, and Cron Jobs	28

7.7	Sending to Digest Subscribers	29
7.8	Surgery	29
8	Installation and Customization of the Xnetlib Client	29
8.1	Acquiring the Xnetlib Software	29
8.2	System Requirements	30
8.3	Building Xnetlib	30
8.4	Customization of Xnetlib	31
8.4.1	X Resources	31
9	Installing and Running Nlreexecd	32
9.1	Acquiring the Nlreexecd Software	32
9.2	System Requirements	32
9.3	Building Nlreexecd	32
9.4	Services and Protocol	32
9.4.1	Adding a Service to Nlreexecd	33
9.4.2	Reserved Service Names	34
9.5	Command Line Options	35
9.6	Keyword and Database Lookup Services	35
9.6.1	Keyword Lookup	35
9.6.2	Latent Semantic Indexing	35
9.6.3	Whois Service	35
9.6.4	Performance Database Service	36
9.6.5	Conference Database Service	36
10	Anonymous FTP Server for Netlib	36
11	Netlib Anonymous RCP Implementation	38
11.1	“anon” Account	38
11.2	Invocation of Anonymous RCP	38
11.3	anon-shell	38
11.4	Modified rcp and ls commands	39
11.5	Locations of files	39
A	Netlib Sites	40
A.1	Sites Mirroring the Netlib Repository	40
A.2	Some Sites Using the Netlib Email Server to Distribute Other Types of Software	40
2	References	40

NETLIB SERVICES AND RESOURCES

Shirley V. Browne
Jack J. Dongarra
Stan C. Green
Keith Moore
Thomas H. Rowan
Reed C. Wade

Abstract

The Netlib repository, maintained by the University of Tennessee and Oak Ridge National Laboratory, contains freely available software, documents, and databases of interest to the numerical, scientific computing, and other communities. This report includes both the Netlib User's Guide and the Netlib System Manager's Guide, and contains information about Netlib's databases, interfaces, and system implementation. The Netlib repository's databases include the Performance Database, the Conferences Database, and the NA-NET mail forwarding and Whitepages Databases. A variety of user interfaces enable users to access the Netlib repository in the manner most convenient and compatible with their networking capabilities. These interfaces include the Netlib email interface, the Xnetlib X Windows client, the netlibget command-line TCP/IP client, anonymous FTP, anonymous RCP, and gopher.

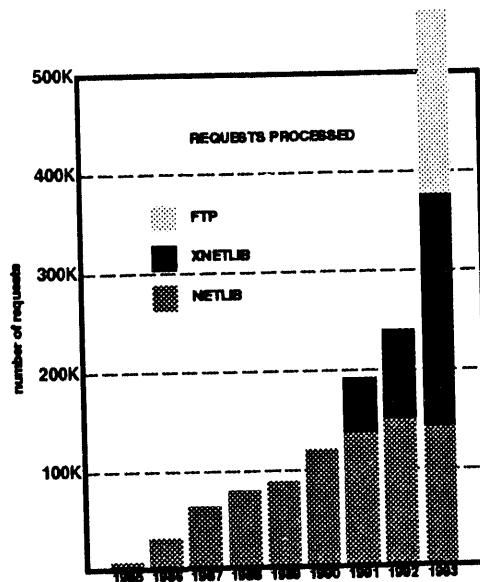


Figure 1: Netlib Requests

1. Introduction

Netlib began services in 1985 to fill a need for cost-effective, timely distribution of high-quality mathematical software to the research community. Netlib sends, by return electronic mail, requested routines together with subsidiary routines and any requested documents or test programs supplied by the software authors [5]. Xnetlib, a recently developed interactive tool for software and document distribution [4], uses an X Window interface and TCP/IP connections to allow users to receive replies to their requests within a matter of seconds. The interface provides a number of modes and searching mechanisms to facilitate searching through a large distributed collection of software and documents. The `netlibget` command-line interface and anonymous FTP and RCP provide services to users who do not need a sophisticated interface. Figure 1 shows the growing number of requests for Netlib services.

A new hypertext/hypermedia version of Xnetlib, currently under development, will interoperate with other information services such as gopher, WAIS, and *World Wide Web*. It will incorporate a new type of executable document, called an *active object*, that will greatly enhance the flexibility and adaptability of Xnetlib by allowing runtime binding of functionality.

Although the original focus of the Netlib repository was on mathematical software, the collection has grown to include other software (such as networking tools and tools for visualization of multiprocessor performance data), technical reports and papers, a Whitepages Database, benchmark performance data, and information about conferences and meetings. The number of Netlib servers has grown from the original two, at Oak Ridge National Laboratory (initially at Argonne National Laboratory) and AT&T Bell Laboratories, to several sites worldwide, including servers in Norway, the United Kingdom, Australia, Japan, and Taiwan. A mirroring mechanism keeps the repository

contents at the different sites consistent on a daily basis. This mechanism provides redundancy in case of computer or network failures, shares the workload, and broadens human contacts for identifying software to add to the collection.

Netlib differs from other publicly available software distribution systems, such as **Archie**, in that the collection is moderated by an editorial board and the software contained in it is widely recognized to be of high quality. The user is assured of getting an up-to-date copy of the master version of the requested software. We log requests so we can send bug reports and updates to users of our software. However, the Netlib repository is not intended to replace commercial software. Commercial software companies provide value-added services in the form of support. Although the Netlib collection is moderated, its software comes with no guarantee of reliability or support. Rather, the lack of bureaucratic, legal, and financial impediments encourages researchers to submit their codes by ensuring that their work will be made available quickly to a wide audience.

Requests for consideration of software and document submissions to Netlib, as well as questions, comments, and problems concerning Netlib, should be sent to the following address:

netlib_maintainers@netlib.ornl.gov

Netlib User's Guide

2. Repository Contents

2.1. Software and Documents

The Netlib repository contains a large collection of high-quality public-domain mathematical software. In addition, the repository contains other material of interest to the scientific computing community, including software documentation, test data, technical papers, and reports. Most of the software is written in Fortran, but programs in other languages, such as C, are also available. Netlib is intended primarily for users who wish to obtain individual routines or pieces of a package. Netlib supports dependency checking, so that all of the routines a particular routine depends on (i.e., calls) can be retrieved with it.

The software is organized as a UNIX directory tree. The subdirectories under the Netlib root are called *libraries*. Each library contains an index file describing the library contents: the files for the individual routines and any subdirectories the library may have. Software routines from a variety of sources are currently available from Netlib. Some of the libraries Netlib distributes – such as EISPACK, LINPACK, FFTPACK, LAPACK, algorithms from the ACM *Transactions on Mathematical Software*, and algorithms from the book by Forsythe, Malcolm, and Moler [6] – have long been used as important tools in scientific computation and are widely recognized to be of high quality. The Netlib collection also includes a large number of newer, less well-established codes.

2.1.1. Software and Document Submissions

To submit software or documents for inclusion in the Netlib repository, follow the guidelines in `/netlib/misc/contrib`, which you can obtain either by email as follows:

```
mail netlib@ornl.gov
send contrib from misc
```

or by downloading the `contrib` file from the `misc` library using the Xnetlib client. Send your submission to the following address:

```
netlib_maintainers@netlib.ornl.gov
```

2.2. Performance Database

The Performance Database is a publicly-accessible central repository of performance data for all ranges of machines, from personal computers to supercomputers. It provides an on-line catalog of public-domain computer benchmarks such as the LINPACK Benchmark, Perfect Benchmarks, and the NAS Parallel Benchmarks. The benchmark data are presented in a way that facilitates objective interpretations of machine performance. The Performance Database allows all branches of the computing community to archive performance metrics and makes them readily available to the public. For further details, see [1].

The performance data are stored in the **performance** directory in the Netlib repository. Although it is possible to download performance data using anonymous FTP or RCP, the **Performance** button on the Xnetlib X Windows client provides browsing and keyword searching mechanisms, as well as formatted display of the data. See Section 3.2 for further information.

2.3. Conferences Database

The Conferences Database contains conference and meeting announcements. Although it is possible to download conference descriptions using anonymous FTP or anonymous RCP, use of the **Conferences** button on the Xnetlib client allows searching separately by title, dates, location, or description keywords, as well as viewing of the results in a window. Furthermore, conferences can be submitted only through the Xnetlib client.

The conference description files are stored in the Netlib repository in the **confdb** directory. The title, start_date, end_date, location, submitter, and submit_date fields are stored in a relational database, currently Postgres. The contents of the title, location, and description fields are added to a full-text index when a conference is entered into the database. A geographical database handles location name aliases and hierarchical geographical relationships. See Section 3.2 for further information.

2.4. Numerical Analysis Net (NA-NET)

NA-NET is a community of numerical analysts and other researchers who communicate through a common email facility. NA-NET is supported by a mail-forwarding database and a Whitepages Database. The mail-forwarding database gives users an easy method of communicating with each other through the use of a uniform email address. This feature avoids the problem of having to remember an individual's specific email address. Mail to an NA-NET member can be addressed to **na.<key>@na-net.ornl.gov**, where **<key>** is the member's NA-NET key. The key is usually the member's last name, possibly prefixed by the first letter of the first name. The mail-forwarding database also serves as the distribution list for the NA-NET News Digest. To join or use NA-NET, use the email interface to NA-NET. See Section 3.1.3 for further details.

The NA-NET Whitepages is a directory service that allows NA-NET members to find out more information about other members. Anyone can join the Whitepages, but NA-NET members are especially encouraged to join. The Whitepages can be accessed by means of the NA-NET Whitepages email interface, the **netlibget** command-line TCP/IP client, or the **Who** button on the Xnetlib client. Currently, the Whitepages can be joined only through the email interface. See Section 3.1.4 for further details.

The SIAM Membership Directory is a separate directory that is copyrighted by the Society for Industrial and Applied Mathematics. Queries to the NA-NET **whois** service that do not find an entry in the NA-NET Whitepages database are referred to the SIAM list.

3. User Interfaces

3.1. Email

Anyone with an email connection to the Internet can access most of the Netlib repository. There are email interfaces to the software and document libraries, and the NA-NET mail-forwarding and Whitepages Databases. To receive more information about the email interface to Netlib, send a message to `netlib@ornl.gov` with the message body `send index`. To receive more information about the email interface to NA-NET, send a message to `na.help@na-net.ornl.gov`.

3.1.1. Support Addresses

There are email support addresses for the different Netlib services. Users with comments, questions, or bug reports should send a message to the appropriate support address listed below.

Netlib repository	<code>netlib_maintainers@netlib.ornl.gov</code>
Xnetlib client	<code>xnetlib@cs.utk.edu</code>
Performance database	<code>utpds@cs.utk.edu</code>
Conferences database	<code>conferences@cs.utk.edu</code>
NA-NET	<code>nanet@na-net.ornl.gov</code>
Gopher server	<code>gopher@netlib.ornl.gov</code>

3.1.2. Email Interface to the Netlib Repository

Netlib email addresses

The Internet address `netlib@ornl.gov` refers to a gateway machine at Oak Ridge National Laboratory in Oak Ridge, Tennessee. This address should be understood on all the major networks through the normal Domain Name System name resolution.

For access outside the United States, you may want to use one of the repositories that mirror the UT/ORNL repository. See Appendix A for a list of these other Netlib sites.

Request syntax

A valid Netlib email request has a message body that is of one of the following basic forms:

```
send index
send index from <library>
send <file(s)> from <library>
find <keywords>
whois <name>
mailsize <size>
```

Here are examples of the various kinds of requests.

