An Internet Gopher for the University of Florida*

Bradley C. Spatz
Computer & Information Sciences
bcs@ufl.edu

August 31, 1992

The Internet Gopher is a distributed document search and retrieval system developed by the University of Minnesota and used at a growing number of sites around the world. The term "Gopher" refers to both a protocol and a suite of public domain software applications that implement the protocol.

Gopher supports a simple, efficient method for making documents and other network resources available to a wide audience. I believe that Gopher has tremendous potential for UF and that we (the University of Florida) should consider implementing it at our university.

This document is directed to anyone with an interest in new networked information technologies and, in particular, how these may affect the function of the University. The document is divided into three main sections:

- The first section describes Gopher generically as a networked application and as an information server.
- The next section explains why an Internet Gopher should be implemented at the University of Florida.
- The last section discusses the current state of Gopher development at UF.

1 The Internet Gopher

1.1 What's an Internet Gopher?

The University of Minnesota's Microcomputer and Workstation Networks Center describes Gopher in this way:

> *The Internet Gopher uses a simple client/server protocol that can be used to publish and search for information held on a distributed network of hosts. Gopher clients have a seamless view of the information in the Gopher world even though the information is distributed over many different hosts. Clients can either navigate through a hierarchy of directories and documents – or – ask an index server to return a list of all documents that contain one or more words. Since the index server does full-text searches, every word in every document is a keyword.*

That's the thumbnail sketch. To get a better sense of the way Gopher works, let's first consider how the parts fit together physically and then how the application functions for the user.

1.2 Gopher as a Networked Application

The Gopher server can be an IBM (CMS) or VAX (VMS) mainframe, a Unix workstation, a NeXT workstation, a PC-compatible microcomputer, or a Macintosh. The server is connected in some way to the Internet.

The Gopher server software application runs on this machine and manages access to various resources that are located either on the server itself or on other machines[1] to which the server can connect.

*Portions of this document are taken from *An Internet Gopher for the University of Victoria*.

[1] These "other machines" comprise what the University of Minnesota documentation refers to as "a distributed network of hosts."

The user runs a Gopher client on an IBM (CMS/MVS) or VAX (VMS) mainframe, Unix or NeXT workstation, PC, or Mac. The client software establishes a connection with the server and provides an interface to the host application. Any Gopher client can access any Gopher server from anywhere on the Internet.

In its simplest form, an Internet Gopher could look like this:

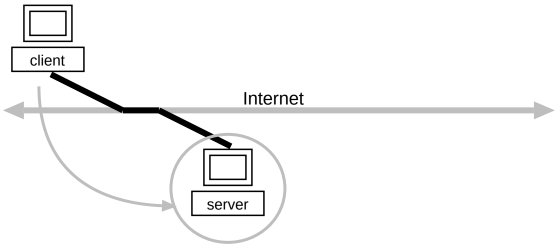

A Gopher server also can provide links to other Gopher servers. In this configuration, the server to which the user's client connects is referred to as the top-level server. The diagram below shows the client connecting to a top-level server(1) and accessing a resource – via an intermediate machine – which is located on a third server(2).

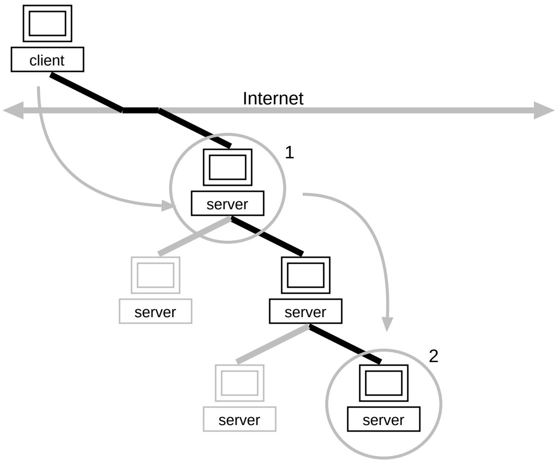

Other Gopher servers may be located within a particular top-level server's domain (for example, distributed across a university campus), or they may be in other domains. A Gopher server(1) also can include links to other non-Gopher hosts – for instance, an Internet-accessible weather service(3).

