

Computer Networking For Michigan Economists

Charles Taragin*

August 18, 2005

Introduction

While many doctoral students have used computers extensively, few have taken advantage of the assorted capabilities that modern computer networks allow. In particular, many students believe that computer networking *only* enables users to surf the web and write email. In fact, computer networking can be used to

- Access your data from **any** computer with an internet connection.
- Use proprietary software that isn't installed on your personal computer.
- Backup and restore data.
- Maintain a personal web page.

Part 1 of this tutorial will describe the fundamentals of networking, including what a network is, and how it works. Part 2 will discuss how to access your network accounts. Parts 3 and 4 will discuss some advanced networking topics including an introduction to the Institutional File System (IFS)¹, restoring backup files, an introduction to the X windows system², and creating your personal web page.

*I want to thank Taryn Dinkelman, Ann Ferris, Andreas Pape, and Amy Gass-Kandilov for helping me transform my ramblings into something comprehensible.

¹The IFS is the network file system where all ITCS affiliated computer accounts are housed.

²In a nutshell, the X windows system allows you to run the graphical front-end of any server-based program on your local machine.

1 Networking Fundamentals

A **computer network** is a system for communication among two or more computers. One popular network configuration is known as the **client-server architecture**, in which each computer or process (program) on the network is either a client or a server. Servers are powerful computers or processes dedicated to managing disk drives (file servers), printers (print servers), or network traffic (network servers). Clients are PCs or workstations on which users run applications. Clients rely on servers for resources, such as files, devices (like cd burners and scanners), and even processing power. In other words, servers are programs that provide some sort of service to a network of computers, while clients are programs on the network that make use of that service³. Most institutions that have invested extensively in their computer networks




Figure 1: The Client-Server Architecture

have established procedures allowing only select individuals access to their networks. At the University of Michigan (and indeed most academic environments) access is managed through an **account** system where students and faculty are assigned a login name and allowed to choose a password meant to verify their identity. In addition to the computer network maintained by the University, colleges within the U of M (like LSA) and even departments with special needs also maintain their own private networks. Access to these networks is typically restricted to members of those colleges or departments, and is enforced by creating additional accounts with their own login names and passwords. For convenience, the administrators of these smaller networks use the same login name as the U of M's, which frequently confuses students that are uncomfortable with computer networking.

1.1 Your ITCS Account

As doctoral students in the Economics program at the U of M, we are endowed with a single computer account, which is maintained by the University of Michigan's

³It is possible for the same computer to be both server and client.

Information Technology Central Services (ITCS). Your ITCS account includes the following features:

- an email account.
- 1 gigabyte of network storage space⁴.
- a personal webpage.
- access to university-wide computational servers. These servers have many useful programs, including SAS, stata, ArcView, StatTransfer, SPSS, and R⁵.

In short, you will want to use this account to save files related to your web page, download electronic journal articles from Jstor, keep electronic copies of your instructors' notes, store and execute code for your course assignments, and check email.

1.2 Economics Computational Servers

In addition to the services provided by ITCS, the Economics department maintains computational servers exclusively for Economics graduate students and instructors. Like the university-wide computational servers, the economics servers are computational powerhouses, containing a wide selection of programs including SAS, stata, matlab.

2 Accessing Your Accounts

Once your computer account has been activated, four things are necessary to connect to the desired account

- A computer with an internet connection.
- the **host name** or **ip address** of the computer that you wish to connect to. A computer's host name is the unique name that identifies a computer on a network. Below is a table listing host names that grant access to various campus-wide computer resources.

⁴Additional storage space may be purchased from ITCS at a rate of \$12 per gigabyte per year. See <http://www.itd.umich.edu> for more information.

⁵A complete listing of the software available on this server may be found at <http://www.umich.edu/~gpcc/scs/>.

Host Name	Manager	Purpose	Selected Applications
login.itd.umich.edu	ITCS	Email	rsync, pine, L ^A T _E X, Acrobat Reader
sftp.itd.umich.edu	ITCS	FTP	scp, rsync L ^A T _E X, Acrobat Reader
scs.itd.umich.edu	ITCS	Computational	stata, matlab, SAS, stattransfer
econstat1.econ.lsa.umich.edu	Econ	Computational	stata, matlab, SAS

It is important to note that email may only be checked on `login.itd.umich.edu`, files may be uploaded and downloaded only through `sftp.itd.umich.edu`⁶, and statistical programs may only be executed on `scs.itd.umich.edu` and `econstat1.econ.lsa.umich.edu`.

- The login name and password that grant you access to **that** account⁷.
- Software that implements a protocol⁸ used to securely connect the client to the server.