- * To get the master index for netlib:
`send index`
- * To get the full index for a library:
`send index from eispack`
- * To get a particular routine and all it depends on:
`send dgeco from linpack`
- * To get just the one routine, not subsidiaries:
`send only dgeco from linpack`
- * To get dependency tree, but excluding a subtree:
`send dgeco but not dgefa from linpack`
- * To just tell how large a reply would be, don't actually send the file:
`send list of dgeco from linpack`
- * To get a list of sizes and times of all files in a library:
`send directory for benchmark`
- * To search for somebody in the SIAM membership list:
`whois gene golub`
- * To do a keyword search for Netlib software:
`find cubic spline`
- * To do a bibliographic search:
`find schumaker from approximation`
`find aasen from linalg`
- * To set the chunk size used for reply:
`mailsize 100k`
- * (optional) End of request:
`quit`

3.1.3. Email Interface to NA-NET

Individual (unicast) messages

Sending email to an individual NA-NET member is the most frequently used feature of NA-NET. Each NA-NET member has a unique NA-NET name, or key. Mail can be sent to an NA-NET member by addressing it to `na.<key>@na-net.ornl.gov`, where `<key>` is the member's NA-NET name. The NA-NET name is usually the member's first initial prepended to her last name, the member's last name, or the member's first

name followed by the first letter of her last name. For example, possible NA-NET names for Joan Smith would be `jsmith`, `smith`, and `joans`.

NA-NET News Digest

Any mail sent to `na.digest@na-net.ornl.gov` will be considered for distribution to all members of NA-NET. Once a week, we send out a digest of information contributed by users of NA-NET. The editor of the NA-NET News Digest goes over the messages that have been received, picks out the ones thought to be of general interest to the numerical analysis community, combines them in the News Digest format, and mails the Digest to everyone on the mailing list.

Joining NA-NET

To join NA-NET, send mail to `na.join@na-net.ornl.gov`. In the message body, specify the following three fields:

`Lastname: <your last name>`
`Firstname: <your first name>`
`E-mail: <your e-mail address>`

The values can be specified in any order. The subject line of your message will be ignored. An attempt will be made to assign to you a unique NA-NET name consisting of your first initial prepended to your last name, your last name, or your first name followed by the first letter of your last name. If at least one of these keys is not already in use, you will receive a message indicating that your join attempt succeeded and telling you which key has been assigned. If all three of these keys fail to be unique, you will receive an error message indicating that your join attempt failed. In case of failure, send a message to `nanet@na-net.ornl.gov`, and you will be assigned a unique key manually.

Removing membership

To remove your membership from NA-NET, send mail to `na.remove@na-net.ornl.gov`. In the message body, specify the following two fields:

`Lastname: <your last name>`
`Firstname: <your first name>`

The values can be specified in any order. The subject line of your message will be ignored. NA-NET will send an acknowledgment message to both the deleted address and the address making the request. If more than one entry exists with the same first and last name, you will receive a message indicating that your removal attempt failed. In this case, you can resubmit the removal request with the additional line:

`Key: <your NA-NET key>`

Changing your email address

To change your email address in the NA-NET mail-forwarding database, send mail to `na.change@na-net.ornl.gov`. In the message body, specify the following three fields:

`Lastname: <your last name>`
`Firstname: <your first name>`
`New-address: <your new e-mail address>`

The values can be specified in any order. The subject line of your message will be ignored. An acknowledgment message will be sent to both the old email address as well as the new address informing you that the change has taken place.

If more than one entry exists with the same first and last name, you will receive a message indicating that your change attempt failed. In this case, you can resubmit the change request with the additional line:

`Key: <your NA-NET key>`

Help with NA-NET

Questions and comments about NA-NET should be addressed to `nanet@na-net.ornl.gov`. Mail sent to `na.help@na-net.ornl.gov` will receive a reply message describing both NA-NET and the Whitepages.

Current member list

Mail sent to `na.sendlist@na-net.ornl.gov` will receive a reply message being sent back to you containing the email addresses of all members of NA-NET.

3.1.4. Email Interface to the NA-NET Whitepages

Querying the Whitepages Database

To find out information about a person, send mail to `na.whois@na-net.ornl.gov`. In the message body or on the subject line specify the person's first name and last name, or just the last name. The order of first name and last name does not matter. For example, to find out more about Jack Dongarra:

`mail to: na.whois@na-net.ornl.gov`
`Subject:`

`Jack Dongarra`

`or`

`mail to: na.whois@na-net.ornl.gov`
`Subject: jack dongarra`
`<null body>`

Keyword searching is also possible. For example, to find out more information about all people who are interested in *parallel*, send the following message:

```
mail to: na.whois@na-net.ornl.gov
Subject:
```

```
Keyword: parallel
```

This query does a string search for the pattern *parallel* on all fields other than the name fields.

As another example, to find out more about all people who live in Knoxville:

```
mail to: na.whois@na-net.ornl.gov
Subject:
```

```
Keyword: Knoxville
```

This query does a string search on all fields other than the name fields for the pattern Knoxville.

Joining the Whitepages

To join the NA-NET Whitepages, send mail to `na.join-wp@na-net.ornl.gov`. In the message body, specify the two mandatory fields and as many of the optional fields as you want.

Mandatory

```
Last_name: <your last name>
```

```
First_name: <your first name>
```

Optional

```
Middle_name:
```

```
Other_name:
```

```
Affiliation:
```

```
Office_address:
```

```
City_state_zip:
```

```
Country:
```

```
Office_phone:
```

```
Research:
```

```
Home_address:
```

```
Home_phone:
```

```
Fax:
```

```
E_mail_address:
```

```
Other:
```

The fields can be specified in any order. The subject line of your message will be ignored. All fields are entered into the database as characters, so spaces can be used for readability. All fields except first_name, last_name, and middle_name can be multiple lines. A multiple-line field ends when the next keyword (e.g., "Country:") is encountered. Each line should end with a carriage return. If your first_name and last_name combination is not unique, send mail to `nanet@na-net.ornl.gov`, and your name will be manually inserted into the Whitepages Database in spite of the duplication. Such duplication will not cause any problems for people querying the Whitepages Database because the database is set up to return information on all people with a given last name, first_name, or combination. An acknowledgment to your join request will be sent back to you confirming that the operation was successful.

Removing your Whitepages entry

To remove your entry from the NA-NET Whitepages Database, send mail to `na.remove-wp@na-net.ornl.gov`. In the message body, specify the following two fields:

```
Last_name: <your last name>
First_name: <your first name>
```

The values can be specified in any order. The subject line of your message will be ignored. An acknowledgment message will be sent to both the address requesting the removal and to the address listed in the Whitepages Database.

Changing fields

To change the value of a field, to add a field, or to delete a field, send mail to `na.change-wp@na-net.ornl.gov`. In the message body, specify the following two fields:

```
Last_name: <your last name>
First_name: <your first name>
```

plus the fields to be added, changed, or dropped. The fields can be specified in any order. The subject line of your message will be ignored. You can not change your name. If you need to change your name, first remove your entry and then rejoin with the new name. If you wish to clear the value of a field, simply include the field with no value. An acknowledgment message will be sent back to you confirming that the operation was successful.

3.2. X Windows Interface – the Xnetlib Client

Xnetlib is an X Window System application that provides interactive file access and database query processing from multiple servers through TCP/IP connections. Xnetlib currently provides access to the Netlib software and document repository, the NA-NET Whitepages Database, the Performance Database, and the Conferences Database. Future releases of Xnetlib will provide additional features, such as access to remote

execution facilities and interoperability with other information services, such as **gopher**, **WAIS**, and **World-Wide-Web**.

3.2.1. Acquiring the Xnetlib Software

To acquire the software for the Xnetlib client send email to **netlib@ornl.gov** with the line

```
send xnetlib.shar from netlib
```

as the body of the message. Netlib will return the file **xnetlib.shar** by email.

Xnetlib is available by anonymous FTP from **netlib2.cs.utk.edu** in the **xnetlib** directory. Both executables and source are available.

To use anonymous FTP to retrieve an executable file, type

```
ftp netlib2.cs.utk.edu
anonymous
<your email address>
cd xnetlib
binary
get xnetlib.<arch>.Z
bye
```

where **<arch>** is your machine architecture (**alpha**, **hp9000**, **next**, **pmax**, **rs6000**, or **sun4**). To use anonymous FTP to retrieve the Xnetlib source, type

```
ftp netlib2.cs.utk.edu
anonymous
your email address
cd xnetlib
binary
get xnetlib.shar.Z
bye
```

Then move the executable file, named **xnetlib** to where you want it. See Section 3.2.3 for more information about building Xnetlib.

If you retrieve an executable, you may also wish to retrieve the shar file **xnetlib3.4.doc.shar**, which contains the Xnetlib man page and quick reference card. These documents are already included in the source code shar file.

3.2.2. System Requirements

Xnetlib can be built on nearly any UNIX system. It runs under the X Window System, version 11, from MIT. It requires release level 4 or greater and the Athena widget libraries as supplied by MIT.

3.2.3. Building Xnetlib

If you retrieved an executable (e.g., **xnetlib.sun4**), install it by uncompressing it and changing the mode to executable, e.g.,

```
uncompress xnetlib.sun4.Z
chmod 755 xnetlib.sun4
```

Xnetlib makes use of the **imake** facility (via **xmkmf**) that comes with standard X Windows distributions. A generic Makefile is provided and can be used if **imake** is not present on your system.

If you retrieved a compressed shar file of the Xnetlib source code, extract the files and build Xnetlib by typing

```
uncompress xnetlib.shar.Z
sh xnetlib.shar
cd xnetlib3.4/src
xmkmf
make
```

After the executable is built, install it by copying the file **xnetlib** to an appropriate directory. There is no application defaults file to install.

With **imake** (**xmkmf**), type

```
xmkmf
make
```

Without **imake** (**xmkmf**), first edit **Makefile.std** to reflect your system characteristics and then type

```
make -f Makefile.std
```

xmkmf should be installed on your system as part of the X distribution. If you get an error on the **xmkmf** command, check your command search path or talk to your system manager. If you are using an IBM/RS6000, you may need to refer to the information in **xnetlib3.4/doc/README.AIXv3**.

For further information on system-wide installation of the Xnetlib client and on customization of the Xnetlib client, see the section in the Netlib Manager's Guide on the Xnetlib client (Section 8).

3.2.4. Xnetlib Man Page and Quick Reference Card

The source code distribution in the **xnetlib3.4/doc** directory include the man page for Xnetlib. To view it, **cd** to this directory and type **nroff -man xnetlib.man**. You may wish to have your system manager install the Xnetlib man page on your system.

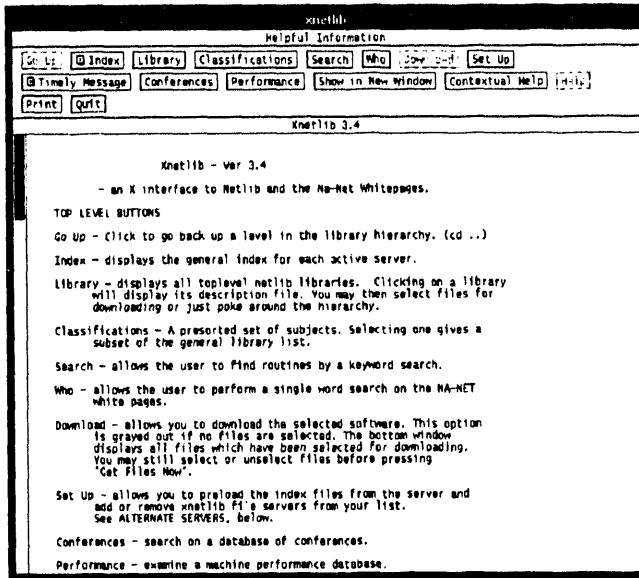


Figure 2: Help Screen

3.2.5. Xnetlib Client Program Operation

Start the Xnetlib client program by typing **xnetlib**. When Xnetlib starts up, it attempts to connect to each server on its list of active hosts. The default active host list is **netlib.ornl.gov**, **netlib.brl.mil** and **softlib.cs.rice.edu**. When these connections have been tested, a window will appear displaying the Xnetlib help screen (Figure 2).

The top row of buttons controls mode changes. In general, Xnetlib reuses the bottom portion of the main window for displays instead of popping up new windows.

Clicking on **Library** displays the top-level listing of libraries available from the Netlib repository (Figure 3). You can think of this top-level listing as a set of UNIX-style directories. Clicking on a library name displays a description of the library's contents. For example, clicking on **lapack** displays the contents of the LAPACK library in library selection mode (Figure 4).

There may be further subdirectories, which are indicated by the folder icon. Files are indicated by the dog-eared page icon. You can select files to be downloaded by clicking on them. Move up the library hierarchy by using the **Go Up** button and move down by clicking on a library name.

When you have selected one or more files to download from the server, click the **Download** button to enter download mode (Figure 5). A list of the files you selected will be displayed. You can alter your selections if desired. Click **Get Files Now** to begin file downloading.

The default directory in which download files are placed is **xnlFiles** in your home directory. You can change the default directory by clicking on the **Download path** button. You can choose to have dependency checking either on or off (default is on) by clicking the **Dependency checking** button. If dependency checking is on, routines

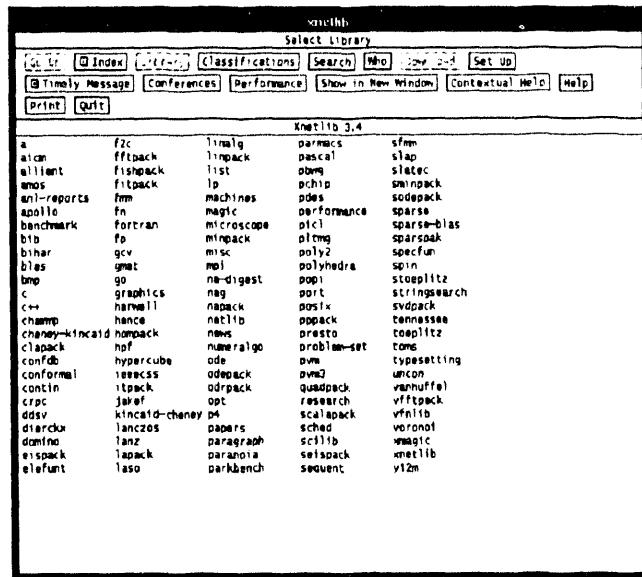


Figure 3: Library List

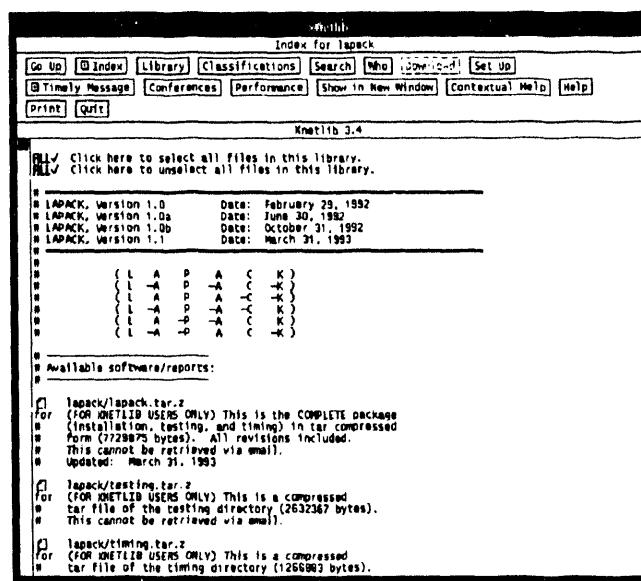


Figure 4: Library Selection Mode

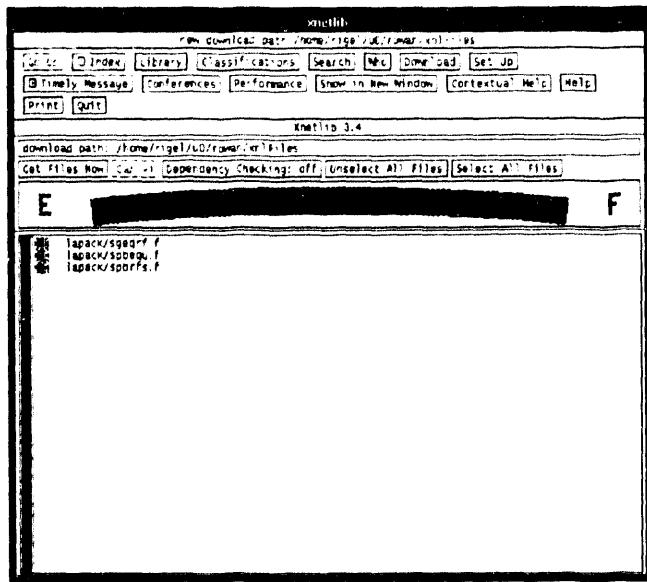


Figure 5: Download Mode

required by the requested file will be appended to the file before it is downloaded.

As you traverse the Netlib tree to examine libraries, the Netlib server at UT/ORNL downloads index files for the libraries to your site. Depending on your network connection, this file transfer may be hardly noticeable or may cause a significant delay. If you prefer to have all the index files loaded at once instead of as you need them, click **Set Up**, then **Press to Check Each Index File**. This will check every index file older than **indexLifetime** (see Section 8.4.1 on Xnetlib X Resources), and will copy from the server any that are missing or changed. To find out how to have several users at your site use one shared collection of index files, instead of several separate collections, see Xnetlib Installation and Customization in the System Manager's Guide (Section 8).

Clicking on **Classifications** lists library topics (Figure 6). The classification is an augmentation of the top level of the GAMS hierarchy. For more information about GAMS, see [2].

Selecting one of these topics displays a subset of the main library list.

Clicking on **Search** takes you into search mode. Buttons for the various types of searches are displayed, along with an explanation of these search types in the main window. The result of a search is a listing of files, from which you can choose files for downloading.

Clicking on **Who** changes to whois mode and allows you to query the NA-NET Whitepages and SIAM membership list for information about members of these groups.

Clicking on **Conferences** takes you into the Conferences Database mode. This mode is an experimental interface to a more general database service based on the relational database model. The Conferences Database contains conference and meeting announcements entered by UT staff and by Xnetlib users.

Clicking on **Performance** takes you into the Performance Database mode. This

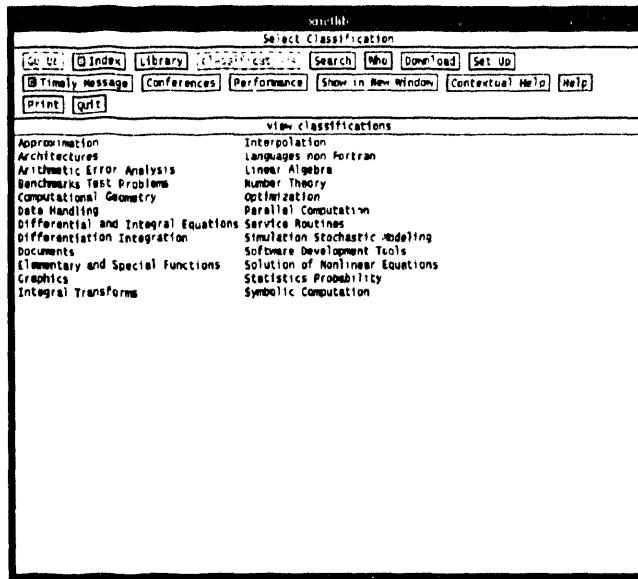


Figure 6: Classifications

mode gives you access to benchmark performance data for a wide range of machines. For more information, see [1].

3.2.6. Command Line Options

Xnetlib understands the normal X Toolkit options (`-display`, `-rv`, etc.) as well as `-help`, and those listed in Figure 7 next to the X resources they affect. See Xnetlib Installation and Customization in the System Manager's Guide (Section 8 in this document) for an more complete explanation of these X Resources.

3.2.7. Application Defaults File

There is no required application defaults file for Xnetlib. However, you can set application defaults for Xnetlib in your `.Xdefaults` file in the usual manner. For example, to set the background color for the Xnetlib client to seashell, include the following line in your `.Xdefaults` file:

```
xnetlib*background: seashell
```

For further information about setting application defaults, see [8].

3.2.8. Support Address

Please send questions, comments, suggestions, or bug reports concerning Xnetlib to the following email address:

`xnetlib@cs.utk.edu`

Option	Resource	
-dw	dontWarp	no gratuitous cursor motion
-nocq	confirmQuit	no confirm on quit
-cq	confirmQuit	confirm on quit
-showdown	showDownloadAnyway	download button always on
-noshowdown	showDownloadAnyway	download button not always on
-showwho	showWhoInfo	show help in who mode
-noshowwho	showWhoInfo	don't show help in who mode
-pub[lic]	publicData	write index files world writable
-nopub[lic]	publicData	don't write index files world writable
-dc	depCheck	dependency checking on
-nadc	depCheck	dependency checking off
-path <path>	xnlDownloadPath	where files are placed
-dp <path>	xnlPath	where index files are cached
-es <string>	entrySep	entry separator in who mode
-pc <string>	printCommand	printing command
-life <number of days>	indexLifetime	how long to keep index files
-email <email address>	email	your email address for our logs
-f[file] <file name>	hostFile	xnetlib server list

Figure 7: Xnetlib options and X resources

3.3. Netlibget, a Command-line TCP/IP Client

Users who do not have X Windows may use **netlibget**, a TCP/IP command-line interface to the Xnetlib servers. This command-line client is included with the Xnetlib release. To build it from the Xnetlib source, type the following:

```
make -f Makefile.std netlibget
```

Running **netlibget** without any options will display the following usage message:

```
usage: netlibget [-v|-q] [-e your_email_address] [-s server[+port]]  
                  [-d|-k|-w] file or key  
-d to get dependent files  
-k for keyword search  
-ko for keyword-or search  
-ka for keyword-and search  
-w for whois search  
-v for verbose; -q for quiet
```

For example, to retrieve **dgeco.f** from **linpack** with dependent routines, type the following:

```
netlibget -d linpack/dgeco.f
```

For a **keyword-or** or **keyword-and** search, the list of keywords should be enclosed in single quotes. For example, to search for files whose descriptions contain the words **condition** and **number**, type the following:

```
netlibget -ka 'condition number'
```

3.4. Anonymous Access

For users who do not need the search capabilities provided by Xnetlib, anonymous access in the form of anonymous FTP and anonymous RCP are provided. For more information about anonymous access, try the following:

```
finger anon@netlib.ornl.gov
```

3.4.1. Instructions for Anonymous FTP

You can use anonymous FTP to retrieve Netlib files. Ftp to `netlib2.cs.utk.edu` and log in as anonymous. Use your complete email address as the password.

The Netlib `ftp` server supports automatic creation of compressed and/or tar'ed versions of files and directories. Just ask for `<filename>.Z` instead of `<filename>`, or `<directory>.tar.Z` instead of each file in the directory, and it gets packaged up on the fly.

3.4.2. Instructions for Anonymous RCP and RSH

You can use anonymous RCP to copy files and anonymous RSH to list directories.