Basically, there are two important protocols that you should be aware of:

- *ssh*, or the secure shell protocol allows you to access and run programs remotely (on a computer that you aren't sitting in front of). This protocol is used by students to check email, conduct basic account maintenance like organizing your files, execute statistical packages, and word processing⁹.
- *sftp*, or the secure file transfer protocol allows you to transfer files between your local computer (the computer that you are sitting in front of) and a remote computer. This extremely useful protocol means that it isn't necessary to carry around USB keys, cd-roms or (ugh) floppy disks. So long as the above conditions are met, you will always be able to access your files.

Like *ssh*, *sftp* is the successor to the *ftp* protocol, but with the added encryption advantage.

⁶this isn't strictly correct; you may use `rsync`, a utility that allows you to synchronize the contents of any collection of folders, on `login.itd.umich.edu`.

⁷I would advise against setting different passwords for your accounts, largely because you will invariably forget which server you are trying to log into and enter the wrong password. Besides, the propagation of additional passwords is bad. Just choose one REALLY difficult password and change it on all your accounts about once a year.

⁸A protocol is a convention or standard that controls or enables the connection, communication, and data transfer between two computing devices.

⁹For the curious, *ssh* is the successor to the protocol *telnet*. *ssh*'s biggest advantage over *telnet* is that it encrypts data sent across a network, which means that it would take a hacker who intercepted your login name and password over a decade on a personal computer to decipher your information.

Both of these protocols should be automatically installed on Unix-style systems like Solaris, Mac OSX, and Linux. On Windows machines, go to <http://www.itd.umich.edu/bluedisc/updates.html>¹⁰ and download “SSH Secure Shell 3.x.x”, or obtain a copy of U of M’s Blue Disc at the campus computer store for \$15.

2.0.1 Example: Checking Email (Mac)

Double-click the “Terminal” icon under Under Applications→Utilities . At the command line type `ssh <username>@login.itd.umich.edu` to log onto `login.itd.umich.edu`, where `<username>` is your ITCS username.

At this point, `ssh` might ask you a question regarding the authenticity of the host that you are trying to connect to. This is a security feature meant to protect you from “Man in The Middle” internet attacks. If you are sure that you typed in the host name correctly, answer “yes”.

Enter your ITCS password when prompted. Assuming that you entered your password correctly, you should be logged on to one of Michian’s servers. To check email, execute the *nix-based mail client `pine` by typing `pine` at the command line.

2.0.2 Example: Checking Email (Windows)

Execute “SSH Secure Shell” by clicking on the appropriate icon under Start→Program Files . Click “Quick Connect”, and in the “Host Name” field type `login.itd.umich.edu`, in the “User Name” field type your username, and in the “Port Number” field type 22.

At this point, `ssh` might ask you a question regarding the authenticity of the host that you are trying to connect to. This is a security feature meant to protect you from “Man in The Middle” internet attacks. If you are sure that you typed in the host name correctly, answer “yes”.

Assuming that you entered your password correctly, you should be logged on to one of Michian’s servers. To check email, execute the *nix-based mail client `pine` by typing `pine` at the command line.

2.1 Accessing email and network space via the World Wide Web

ITCS has created a series of web-based applications allowing you to access both your email and your network (IFS) space through your favorite web browser. To check email, simply go to <http://mail.umich.edu> and enter you ITCS username and password. Be sure to disable your web browser’s popup blocker for this site; if you

¹⁰This page contains many free goodies available for download, including free antiviral software, a must for any Windows user. Login is required.

don't you will not be able to compose emails. Additional documentation, including how to manage spam, is available at

<http://www.itd.umich.edu/itcsdocs/s4297/>

To access your IFS via the web, go to <http://mfile.umich.edu> and log on using your ITCS username and password. Use the menu on the left hand side to upload, download, or manage files. Excellent documentation for mfile is available at

<http://www.itd.umich.edu/itcsdocs/s4311/> .

2.1.1 Example: Uploading a file to your IFS space

Open up your favorite text editor and type “I am the very model of a modern major general” . Save the document as a text file (.txt suffix) on your computer's Desktop. Go to <http://mfile.umich.edu> and log in using ITCS username and password.

If everything worked properly, you should be looking at a list of all the files located on your IFS space. Click “Upload Files” on the left-hand menu. Click ”Browse”, go to the Desktop folder, and select the text file you just created. Click the “Upload File(s)” button. Congratulations! you have just finished uploading your first file to your IFS space

To delete the file you just uploaded, check the box next to that file, and select “Delete Selected Items” from the left-hand menu.

3 The Institutional File System (IFS)

This section introduces IFS¹¹. Having some understanding of how IFS works is important for two reasons. First, if you are storing data on your network space, knowing a bit about how IFS works will make it easier to manage your data. Second, understanding the IFS makes it extremely easy to securely share files with selected members of the University community without going through the bother of setting up a secure web page.