1. To copy a particular file to your system, type

```
rcp anon@netlib.ornl.gov:<filename> <local_filename>
```

where `<filename>` is the filename in the Netlib repository and `<local_filename>` is the local file in which you wish to have the data stored on your system.

For example, to retrieve `sgetrf.f` from `lapack`, type

```
rcp anon@netlib.ornl.gov:lapack/sgetrf.f sgetrf.f
```

2. To list files from a particular directory, type

```
rsh netlib.ornl.gov -l anon ls <directory>
```

(On some machines, the command is `remsh` instead of `rsh`).

Normal `ls` options work; `ls -l` gets the size, etc. You can use metacharacters by putting the argument to `ls` in single quotes.

3.5. Access via Gopher

A `gopher` server is running on a Netlib machine and can be accessed at the address `netlib.ornl.gov` on port 70. The `gopher` URL is `gopher:netlib.ornl.gov:70`. The Netlib libraries are listed as entries in the top-level menu.

4. Future Plans

We plan to expand the Netlib repository from a handful of mirrored servers to a virtual repository consisting of many servers. All of these servers will be accessible from the Xnetlib client program. The services offered will include access to software, documents, and databases, as well as remote execution facilities. A user will not need to know where a particular service is located but will be able to employ both browsing and searching mechanisms to locate and access the desired service.

Future releases of Xnetlib will provide a means whereby newly inaugurated services using the Xnetlib protocol will not require use of a new version of the Xnetlib application. The special-purpose code describing the window layout and functions will be downloaded at run time from the remote service provider. This run-time binding of functionality will greatly enhance the flexibility and adaptability of Xnetlib.

As the software and document base grows in size and complexity, simple syntactic keyword searching mechanisms, such as those provided by WAIS, are likely to prove inadequate. We plan to add new search mechanisms that incorporate semantics and make use of more structured knowledge representations.

Netlib System Manager's Guide

This guide is intended to assist a system administrator in obtaining, installing, and maintaining the files, databases, and client/server software associated with Netlib. It also explains how to set up and maintain the underlying directories and databases.

5. Netlib Repository Setup and Maintenance

The Netlib repository consists of a collection of libraries. A library contains a group of routines that can be retrieved individually. The repository is organized on disk as a UNIX directory tree, hereafter called the "Netlib tree". Assume for purposes of this explanation that the root of the Netlib tree is `/netlib`. Underneath this root are subdirectories for the various libraries. (For example, the EISPACK library is in `/netlib/eispack`). The library subdirectories contain either further subdirectories or files. Although the Netlib tree can be accessed via anonymous FTP and anonymous RCP and can also be browsed via the Xnetlib client, additional searching capabilities were thought to be desirable. Thus, the Netlib tree has been augmented by the inclusion of index, or description, files. The file `/netlib/master/index` contains a listing of all the libraries with descriptions. An index file in the subdirectory for each library lists and describes the library contents. The index files are in a format that is intended to be easily parsed by searching tools. Both the Netlib and the Xnetlib servers use the index files as the basis for their searching mechanisms.

A site may wish to replicate its repository contents, in order to achieve greater reliability and take advantage of load balancing. This replication has been carried out at UT/ORNL, where the Netlib tree is duplicated on two machines, one at UT and the other at ORNL. The copy at UT is the master, and the copy at ORNL is the slave. Whenever someone updates the master copy, that person manually forces propagation of the update to the slave copy using a program that runs the `rdist` command.

Autonomous sites that maintain copies of the same files need a mechanism to keep those copies consistent. Each file has one master copy and some number of slave copies. The site holding the master copy may be different for different files. Netlib has adopted a low overhead repository mirroring scheme, based on checksum files and `ftp`, that keeps slave copies of files consistent with the master copies. For more information on this repository mirroring mechanism, see [7].

5.1. Netlib Index File Format

Netlib is intended to be not just a warehouse but a library, and for that it must have adequate search tools. Each Netlib directory comes with an index file in the format shown below, to promote searchability.

An index file contains one paragraph per file or subdirectory, each separated from the next by an empty line. Each line of the paragraph contains a keyword, then a tab, and finally a corresponding value. If no value is known or applicable, that line is omitted. If all the files in a directory share a value, that line is moved up in the hierarchy to the index file in the parent directory.

An entry for a regular file starts with the keyword *file*, followed by the pathname of the file relative to */netlib*. An entry for a directory starts with the keyword *lib*, followed by the pathname of the directory relative to */netlib*. In schematic form, the keywords are as follows:

```
lib    libraryname
for   {scope}
editor {name <email address>}
master {netlib@...}
      {other fields common across library}
see    related libraries

file   filename.suf
name   {if not same as file}
for    {what problem does it solve?}
alg    {algorithm: what methods does it use?}
by     {names}
ref    {terse citation}
gams   classification codes ("send gams from bib")
size   999 kilobytes
prec
      single      {half}
      double      {full}
      single/double {contains both precisions with #ifdef
                      or other switch}
      real
      complex

rel
      excellent   {widely tested code; firm theory;
                     tractable problem}
      good        {good reputation, but intrinsically
                     hard problem}
      ok          {some counterexamples, but as good as
                     most alternatives}
      weak        {fails without warning; better methods
                     known}

age
      stable      {untouched in years; author still
                     supports it}
      old         {untouched in years; no one supports it}
      research    {believed stable, but still relatively
                     new}
      experimental {known to need polishing}

kind
      experimental {if not the default, a library
                     subroutine}
```

```
command      {standalone program}
data         {input data, measurements, sample
              output, etc.}
text         {documentation}
hdr          {location of function prototype}
obj          {where object code is stored}
overload     {generic name}
#             {miscellaneous comments}
lang         {if not obvious from suffix}

.ada          Ada
.awk          awk
.bas          Basic
.C            C++
.c            C
.f            Fortran77
.ps           PostScript
.r            Ratfor
.tex          TeX (commonly LaTeX)
.bib          BibTeX input
.bbl          BibTeX output
```

Besides the index file, a directory may also have the files

```
readme        general, unstructured information about the
              directory;
dependencies  which symbols are defined in which files;
changes       who did what to which files when;
```

There are two kinds of "test" subdirectories:

```
ex            example "drivers"
chk           self-testing, for use during installation or to check
              compilers, etc.
```

5.2. Repository Replication in Netlib

At UT/ORNL, the Netlib repository contents are replicated on two machines. The master copy resides on `netlib2.cs.utk.edu`, while the slave copy resides on `netlib1.epm.ornl.gov`. Both machines are registered with the Domain Name System as host addresses for the domain name `netlib.ornl.gov`. Theoretically, the load balance for Xnetlib requests between the two machines should be about 50/50. The mail preference is currently set up so that `netlib1` gets most of the email requests. If `netlib1` is down, however, email requests should get sent to `netlib2`.

Updates to the repository must be made on the master copy on `netlib2`. When a staff person makes an update, either by installing a new library or by changing or adding files in an existing library, that person is supposed to run a script that logs

the change and asks him or her whether to propagate the change to **netlib1**. When the person answers yes, the update is propagated by running the **rdist** command. A nightly cron job runs an **rdist** job using a distfile that lists all the libraries that are replicated. The **rdist** job checks for discrepancies, propagates any changes, and notifies the Netlib maintenance staff of any such changes.

6. Email Netlib Server Installation

The Netlib email server software is available from the Netlib repository itself. The software provides mechanisms for processing user email requests, which may involve keyword searching and file retrieval. The Netlib email server is installed at UT/ORNL and a number of other sites for the purpose of distributing numerical software. The software can be adapted easily, however, to distribute other types of software. For example, slightly modified versions are in use to distribute statistical software from **statlib@temper.stat.cmu.edu** and by the TeX User Group to distribute TeX-related software from **tuglib@math.utah.edu**. Appendix A lists Netlib sites and sites known to be running the Netlib email server.

6.1. Acquiring and Installing the Netlib Software

CAUTION: The most common problems with the email interface to Netlib are corrupted mail addresses, network errors, and so on. You should be reasonably expert with email before installing the Netlib email server.

To obtain the Netlib email server software, send the following message to **netlib@ornl.gov**:

```
send netlib from misc
```

Netlib will return a shar file with instructions at the beginning as to how it should be unbundled. This guide assumes that the software is unbundled in a directory named **/netlib/admin**. After unbundling, carry out the following steps:

1. Edit the call to **chdir** in **bin/reply.c** to point to the place where the source code is stored on your system. **/netlib** is assumed for illustration here.
2. Edit **/netlib/admin/LIBS** and **/netlib/master/index** to reflect what you are distributing. Items strictly for local users go in **LIBS.lcl** and **index.lcl**; edit function **groupid()** in **reply.c** to define "local users". If everything is public, **groupid()** can simply return 0. As it currently stands, "local" addresses are those whose machines are in **/netlib/admin/groups/local**; if any machine name in **groups/enemies** is found, processing is aborted.

Suppose you have a library **eispack** that you want to install. Make the directory **/netlib/eispack**; copy in files such as **rg.f**; create a file **index** there; add some lines to **/netlib/admin/LIBS** like

```
eispack => eispack
eispac => eispack
```

(That second line is to allow for misspellings; use your imagination and watch the logfiles for common mistakes.) Be sure that the line

master => master

is in LIBS, so that people can get the main index.

3. Edit the various disclaimers in **/netlib/admin/mess**. You may also wish to add disclaimer files in the source directories.

To activate mail processing:

- * if you run 9th Edition UNIX, put
Pipe to **/netlib/admin/bin/reply**
in **/usr/spool/mail/netlib** and make that file owned by "netlibd",
and similarly put
Pipe to **/netlib/admin/bin/netlibd**
in **/usr/spool/mail/netlibd** and make that file owned by yourself;
- * if you run BSD UNIX, put
netlib: "|/netlib/admin/bin/reply"
netlibd: "|/netlib/admin/bin/netlibd"
into **/usr/lib/aliases** and execute newaliases;
- * else if your system has no equivalent mechanism, try the
daemon in **/netlib/admin/bin/Old-mail-sys**.

The script **admin/bin/netlibd** contains (on line 3) **cd /netlib**, which you may need to change. Because Berkeley's alias facility provides no way to set the userid, you probably should put your name and address in the message so people know who is actually sending the mail.

To try the system out, **echo send index | mail netlib** and expect return mail in a couple minutes. A line should be added to **/netlib/admin/log** and **admin/stderr** should remain empty.

Once the basics are working, you can polish things a bit.

- Set up a nightly process to run **/netlib/admin/bin/mkdir**ectory. Using the **lo** command (with source in **bin/lo.src**) and adapting **bin/mkfile** you can create "dependencies" files in each Netlib directory to represent the relationship between source files there. Not only is this file itself useful to browsers, but by changing the lines in LIBS to the form

eispack => -leispack

you can teach Netlib to respond to a request for **rg** from **eispack** by sending not just **rg.f**, but also **balbak.f**, **hqr2.f**, and so on. To reduce the size of the basic netlib distribution, the **lo** source (all written by David Gay at Bell Labs) is now obtained by **send lo from misc**.

- You should permanently save `/netlib/admin/log` so that bug fixes can be distributed, traffic measurements made, and annual summaries sent to code authors. The format of the log is: date time [address] bytes-sent library/item possibly followed by: L = list, F = find. The [address] is followed by "l" if the address was recognized as local. In contrast, the copy of incoming messages kept in `/tmp/netreq` is for debugging mail headers and monitoring illegal request syntax. Discard when convenient (perhaps by an `rm` in `/etc/rc`) or, if you prefer, comment out the line in `reply.c:handle()` that writes the file.

NOTE: There are a few things in `reply.c`, such as `index/strchr` and `time.h`, that aren't standard across different versions of UNIX. We have tried to deduce which system you run and generate appropriate code. Send reports of failures to Eric Grosse at `ehg@research.att.com`.

7. NA-NET Database Setup and Maintenance

7.1. Acquiring the NA-NET Software

The NA-NET software in use at UT/ORNL is not yet in general release. For more information about a possible future release date, send email to `nanet@na-net.ornl.gov`.

7.2. The NA-NET Program

The NA-NET program handles the following:

- receipt of incoming mail for all NA-NET recipient addresses, including ordinary subscriber addresses (e.g., `na.joe`) and special functions (e.g., `na.help`, `na.join`, `na.whois`, etc.)
- receipt of bounced mail from attempts to send na-digests. These failed attempts are logged.

The NA-NET program is invoked by a modified version of sendmail whenever the modified sendmail receives a message addressed to: `<something>@na-net.ornl.gov`. Other machines know to send mail addressed to `na-net.ornl.gov` to `netlib2.cs.utk.edu` because the latter machine is registered as the "mail exchanger" for `na-net.ornl.gov` with the Internet Domain Name System.

7.3. NA-NET Files

The NA-NET programs and files are in the directory `/usr/local/na-net`. The NA-NET main program is in `/usr/local/na-net/na-net`. The NA-NET program gets the location of all files from a config file that is passed to NA-NET on the UNIX command line when NA-NET is called by sendmail to deliver mail. The config file is in `/usr/local/na-net/config`. It contains the following:

`libdir=/usr/local/na-net/lib`

```
workdir=/var/spool/na-net
prefix=na
helpfile=/usr/local/na-net/help.txt
human=na-net@netlib2.cs.utk.edu
mail_domain=na-net.ornl.gov
errlog1=na.errlog1@na-net.ornl.gov
errlog2=na.errlog2@na-net.ornl.gov
master_file=/usr/local/na-net/nanet_names.master
digest_file=/var/spool/na-net/digest
logfile=/var/log/na-net/na-net.log
digest_ack_file=/usr/local/na-net/digest-ack.txt
delivery_error_file_1=/var/log/na-net/misc-delivery-errors.log
delivery_error_file_2=/var/log/na-net/delivery-errors.log
white_pages_file=/usr/local/na-net/whitepages.database
-----
```

The purpose of these entries is as follows:

libdir - location of subsidiary programs to be called by NA-NET. (currently unused)

workdir - a scratch directory for NA-NET to write temporary files, etc.

prefix - This is the prefix that NA-NET will accept at the beginning of addresses.

NA-NET will remove it from the local part of a recipient address before processing.

The prefix is used to distinguish NA-NET addresses from ordinary local addresses on systems where the NA-NET system shares a mail domain with local users.

Normally the prefix doesn't change. However, you could run multiple NA-NET domains with different prefixes, simply by prefixes, simply by having sendmail call NA-NET with a different config-file for each prefix.

helpfile - location of the text file sent in response to a message addressed to **help**

human - Internet email address (not an NA-NET address) of a human being. This is used for mail to any of the following:

```
na.net@na-net.ornl.gov
na-net@na-net.ornl.gov
nanet@na-net.ornl.gov
postmaster@na-net.ornl.gov
```

It is also used as a reply address on responses to help, join, change, remove, sendlist, whois, join-wp, change-wp, and remove-wp commands, and digest submissions. (But not as the reply address on digests sent to subscribers!)

mail_domain -

Internet mail domain for NA-NET addresses. This is used in error messages and also so that NA-NET will recognize addresses such as `user%na-net.ornl.gov@na-net.ornl.gov` which shouldn't occur, but do.

errlog1 - This address is used for the return address on most responses to NA-NET commands. If a response bounces, the message will be sent back to this address. Normally `errlog1` is set up to feed back into `errlog1@na-net.ornl.gov`, so that we can save the returned message. Unfortunately, if one of these response messages bounces, we can't do much about it; but the logs are sometimes useful to help answer questions or diagnose problems with someone's mail system. (See `delivery_error_file_1`, below).

errlog2 - This should be an Internet email address. NA-NET will use this address as the envelope return-address on all outgoing na-digests. Normally this will point back to `<prefix>.errlog2@<mail_domain>`, which allows NA-NET to process and log bounced mail messages. (Unlike messages sent in response to commands, we can use the bounced digests to let us know whose addresses are no longer valid.) (See `delivery_error_file_2`, below).

master_file - location of the NA-NET subscriber database. See "File Formats" below.

digest_file - This is where incoming digest articles are stored. Each article is separated from the others by a line of the form "From sender date". Two blank lines are added before the "From" line and after each message.

logfile - This is a log of messages sent to NA-NET. Each line is of the form:

date time sender recipient status-code status text...

digest_ack_file - This file contains a message that is sent to people who send mail to the na.digest.

delivery_error_file_1 - This is where bounced messages in response to NA-NET responses get filed.

delivery_error_file_2 - This is where bounced digests get filed.

white_pages_file - location of the NA-NET Whitepages Database.

7.4. File Formats

1. na-digest subscriber database

The subscriber database is an ordinary text file consisting of lines of the form

lastname, firstname (na.key) email-address

These lines are sorted lexicographically by lastname then firstname.

2. Whitepages Database

The format of the Whitepages Database is as follows:

1. Each record is separated from the next by a newline.
2. Each field of the record is separated from the next field by a CR (carriage return, control-M)
3. Any newlines within a field are represented as control-A.

It's difficult to read the file, but if the need arises it can be edited with a text editor.

The fields are, in order:

```
last name
first name
middle name
other name
affiliation
office address
city_state_zip
country
office phone
research
home address
home phone
fax
email address
other
date           -- preferably in rfc 822 format
added by      -- email address of whoever added
                  the record
```

7.5. NA-NET Source Files

The source files are in `/usr/local/src/na-net`. Simply typing `make` should rebuild them.

Notes of any changes are in the file `ChangeLog`. The software is currently stable, but changes are made occasionally to fix (hopefully minor) bugs, to make the software more tolerant, easier to use, or easier to maintain.

7.6. Database Changes, Backups, and Cron Jobs

Database updates are currently performed by sequentially copying the `nanet_names.master` or `whitepages_database` files to "new" files (those ending in `.new`). If the file copy happens successfully, the current database is linked to

`<filename>.old` and the `.new` file is renamed to `nanet.names.master` or `whitepages_database`, respectively. The previous version of the file remains in `<filename>.old`.

The new files are locked during updates so that two concurrent writes cannot happen, though others can read the database while it is being updated. The actual update - replacement of the old database file with the new - is atomic; queries can never see an inconsistent copy of the database file.

In order to prevent race conditions or locking the database for an excessive period of time, some operations may have to be backed out. For example, if someone tries to delete a subscriber record using the first and last names of the subscriber, and more than one subscriber has those names. In this case NA-NET will detect the condition, issue an appropriate message, and discard the new copy of the database (which may already have some records deleted) instead of replacing the old file. (Thus, the presence of a `.new` file does not mean that file contains more current information.)

In addition to the normal system backups, and the `.old` files, there is a cron job (`/usr/local/na-net/rotate.sh`) that gets run once per day that saves an extra copy of the NA-NET databases. Currently these are kept for five days.

A summary of NA-NET activity for a given day is run from cron at 11:59 pm. The summary script is in (`/usr/local/na-net/summary.sh`)

There is currently no mechanism for rotating log files.

7.7. Sending to Digest Subscribers

Digest mailings are accomplished by sending a message to a special NA-NET address known to the digest moderator. Replies to digests are currently sent to the `na.digest` address.

7.8. Surgery

If it is necessary to change NA-NET databases "by hand", the NA-NET system should be suspended. In the directory `/usr/local/na-net` there is a script called `na-net.sh`; this does nothing but exit with a "temporary failure" status. If the normal NA-NET binary is moved aside and the `na-net.sh` file put in its place, sendmail will queue any messages for NA-NET and keep trying every half hour or so. After renaming the `na-net.sh` file, wait a few minutes so that any NA-NET processes can finish up before editing the database. After making whatever changes are necessary, don't forget to rename the real NA-NET program back to `na-net`.

8. Installation and Customization of the Xnetlib Client

8.1. Acquiring the Xnetlib Software

To acquire the Xnetlib client software from Netlib send email to `netlib@ornl.gov` with the line

`send xnetlib.shar from xnetlib`

in the message. Netlib will return the `xnetlib.shar` file by email.

Xnetlib is available by anonymous FTP from `netlib.ornl.gov` in the `xnetlib` directory.

8.2. System Requirements

Xnetlib will build on nearly any UNIX system. It runs under the X Window System, version 11, from MIT. It requires release level 4 or greater and the Athena widget libraries as supplied by MIT.

Xnetlib is known to run on the following systems:

- Convex.
- DECStation running Ultrix 4.1, 4.2, and 4.2A, and DecWindows. (For Xnetlib to compile with the vendor-supplied X Windows libraries, you must have the "Unsupported X11 Components" software subset loaded.) Xnetlib should also build and run with MIT's X11R4 or X11R5, but this configuration has not been tested.
- HP 9000 and MIT's X11R5. (Xnetlib may work with the vendor-supplied X Windows libraries, but this is not recommended.)
- IBM RS/6000 running AIX 3.1, 3.2, and MIT's X11R5. (Xnetlib may work with X11R4, but R5 is preferred.)
- NeXT Dimension and Co-Exist X11R4.
- Sequent Symmetry and MIT's X11R4.
- SGI 4D/25 running IRIX 3.3.3 and X11R4.
- Stardent (Kubota) Titan and vendor-supplied X11R4.
- Sun 3 running SunOS 4.1 and X11R4.
- Sun 4 running SunOS 4.1 or later and X11R4 or X11R5.

8.3. Building Xnetlib

Xnetlib makes use of the `imake` facility (via `xmkmf`) that comes with standard X Windows distributions. A generic Makefile is provided and can be used if `imake` is not present on your system.

After the executable is built, install it by copying the file `xnetlib` to an appropriate directory. There is no application defaults file to install.

With `imake` (`xmkmf`), type

```
xmkmf
make
```

Without `imake` (`xmkmf`), first edit `Makefile.std` to reflect your installation and then type

```
make -f Makefile.std
```

8.4. Customization of Xnetlib

8.4.1. X Resources

publicData - When a user runs Xnetlib, the index files for the libraries are downloaded from the server and cached in a special directory. For sites where many people use Xnetlib, it will save disk space if users share these index files. This can be accomplished by setting the **xnlPath** resource to some commonly writable directory and by setting the **publicData** resource to **True**. The effect of the **publicData** resource is to cause all index files to be saved world writable so they can be updated by anyone. Some sites set these defaults at compile time by setting the fallback resources for these variables. (Look for **fallback_resources** in **main.c**.)

indexLifetime - controls frequency of client to server communication. This resource sets the number of days an index file will be used before Xnetlib checks with the server to find out if the file is out of date. The default value is seven days.

hostFile - sets the name of host file, which contains the lists of servers to contact. The default is **\$HOME/.xnetlibHosts**.

email - specifies the Internet email address of the user. Xnetlib will attempt to guess the email address but will very often be wrong. The email address is recorded in the server's log and is used to inform users of software bugs and updates.

printCommand - sets the format string of the print command. It should contain a "%s" which is replaced when the command is executed by the name of the temporary file used in printing the text. The default is "lpr %s".

showWhoInfo - determines whether instructions for adding your name to the NA-NET Whitepages Database are shown in Who mode. The default is **True**.

entrySep - sets the string to be displayed between entries in Who mode. The default is "____".

xnlDownloadPath - sets the directory where files selected for downloading are to be placed. The default is **\$HOME/xnlFiles**.

depCheck - sets the default value for dependency checking in Download mode.

confirmQuit - If the **confirmQuit** resource is **True**, you will be asked for confirmation before quitting Xnetlib.

dontWarp - turns off automatic cursor positioning if **True**.

okColor, **badColor**, and **cautionColor** - affect the status message window background. The defaults are green, red, and yellow.

dial, **needle**, **ef**, and **gasGauge.background** - affect features of the gas gauge in Download mode. The defaults are DarkViolet, yellow, red, and white.