IFS is U-M's implementation of the AFS file system. What is the AFS file system, you ask? AFS is a distributed networked file system, which just means that it is a system for organizing files that spans multiple computers. AFS uses the Kerberos protocol for authentication, and implements access control lists on directories for users and groups. More on access control lists later.

¹¹This section revisits material introduced in Section 1, albeit at a more complicated level. Much of this section was shamelessly borrowed from <http://www.itd.umich.edu/itcsdocs/r1070/#connect> .

Like most file systems, IFS is organized by directories. At the AFS root level directory, you can connect to implementations of AFS worldwide. The first layer of directories in the AFS file space contains what are called cells. A cell is an administrative domain and is generally controlled by a company, university, department, or other large group of users. ITCS manages the `umich.edu` cell for the use of the U-M community. There are other cells on campus. For example, Michigan's Engineering Department manages the `engin.umich.edu` cell

Within the `umich.edu` cell are five directories: `class`, `group`, `system`, `um`, and `user` (see Figure 2). Within the `user` directory, there are 26 folders, one for each letter of

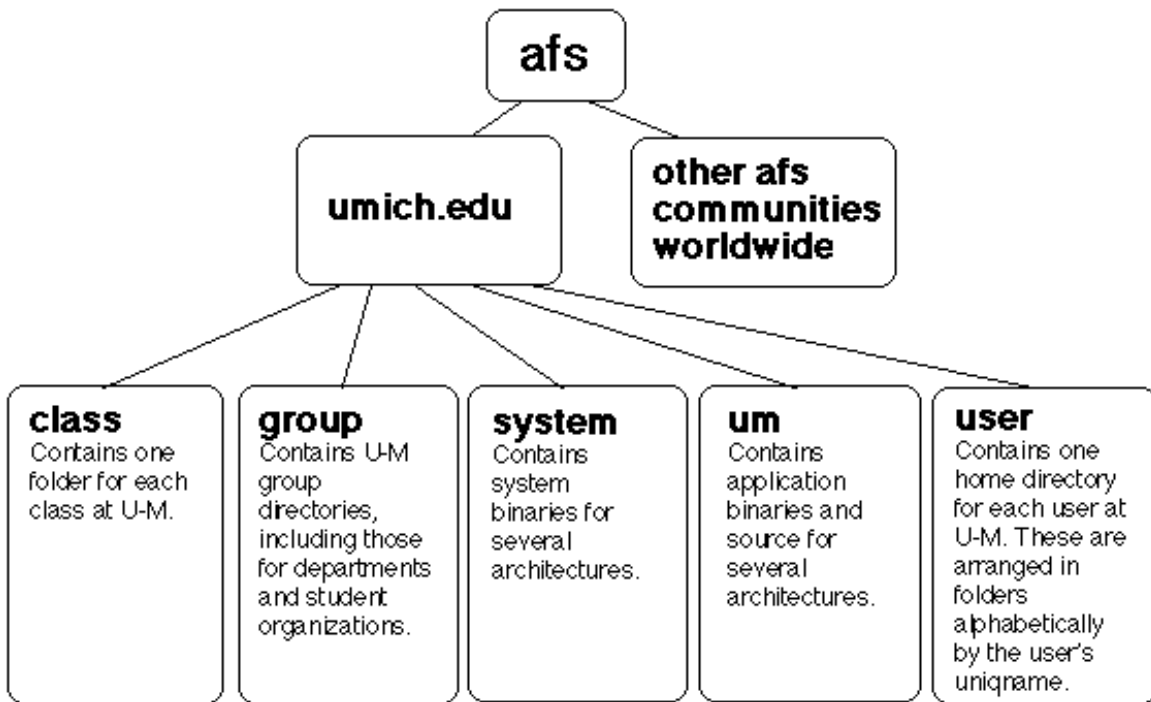


Figure 2: IFS at U of M

the alphabet. Within each of those folders are 26 folders, again one for each letter of the alphabet. Individual user home directories are inside, filed by the first two letters of the person's unickname. For example, to find the home directory of someone whose unickname is `bjensen`, you would go to the `user` directory, then open the `b` directory, then the `j` directory.

People generally talk about locations in IFS in terms of pathnames. A pathname is basically the path one takes (or the path your computer takes) to get to the directory or folder you want. For example, the pathname to `bjensen`'s ITCS home directory would be `/afs/umich.edu/user/b/j/bjensen`, while the pathname to `bjensen`'s Economics home directory would be `/afs/lsa.umich.edu/user/b/j/bjensen`.

3.1 Access Control

You can control whether other people can see the files and documents in the folders inside your IFS home directory, whether they can make changes to them, and more through the use of Access Control Lists (ACLs). An ACL is a list of uniqnames and/or protection groups to which access rights have been assigned. (A protection group – or pts group – is similar to an e-mail group except that it is a list of uniqnames rather than a list of e-mail addresses; it is used to assign permissions to a group of people.) ACLs are set for folders.