9. Installing and Running Nlreexecd

Nlreexecd is the service provider daemon for the Xnetlib services. The nlreexecd daemons running at UT/ORNL currently provide access to the Netlib software and document repository, the NA-NET Whitepages Database, the Performance Database, and the Conferences Database. **Nlreexecd** is written to be a general service provider, and can be configured to offer an arbitrary set of services that understand the Xnetlib protocol. A new service can be added easily by providing the code for the function to be called when the service is invoked.

9.1. Acquiring the Nlreexecd Software

Nlreexecd is not yet in general release. Send email to **xnetlib@cs.utk.edu** for more information.

9.2. System Requirements

Nlreexecd should build and run on any UNIX system supporting TCP/IP domain sockets. Nlreexecd does not require any type of X Windows support.

9.3. Building Nlreexecd

There are two flavors of Nlreexecd. The large server distribution provides file transfer and keyword and database lookup and is the basis for the main Xnetlib server. The small server distribution provides only the file transfer service. Site-specific services can easily be added to either server.

Building the small server requires a file called **nlreexecd_small.tar** and otherwise requires a file called **nlreexecd.tar**.

Untar the file in a suitable area. Examine the Makefile and make any site-specific changes you may require. Type **make** to build **nlreexecd**.

Modules used by the small server are also used by the large server. Code that is specific to one or the other within these common modules is differentiated by the **SMALL_SERVER** preprocessor symbol.

9.4. Services and Protocol

All actions performed by nlreexecd are indicated by a unique service name. The service required is passed by name to nlreexecd after a TCP connection is established from the client.

Xnetlib Protocol	
server (nlreexecd)	client (xnetlib or telnet)
listens on TCP port 5555 service specific internal call is made passing client data and socket descriptor	opens TCP connection to server sends newline terminated email address sends newline terminated service name sends newline terminated service specific data

The nlreexecd protocol describes only what occurs up until the service name and data are correctly specified, after that point the connection is "taken over" by that service.

Note that requiring the client to provide any service specific data is a violation of the spirit of separation between the service and the nlreexec layer. It was done entirely to simplify the job of the service module writer. There are several instances in which the service simply ignores this data. In any case, it may not be omitted by the client.

9.4.1. Adding a Service to Nlreexecd

Adding a new service to the server code involves writing the function to be called when your service is invoked and adding the service name in the main module. The following example shows how to add a service called "howdy".

To add the service name in the main module edit the file **nlreexecd.c**.

1. Declare your function (which should return a **char***) where the other service functions are declared.

```
char *howdy();
```

2. Add the name of your service to the **service_list** structure.

```
"howdy",
```

3. Add the name of your function to the **service_call** structure.

```
howdy,
```

Notice that the position of "howdy" in **service_list** should correspond to the position of **howdy** in **service_call**.

4. Create a new file called **howdy.c**. It should look something like this.

```
#include <stdio.h>
#include "nlreexecd.h"

char *howdy(s, service, extra)
int s ; char *service, *extra;
```

```
{  
    swrite(s, "hello world\n");  
    return "howdy ok";  
}
```

The string returned by a service function is written to the log file.

5. Add `howdy.o` to `OBJS` in the Makefile.

Now type `make`.

To test your new service start the server and use telnet to talk to it.

```
csh> telnet localhost 5555  
Trying...  
Connected to localhost  
Escape character is '^]'.  
wade@cs.utk.edu          -- you type this  
howdy                    -- and this  
nothing                  -- and this  
hello world  
Connection closed.  
csh>
```

9.4.2. Reserved Service Names

To avoid service name conflicts it is intended that Xnetlib service name prefixes will be maintained in a central registry. Entities will be provided a service name prefix that they can use to manufacture unique service names. This would work in a manner similar to the Domain Name System except with the most general qualifier at the beginning of the name instead of the end.

Reserved Service Names	
list-services ?	required by all servers
file-tag file-get file-get-dep	used by all servers
who keyword keyword-or keyword-and keyword-lsi keyword-literal keyword-literal-case	used by large server
performance dataserve f2c	reserved prefixes

9.5. Command Line Options

Option	Explanation	Default
<code>-port PORTNUMBER</code>	port at which nlreexecd will listen for client requests	5555
<code>-dir DIRECTORY</code>	effective root directory (argument to chroot)	<code>"/netlib"</code>
<code>-nochroot</code>	Don't do chroot	
<code>-log PATHNAME</code>	pathname for log file	<code>"/usr/local/logs/nlr.log"</code>

9.6. Keyword and Database Lookup Services

9.6.1. Keyword Lookup

The keyword lookup services `keyword`, `keyword-or`, and `keyword-and` use an ndbm keyword database that is manufactured nightly from the netlib index files. This database is a full-text index and includes all words in the netlib index files except those that are excluded by a common word list filter.

The `keyword-literal` and `keyword-literal-case` searches are slower because they do string matching on the index file descriptions themselves.

The keyword searching mechanisms are expected to be changed to use WAIS indexing and searching in the near future.

9.6.2. Latent Semantic Indexing

Latent Semantic Indexing is a method for automatic indexing and retrieval that tries to take advantage of the semantic, or conceptual, content of documents. The particular LSI technique used in Xnetlib at UT/ORNL employs singular-value decomposition (SVD) to take a large matrix of term-document association data (in the case of Xnetlib, the documents are the netlib index files) and construct a "semantic" space wherein terms and documents that are closely associated are placed near one another. LSI tries to tackle the problems of *synonymy* (many ways to refer to the same object) and *polysemy* (more than one meaning for a term), so as to improve the recall and precision of retrieval. In fact, terms that do not actually appear in a document may still end up close to the document, if that is consistent with the major patterns of association in the data. Retrieval is carried out by using the terms in a query to identify a point in the semantic space and by returning documents in the neighborhood of this space. At UT/ORNL, the SVD is done periodically on a matrix constructed from the netlib index files to produce a semantic space for the netlib repository. The `keyword-lsi` service invoked from nlreexecd carries out retrieval. For more information about LSI, see [3].

The particular LSI technique currently used in Xnetlib is patented and proprietary and can be used only with the written permission of Bell Communications Research.

9.6.3. Whois Service

The nlreexecd whois service does a lookup on the NA-NET Whitepages Database, augmented by the SIAM Membership List. Nlreexecd opens the file

`/netlib/whois/whitepages.database` and does a linear search for matches to the user's query.

9.6.4. Performance Database Service

The Performance Database service uses the public-domain RDB relational DBMS developed by Walter Hobbs of Rand Corporation. RDB tables are stored as regular UNIX ASCII files and thus can be manipulated by the normal UNIX utilities. The RDB tables for UT/ORNL Performance Database are stored in the `performance` library in the Netlib tree. The `performance-query`, `performance-or-search`, and `performance-and-search` services invoked from `nlrexecd` carry out searches on the RDB tables.

9.6.5. Conference Database Service

The Conference Database service currently uses the Postgres extended relational database system. Postgres is in the public domain and is available via anonymous FTP from the University of California at Berkeley. The function called by `nlrexecd` for this service is `dataserve`. `dataserve` takes a database-system-independent client request and translates it into the appropriate Postquel language queries which are then executed on the Postgres database.

The conference description files are stored in the `confdb` library in the netlib tree. The filenames for these descriptions are the Postgres Object IDs for the corresponding entries in the Postgres database. Although the descriptions themselves are not stored in Postgres, a full-text index derived from these descriptions is stored in a Postgres relation. Postgres is used to handle searching by dates, keywords, and location. A geographical database stored in a Postgres relation handles geographical aliasing and hierarchical geographical relationships (e.g., retrieving entries for Belgium when asked for those in Europe).

10. Anonymous FTP Server for Netlib

An anonymous FTP server has been installed on `netlib2.cs.utk.edu` to provide anonymous FTP access to the Netlib directory tree. The Netlib anonymous FTP server is based on the Washington University `ftp` daemon, which in turn is based on the 4.3-Reno BSD `ftp` daemon. Minor modifications were necessary to get it to compile in a vanilla SunOS environment. The server supports several useful features including per-site, per-user, and per-directory access control, extensive logging, automatic display of "readme" messages when you `cd` to a particular directory, and automatic creation of compressed and/or tarred versions of files and directories. (Just ask for `<filename>.Z` instead of `filename`, or `<directory>.tar.Z` instead of each file in the directory, and it gets packaged up on-the-fly.)

In "anonymous" mode, `ftp` logs into the "ftp" account and changes its root directory to that account's home directory, so that it becomes impossible to access files outside of that directory. On `netlib2`, the "ftp" account's home directory is `/netlib`.

Configuration files are as follows:

/etc/ftpusers

This file contains a list of users (like "nobody") who are not permitted to log in via ftp.

/usr/local/etc/ftpaccess

This file specifies who can use the ftp server (you can create classes based on where someone logs in from and who they say they are), how many members of each class can log in at once, which files get printed out when you cd to a particular directory, a message to be displayed at login time, whether a class of user can request auto compression or auto-tar, what kinds of things can be logged, who can "upload" files, and where warning messages get mailed.

See "man ftpaccess" for more information.

/usr/local/logs/ftpd.log

This is where transactions get logged.

There are other config files, detailing other features of ftp, but they are not used by the netlib implementation.

For more information, see the man pages for **ftpaccess(5)**, **ftpconversions(5)**, **ftpcount(1)**, **ftpd(8)**, **ftphosts(5)**, **ftpshut(8)**, **ftpwho(1)**, and **xferlog(5)**. (On **netlib2**, these man pages are installed in **/usr/local/man**; make sure your **MANPATH** environment variable contains **/usr/local/man** ahead of **/usr/man**).

Netlib's ftp server is installed in **/usr/local/etc/ftpd**. The Sun-supplied binary is in **/usr/etc/in.ftpd**; but the file **/etc/inetd.conf** has been changed to point to **/usr/local/etc/ftpd**.

A few other programs are also needed to make anonymous FTP work properly - special (statically-linked) versions of **ls**, **compress**, and (GNU) **tar**. These are installed on **netlib2** in **/netlib/bin**. On **netlib2** there is also a dummy **/netlib/etc/passwd** file, which contains dummy entries for **root** and **ftp** - with fake passwords. These are so that the output of **ls -l** can use meaningful user names rather than uid numbers.

The sources are in **/usr/local/src/wu-ftpd-2.1a**. There are some documentation files there that detail how to configure it. For these instructions see **/usr/local/src/wu-ftpd-2.1a/INSTALL** and **/usr/local/src/wu-ftpd-2.1a/NOTES**.

11. Netlib Anonymous RCP Implementation

11.1. "anon" Account

An "anon" account exists on the machines `netlib1.epm.ornl.gov` and `netlib2.cs.utk.edu` for the purpose of allowing anonymous RCP access. Normal logins to this account are disabled by giving it a passwd field of "*". In addition, the shell for this account is a special shell named "anon-shell".