For example, you might create a folder in your IFS home directory that you want to use for a group project. You could then set ACLs for that folder to allow only your group members to see what is inside it and make changes. Relatively clear instructions for manipulating ACLs are available from

<http://www.itd.umich.edu/itcsdocs/s4111/>.

3.2 Restoring Automated Backups

One important feature available on your ITCS account is an automatic file archival system. This system is incredibly useful if you accidentally delete files, or just want to go back to an older version of a file.

Your ITCS account archives files both daily and weekly, with the weekly backups stored for 6 months. You may access the daily backups through the `.oldfiles` directory in your ITCS account. Restoring the weekly backups costs \$10, and can only be done by contacting Information Services at `ifs-restore@umich.edu`. Detailed information regarding this service is available at

<http://www.itd.umich.edu/itcsdocs/s4110/> .

4 Advanced Topics

4.1 The X Window System

The X Window System (commonly X11 or X) is the standard graphical interface on Unix and Unix-like operating systems. X provides the basic framework for a graphical user interface (GUI) environment: drawing and moving windows on the screen and interacting with a mouse and keyboard. X features network transparency: the machine where application programs (the clients) run need not be the user's local machine (the display server). X's usage of the terms "client" and "server" is the reverse of what people often expect, in that the "server" is the user's local display rather than the remote machine.

In other words, the X Window system allows you to execute programs with graphical user interfaces (GUIs) remotely.

X is freely available for Unix, Linux and Mac OSX. A one-year license for Exceed, the Windows implementation of X, may be purchased at the university computer store for \$35.

4.1.1 Example: Running Matlab from scs on your Macintosh (with OSX Panther or better)

Suppose you want to analyze some data using Matlab, but have spent all your cash on your brand new Macintosh. Normally, you would be s*** out of luck, but happily, you're a member of the fabulous U of M Economics Department with access to the scs server and its copy of Matlab.

You could just use `ssh` to log onto scs and type `matlab` at the command prompt, but if you did, you wouldn't get Matlab's cool graphical interface. bummer.

Here is where X comes in to play. First, start X up on your computer by clicking on Macintosh HD → Applications → Utilities → X11 (or use Finder)¹². Once X is running, open a terminal window by clicking on the Applications menu and selecting Terminal. At the prompt, type

```
ssh -X <uniqname>@scs.itd.umich.edu
```

This command establishes a secure connection between your computer and scs. the `-X` switch tells `ssh` that you want it to render graphics sent from scs onto your Macintosh. Now, type `matlab` at the command prompt, and ... voila! Matlab's graphical interface should appear. Remember, although it may seem that you're working on the computer in front of you, really your not. This means that in order for Matlab to access any of your data, it must be located somewhere accessible from your ITCS account.

Also, note that X11 isn't just restricted to Matlab (or to scs for that matter), but can be used with *any* program on *any* accessible server. Indeed, much of this document was written on the Angell Hall Macs while accessing `login.itd.umich.edu`'s copy of Emacs.

Here is where you can see the true power of computer networking. You can sit down in front of any computer with an internet connection and access whatever files or programs that are available through your ITCS accounts. That being said, please heed the following

WARNING: Running programs interactively with X uses more server resources than running programs in batch (noninteractively). Please make every attempt to conserve server resources by running programs in batch mode when possible. See footnote ?? for further details.

¹²If X11 isn't installed on your computer, you can download it from <http://www.apple.com/support/downloads/x11formacosx.html>.

4.2 Creating Your Personal Web Page

A personal web page is a great place to post your *Curriculum Vitae* (CV), research papers, course handouts, or pictures.

To set up your web page, follow the instructions at

<http://www.umich.edu/~umweb/how-to/homepage.html> .

Pay special attention to the links at the bottom page, particularly the link relating to CGI scripts. CGI scripts are programs that add functionality to websites. For example, if you want to restrict access to your web site, you could do so by downloading the `htaccess` CGI script.

Modern word processing programs have made it fairly easy to build your own web page, even with little or no training in Hypertext Markup Language (HTML). Perhaps the easiest way to generate a web page is using an office suite like Microsoft Word¹³. Open a word document and type whatever you like. You can even use Word to add hyper-links (the blue underlined text in a browser) and reference files. When you've finished, save the document as an HTML file and upload it to the appropriate directory in your ITCS account.¹⁴

¹³Don't go out and buy MS Word if you don't have it. Instead, download the excellent OpenOffice suite from <http://www.openoffice.org>. It is free, has most of the functionality of MS Office, and can read and write documents in the various Office formats. Did I mention that it is free?

¹⁴See the above link for more detailed information. It is important to note that the default file that is read when your web page is accessed is called `index.html`. Hence, your first order of business should be to create this file.