11.2. Invocation of Anonymous RCP

The remote client's `rcp` command invokes the `rcp` command on the netlib server machines via the remote shell service (`rshd`). `rshd` on the Netlib machines has been specially modified to accept the following syntax in `.rhosts` files:

- If the remote-user field is "*", and the remote-machine field is filled in, any remote user at that machine can execute commands.
- If the remote-machine field is "*", and the remote-user field is filled in, any user by that name on any machine can execute commands.
- If both fields are "*", anyone can execute commands for this particular user.

The `.rhosts` file for user `anon` on the netlib machines currently consists of the following line:

* *

The `rshd` program defines the environment variables `REMOTE_HOST` and `REMOTE_USER` for use by programs that it runs.

11.3. anon-shell

"anon-shell" is a very primitive command parser. Basically, it understands c-shell style quoting and globbing (wildcard expansion). When given a command, it splits it up into arguments, expands wildcards on each argument, and then attempts to execute that command. It has a built-in table of commands that it will attempt to run. It will refuse to run any commands that are not in its hard-coded table.

"anon-shell" also logs every command executed, along with the remote user and host, via the syslog facility. Currently it uses `LOG_DAEMON` and `LOG_INFO`. `netlib2`'s syslog currently stores such entries in `/usr/adm/anon-rcp-log`.

"anon-shell" currently has two commands: `rcp` and `ls`. `rcp` is used by the client `rcp` program to retrieve remote files. `ls` can be used to browse directories.

IMPORTANT NOTE: "anon-shell" runs set-uid to root and passes root privilege to any commands that it runs. Thus, it is dangerous to add new commands without going over them carefully.

11.4. Modified **rcp** and **ls** commands

The versions of **rcp** and **ls** installed on **netlib1.epm.ornl.gov** and **netlib2.cs.utk.edu** have been modified as follows:

- Both **rcp** and **ls** immediately do the following:

```
chdir ("/netlib");
chroot ("/netlib");
setuid (getuid());
```

thus limiting their view of the world to everything under **/netlib**, and turning off any special privileges.

- The modified **rcp** can deal with not having an **/etc/services** or **/etc/passwd** file.
- The modified **rcp** has all calls to **mkdir()** and **open(..., O_CREAT, ...)** #ifdefed out and replaced with code that prints “Permission denied”. In general, the file receiving code is disabled, but it will talk protocol with the client **rcp** and return error messages. The file sending code works normally.
- **rcp** and **ls** are statically-linked binaries, since they have no access to system shared libraries.

11.5. Locations of files

The source to the modified **rsh** program is in the directory **/usr/local/src/rsh**.

The specially modified **rshd** program is installed in **/usr/etc/in.rshd**. (the original one is in **/usr/etc/in.rshd.ORIG**)

The sources to the other commands are in **/usr/local/src/anon-rcp** and its subdirectories. The subdirectories **ls**, **shell**, and **rcp** contain the sources to **anon-ls**, **anon-shell**, and **anon-rcp**. These are installed in **anon**, which is currently **/usr/local/homes/anon**.

Logs of transactions are currently kept in **/var/adm/anon-rcp-log**.

A. Netlib Sites

A.1. Sites Mirroring the Netlib Repository

New Jersey	netlib@research.att.com
Tennessee	netlib@ornl.gov
Norway	netlib@nac.no
England	netlib@ukc.ac.uk
Germany	eLib@sc.zib-berlin.de
Taiwan	netlib@nchc.edu.tw
Australia	netlib@draci.cs.uow.edu.au

A.2. Some Sites Using the Netlib Email Server to Distribute Other Types of Software

The software that runs the Netlib email server is available from the Netlib repository. It can be retrieved by sending the message `send netlib from misc to netlib.ornl.gov`. A number of groups have acquired the email distribution software, and a few of these are listed below.

The former `netlib/matlab` directory is now maintained at
`matlib@mathworks.com`.

A collection of statistical software is available from
`statlib@temper.stat.cmu.edu`.

The TeX User Group distributes TeX-related software from
`tuglib@math.utah.edu`.

The symbolic algebra system REDUCE is supported by
`reduce-netlib@rand.org`.

Parallel software and information about parallel processing
is available from
`parlib@hubcap.clemson.edu`

2. References

- [1] Michael W. Berry, Jack J. Dongarra, and Brian H. Larose. PDS: A performance database server. *Scientific Computing*, 1994. (to appear).
- [2] Ronald F. Boisvert, S. E. Howe, and D. K. Kahaner. The Guide to Available Mathematical Software problem classification system. *Comm. Stat. - Simul. Comp.*, 20(4):811-842, 1991.
- [3] S. Deerwester, S. Dumais, G. Furnas, T. Landauer, and R. Harshamn. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391-407, September 1990.
- [4] Jack Dongarra, Tom Rowan, and Reed Wade. Software distribution using XNETLIB. Technical Report CS-93-191, University of Tennessee-Knoxville, June 1993.

- [5] Jack J. Dongarra and Eric Grosse. Distribution of mathematical software via electronic mail. *Communications of the ACM*, 30(5):403-407, May 1987.
- [6] George E. Forsythe, Michael A. Malcolm, and Cleve B. Moler. *Computer Methods for Mathematical Computations*. Prentice Hall, Inc., 1977.
- [7] Eric Grosse. Repository mirroring. (submitted for publication), 1994.
- [8] Adrian Nye and Tim O'Reilly. *X Window System User's Guide*. O'Reilly and Associates, Inc., 1993.

Index

- anonymous FTP, 18
 - acquiring Xnetlib by, 11
 - server, 36
- anonymous RCP, 18
 - implementation, 38
- anonymous RSH, 18
- benchmarks, 3
- Conferences Database, 4, 15
 - server, 36
- contacts, 5
- dependency checking, 13
 - Xnetlib, 31
- downloading, 13
 - xnlDownloadPath, 31
- index file
 - format, 20
- index files, 20
 - Xnetlib, 15
 - indexLifetime, 15, 31
 - shared, 31
- Latent Semantic Indexing, 35
- LSI, 35
- mirroring, 20
- NA-NET, 4
 - email interface, 6
 - file formats, 27
 - files, 25
 - joining, 7
 - name, 6
 - News Digest, 7
 - software, 25
 - source code files, 28
 - Whitepages, 15, 35
 - Whitepages email interface, 8
- netlib
 - email request syntax, 5
- Netlib repository, 1, 3, 20
- nlreexecd, 32
- adding a service to, 33
- Performance Database, 3, 15
 - server, 36
- Postgres, 4, 36
- RDB, 36
- replication, 20, 22
- searching, 15
 - netlibget, 17
- SIAM membership list, 4, 15
- support, 5
- TCP/IP, 10
- user interface
 - anonymous, 18
 - command-line, 17
 - email, 5
 - X Window, 10
- Whitepages, 15
- X Window System, 10
- Xdefaults, 16
- Xnetlib, 10
 - X resources, 31
 - acquiring, 11, 29
 - installation, 12, 30
 - man page, 12
 - operation, 13
 - protocol, 32
 - reference card, 12
 - server, 32
 - services, 32
 - system requirements, 11, 30

ORNL/TM-12680

INTERNAL DISTRIBUTION

1. B.R. Appleton	26-30. R.F. Sincovec
2. T.S. Darland	31-35. R.C. Ward
3. E.F. D'Azevedo	36. D.W. Walker
4. J.M. Donato	37. P.H. Worley
5-9. J.J. Dongarra	38. Central Research Library
10. G.A. Geist	39. ORNL Patent Office
11. S.L. Lee	40. K-25 Applied Technology Li-
12. N. Nachtigal	brary
13. E.G. Ng	41. Y-12 Technical Library
14. C.E. Oliver	42. Laboratory Records Department
15. P.M. Papadopoulos	- RC
16-20. S.A. Raby	43-44. Laboratory Records Department
21-25. T.H. Rowan	

EXTERNAL DISTRIBUTION

45. Cleve Ashcraft, Boeing Computer Services, P.O. Box 24346, M/S 7L-21, Seattle, WA 98124-0346
46. Donald M. Austin, 6196 EECS Bldg., University of Minnesota, 200 Union St., S.E., Minneapolis, MN 55455
47. Robert G. Babb, Oregon Graduate Institute, CSE Department, 19600 N.W. von Neumann Drive, Beaverton, OR 97006-1999
48. Lawrence J. Baker, Exxon Production Research Company, P.O. Box 2189, Houston, TX 77252-2189
49. Jesse L. Barlow, Department of Computer Science, Pennsylvania State University, University Park, PA 16802
50. Edward H. Barsis, Computer Science and Mathematics, P.O. Box 5800, Sandia National Laboratories, Albuquerque, NM 87185
51. Chris Bischof, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
52. Ake Bjorck, Department of Mathematics, Linkoping University, S-581 83 Linkoping, Sweden
53. Jean R. S. Blair, Department of Computer Science, Ayres Hall, University of Tennessee, Knoxville, TN 37996-1301
54. Roger W. Brockett, Pierce Hall, 29 Oxford Street, Harvard University, Cambridge, MA 02138
55. James C. Browne, Department of Computer Science, University of Texas, Austin, TX 78712
- 56-60. Shirley V. Browne, Department of Computer Science, Ayres Hall, University of Tennessee, Knoxville, TN 37996-1301

61. Bill L. Buzbee, Scientific Computing Division, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307
62. Donald A. Calahan, Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109
63. Ian Cavers, Department of Computer Science, University of British Columbia, Vancouver, British Columbia V6T 1W5, Canada
64. Tony Chan, Department of Mathematics, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90024
65. Jagdish Chandra, Army Research Office, P.O. Box 12211, Research Triangle Park, NC 27709
66. Eleanor Chu, Department of Mathematics and Statistics, University of Guelph, Guelph, Ontario, Canada N1G 2W1
67. Melvyn Ciment, National Science Foundation, 1800 G Street N.W., Washington, DC 20550
68. Tom Coleman, Department of Computer Science, Cornell University, Ithaca, NY 14853
69. Paul Concus, Mathematics and Computing, Lawrence Berkeley Laboratory, Berkeley, CA 94720
70. Andy Conn, IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598
71. John M. Conroy, Supercomputer Research Center, 17100 Science Drive, Bowie, MD 20715-4300
72. Jane K. Cullum, IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598
73. George Cybenko, Center for Supercomputing Research and Development, University of Illinois, 104 S. Wright Street, Urbana, IL 61801-2932
74. George J. Davis, Department of Mathematics, Georgia State University, Atlanta, GA 30303
75. Tim A. Davis, Computer and Information Sciences Department, 301 CSE, University of Florida, Gainesville, FL 32611-2024
76. Donald J. Dudziak, Department of Nuclear Engineering, 110B Burlington Engineering Labs, North Carolina State University, Raleigh, NC 27695-7909
77. Iain Duff, Atlas Centre, Rutherford Appleton Laboratory, Didcot, Oxon OX11 0QX, England
78. Patricia Eberlein, Department of Computer Science, SUNY at Buffalo, Buffalo, NY 14260
79. Stanley Eisenstat, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
80. Lars Elden, Department of Mathematics, Linkoping University, 581 83 Linkoping, Sweden
81. Howard C. Elman, Computer Science Department, University of Maryland, College Park, MD 20742

82. Albert M. Erisman, Boeing Computer Services, Engineering Technology Applications, ETA Division, P.O. Box 24346, MS-7L-20 Seattle, WA 98124-0346
83. Geoffrey C. Fox, Northeast Parallel Architectures Center, 111 College Place, Syracuse University, Syracuse, NY 13244-4100
84. Paul O. Frederickson, NASA Ames Research Center, RIACS, M/S T045-1, Moffett Field, CA 94035
85. Fred N. Fritsch, L-316, Computing and Mathematics Research Division, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550
86. Robert E. Funderlic, Department of Computer Science, North Carolina State University, Raleigh, NC 27650
87. K. Gallivan, Computer Science Department, University of Illinois, Urbana, IL 61801
88. Dennis B. Gannon, Computer Science Department, Indiana University, Bloomington, IN 47405
89. Feng Gao, Department of Computer Science, University of British Columbia, Vancouver, British Columbia V6T 1W5, Canada
90. David M. Gay, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
91. C. William Gear, NEC Research Institute, 4 Independence Way, Princeton, NJ 08540
92. W. Morven Gentleman, Division of Electrical Engineering, National Research Council, Building M-50, Room 344, Montreal Road, Ottawa, Ontario, Canada K1A 0R8
93. J. Alan George, Vice President, Academic and Provost, Needles Hall, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
94. John R. Gilbert, Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto CA 94304
95. Gene H. Golub, Department of Computer Science, Stanford University, Stanford, CA 94305
96. Joseph F. Grcar, Division 8245, Sandia National Laboratories, Livermore, CA 94551-0969
- 97-101. Stan C. Green, Department of Computer Science, Ayres Hall, University of Tennessee, Knoxville, TN 37996-1301
102. John Gustafson, Ames Laboratory, Iowa State University, Ames, IA 50011
103. Per Christian Hansen, UCI*C Lyngby, Building 305, Technical University of Denmark, DK-2800 Lyngby, Denmark
104. Richard Hanson, IMSL Inc., 2500 Park West Tower One, 2500 City West Blvd., Houston, TX 77042-3020
105. Michael T. Heath, National Center for Supercomputing Applications, 4157 Beckman Institute, University of Illinois, 405 North Mathews Avenue, Urbana, IL 61801-2300
106. Don E. Heller, Physics and Computer Science Department, Shell Development Co., P.O. Box 481, Houston, TX 77001

107. Nicholas J. Higham, Department of Mathematics, University of Manchester, Grt Manchester, M13 9PL, England
108. Charles J. Holland, Air Force Office of Scientific Research, Building 410, Bolling Air Force Base, Washington, DC 20332
109. Robert E. Huddleston, Computation Department, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550
110. Ilse Ipsen, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
111. Barry Joe, Department of Computer Science, University of Alberta, Edmonton, Alberta T6G 2H1, Canada
112. Lennart Johnson, Thinking Machines Inc., 245 First Street, Cambridge, MA 02142-1214
113. Harry Jordan, Department of Electrical and Computer Engineering, University of Colorado, Boulder, CO 80309
114. Bo Kagstrom, Institute of Information Processing, University of Umea, 5-901 87 Umea, Sweden
115. Malvyn H. Kalos, Cornell Theory Center, Engineering and Theory Center Bldg., Cornell University, Ithaca, NY 14853-3901
116. Hans Kaper, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Bldg. 221, Argonne, IL 60439
117. Linda Kaufman, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
118. Robert J. Kee, Division 8245, Sandia National Laboratories, Livermore, CA 94551-0969
119. Kenneth Kennedy, Department of Computer Science, Rice University, P.O. Box 1892, Houston, TX 77001
120. Thomas Kitchens, Department of Energy, Scientific Computing Staff, Office of Energy Research, ER-7, Office G-236 Germantown, Washington, DC 20585
121. Richard Lau, Office of Naval Research, Code 111MA, 800 Quincy Street, Boston Tower 1, Arlington, VA 22217-5000
122. Alan J. Laub, Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106
123. Robert L. Launer, Army Research Office, P.O. Box 12211, Research Triangle Park, NC 27709
124. Charles Lawson, MS 301-490, Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109
125. Peter D. Lax, Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012
126. James E. Leiss, Rt. 2, Box 142C, Broadway, VA 22815
127. John G. Lewis, Boeing Computer Services, P.O. Box 24346, M/S 7L-21, Seattle, WA 98124-0346

128. Jing Li, IMSL Inc., 2500 Park West Tower One, 2500 City West Blvd., Houston, TX 77042-3020
129. Heather M. Liddell, Center for Parallel Computing, Department of Computer Science and Statistics, Queen Mary College, University of London, Mile End Road, London E1 4NS, England
130. Arno Liegmann, c/o ETH Rechenzentrum, Clausiusstr. 55, CH-8092 Zurich, Switzerland
131. Joseph Liu, Department of Computer Science, York University, 4700 Keele Street, North York, Ontario, Canada M3J 1P3
132. Robert F. Lucas, Supercomputer Research Center, 17100 Science Drive, Bowie, MD 20715-4300
133. Franklin Luk, Department of Computer Science, Amos Eaton Building - #131, Rensselaer Polytechnic Institute, Troy, NY 12180-3590
134. Thomas A. Manteuffel, Department of Mathematics, University of Colorado - Denver, Campus Box 170, P.O. Box 173364, Denver, CO 80217-3364
135. Consuelo Maulino, Universidad Central de Venezuela, Escuela de Computacion, Facultad de Ciencias, Apartado 47002, Caracas 1041-A, Venezuela
136. James McGraw, Lawrence Livermore National Laboratory, L-306, P.O. Box 808, Livermore, CA 94550
137. Paul C. Messina, Mail Code 158-79, California Institute of Technology, 1201 E. California Blvd., Pasadena, CA 91125
138. Cleve Moler, The Mathworks, 325 Linfield Place, Menlo Park, CA 94025
- 139-143. Keith Moore, Department of Computer Science, Ayres Hall, University of Tennessee, Knoxville, TN 37996-1301
144. Neville Moray, Department of Mechanical and Industrial Engineering, University of Illinois, 1206 West Green Street, Urbana, IL 61801
145. Dr. David Nelson, Director of Scientific Computing, ER-7, Applied Mathematical Sciences, Office of Energy Research, U.S. Department of Energy, Washington, DC 20585
146. Dianne P. O'Leary, Computer Science Department, University of Maryland, College Park, MD 20742
147. James M. Ortega, Department of Applied Mathematics, Thornton Hall, University of Virginia, Charlottesville, VA 22901
148. Charles F. Osgood, National Security Agency, Ft. George G. Meade, MD 20755
149. Chris Paige, McGill University, School of Computer Science, McConnell Engineering Building, 3480 University Street, Montreal, Quebec, Canada H3A 2A7
150. Roy P. Pargas, Department of Computer Science, Clemson University, Clemson, SC 29634-1906
151. Beresford N. Parlett, Department of Mathematics, University of California, Berkeley, CA 94720
152. Merrell Patrick, Department of Computer Science, Duke University, Durham, NC 27706

153. Robert J. Plemmons, Departments of Mathematics and Computer Science, Box 7311, Wake Forest University, Winston-Salem, NC 27109
154. James C. T. Pool, Deputy Director, Caltech Concurrent Supercomputing Facility, California Institute of Technology, MS 158-79, Pasadena, CA 91125
155. Jesse Poore, Department of Computer Science, Ayres Hall, University of Tennessee, Knoxville, TN 37996-1301
156. Alex Pothen, Department of Computer Science, Pennsylvania State University, University Park, PA 16802
157. Yuanchang Qi, IBM European Petroleum Application Center, P.O. Box 585, N-4040 Hafsfjord, Norway
158. Giuseppe Radicati, IBM European Center for Scientific and Engineering Computing, via del Giorgione 159, I-00147 Roma, Italy
159. John K. Reid, Numerical Analysis Group, Central Computing Department, Atlas Centre, Rutherford Appleton Laboratory, Didcot, Oxon OX11 0QX, England
160. Werner C. Rheinboldt, Department of Mathematics and Statistics, University of Pittsburgh, Pittsburgh, PA 15260
161. John R. Rice, Computer Science Department, Purdue University, West Lafayette, IN 47907
162. Garry Rodrigue, Numerical Mathematics Group, Lawrence Livermore Laboratory, Livermore, CA 94550
163. Donald J. Rose, Department of Computer Science, Duke University, Durham, NC 27706
164. Edward Rothberg, Department of Computer Science, Stanford University, Stanford, CA 94305
165. Axel Ruhe, Dept. of Computer Science, Chalmers University of Technology, S-41296 Goteborg, Sweden
166. Joel Saltz, ICASE, MS 132C, NASA Langley Research Center, Hampton, VA 23665
167. Ahmed H. Sameh, Center for Supercomputing R&D, 1384 W. Springfield Avenue, University of Illinois, Urbana, IL 61801
168. Michael Saunders, Systems Optimization Laboratory, Operations Research Department, Stanford University, Stanford, CA 94305
169. Robert Schreiber, RIACS, Mail Stop 230-5, NASA Ames Research Center, Moffet Field, CA 94035
170. Martin H. Schultz, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
171. David S. Scott, Intel Scientific Computers, 15201 N.W. Greenbrier Parkway, Beaverton, OR 97006
172. Lawrence F. Shampine, Mathematics Department, Southern Methodist University, Dallas, TX 75275
173. Andy Sherman, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520

174. Kermit Sigmon, Department of Mathematics, University of Florida, Gainesville, FL 32611
175. Horst Simon, Mail Stop T045-1, NASA Ames Research Center, Moffett Field, CA 94035
176. Anthony Skjellum, Lawrence Livermore National Laboratory, 7000 East Ave., L-316, P.O. Box 808 Livermore, CA 94551
177. Danny C. Sorensen, Department of Mathematical Sciences, Rice University, P.O. Box 1892, Houston, TX 77251
178. G. W. Stewart, Computer Science Department, University of Maryland, College Park, MD 20742
179. Paul N. Swartztrauber, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307
180. Philippe Toint, Dept. of Mathematics, University of Namur, FUNOP, 61 rue de Bruxelles, B-Namur, Belgium
181. Bernard Tourancheau, LIP, ENS-Lyon, 69364 Lyon cedex 07, France
182. Hank Van der Vorst, Dept. of Techn. Mathematics and Computer Science, Delft University of Technology, P.O. Box 356, NL-2600 AJ Delft, The Netherlands
183. Charles Van Loan, Department of Computer Science, Cornell University, Ithaca, NY 14853
184. Jim M. Varah, Centre for Integrated Computer Systems Research, University of British Columbia, Office 2053-2324 Main Mall, Vancouver, British Columbia V6T 1W5, Canada
185. Udaya B. Vernulapati, Dept. of Computer Science, University of Central Florida, Orlando, FL 32816-0362
186. Robert G. Voigt, ICASE, MS 132-C, NASA Langley Research Center, Hampton, VA 23665
187. Phuong Vu, Cray Research, Inc., 19607 Franz Rd., Houston, TX 77084
- 188-192. Reed C. Wade, Department of Computer Science, Ayres Hall, University of Tennessee, Knoxville, TN 37996-1301
193. Daniel D. Warner, Department of Mathematical Sciences, O-104 Martin Hall, Clemson University, Clemson, SC 29631
194. Robert P. Weaver 1555 Rockmont Circle Boulder, CO 80303
195. Mary F. Wheeler, Rice University, Department of Mathematical Sciences, P.O. Box 1892, Houston, TX 77251
196. Andrew B. White, Computing Division, Los Alamos National Laboratory, P.O. Box 1663, MS-265, Los Alamos, NM 87545
197. Margaret Wright, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
198. David Young, University of Texas, Center for Numerical Analysis, RLM 13.150, Austin, TX 78731

199. Earl Zmijewski, Department of Computer Science, University of California, Santa Barbara, CA 93106
200. Office of Assistant Manager for Energy Research and Development, U.S. Department of Energy, Oak Ridge Operations Office, P.O. Box 2001 Oak Ridge, TN 37831-8600
- 201-202. Office of Scientific & Technical Information, P.O. Box 62, Oak Ridge, TN 37831

DATA
MATERIAL

CHANGE/

FILED
DATE