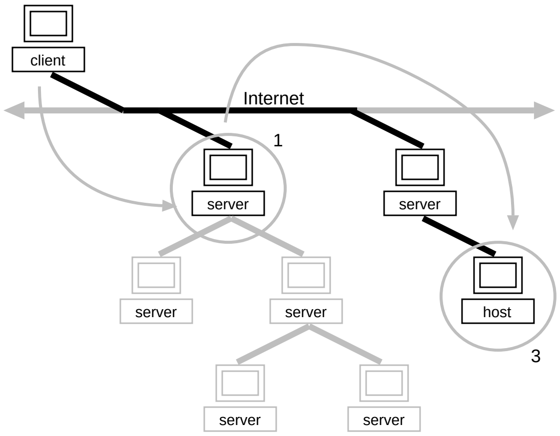

1.3 Gopher as an Information Server

When the user connects to a Gopher server, the client software displays a directory listing as shown below.

```
         Internet Gopher Information Client v1.02

            Root gopher server: gopher.ufl.edu

 -->  1.  About this Gopher.
      2.  What's new in this Gopher.
      3.  Campus Events and Activities/
      4.  Computing and Networks/
      5.  Library Resources/
      6.  Employee Information/
      7.  Student Information/
      8.  Telephone and E-mail Directory <CSO>
      9.  Other Information Services/


Press ? for Help, q to Quit, u to go up a menu    Page: 1/1
```

In *The Internet Gopher Protocol*[1], 11 types of directory list items are defined.

- **0 Item is a file.**
- **1 Item is a directory.**
- 2 Item is a CSO (qi) phone-book server.
- 3 Error.
- 4 Item is a BinHexed Macintosh file.
- 5 Item is DOS binary archive of some sort.
- 6 Item is a UNIX uuencoded file.
- **7 Item is an index-search server.**
- **8 Item points to a text-based telnet session.**
- 9 Item is a binary file.
- \+ Item is a redundant server (same information as the previous server).

More recently, new data types for sounds, images, and complex data types have been proposed and implemented in prototype clients. However, for the purposes of our discussion, only those items printed in bold are described further.

1.3.1 Type 0: A File

Files (or documents) are plain ASCII text files that can be produced with most word processors, database report generators, etc.

To read a document, the user selects the item. Gopher then displays the full text of that item on her screen. A simple text-browsing facility allows her to page backward and forward through the file, perform string searches, etc. When she quits the text browser, she is returned to the directory list from which she made her selection.

1.3.2 Type 1: A Directory

The "root directory" list is the top level of a heirarchy of subdirectories. By selecting a subdirectory name, the user moves "down" to a more narrowly focused range of choices.

For example, the root directory illustrated on the previous page has an item called "Other Information Services." Selecting this leads to a new subdirectory listing:

```
         Internet Gopher Information Client v1.02

                 Other Information Services

      1.  Archie (ftp index) Server <TEL>
      2.  FTP Searches/
      3.  Gopher Servers Worldwide/
      4.  Gopher Servers at UF/
      5.  Library Catalog Services Worldwide/
      6.  Miscellaneous Information Servers Worldwide/
      7.  NERDC Gopher Server/
      8.  Phone and Email Directories Worldwide/
-->   9.  Weather/



Press ? for Help, q to Quit, u to go up a menu    Page: 1/1
```

Selecting "Weather" then brings up a another subdirectory that lists items specific to weather services. And so on.

With the "Up" command, the user can retrace his steps back through the heirarchy and make new selections at higher levels – or he can "Quit" from any level of the heirarchy.

1.3.3 Type 7: An Index-Search Server

An index-search server maintains a full-text index for documents within a given domain. If the user selects a full-text search, he is prompted for search strings, Boolean operators, etc. The server then returns a virtual directory listing of all the documents that match the request (that is, a listing displayed in the same format as any other Gopher directory that may be selected from or reviewed using the same comands).

1.3.4 Type 8: Text-Based TELNET session

The Gopher server can transparently establish telnet sessions with another hosts. In practical terms, this means that:

- *Users do not have to know anything about TELNET in order to access Internet resources.* A user could select "Archie," for example, from the "Other Information Services" directory list, and be immediately connected to an Archie[2] server. When the user selects "Quit" from the Archie system, she would be returned to the "Other Information Services" directory list.

- *Documents and other resources may be located on servers distributed across the Internet.* As a user navigates down through a directory heirarchy, he also may be navigating from one server to another. This is, of course, happening behind the scenes. As far as the user is concerned, all of the documents and resources that can be accessed from a top-level Gopher server are – at least logically – residing on that machine. And this means that...

- *Valuable resources do not have to be duplicated.* Let's say that University of South Florida has developed a particularly good "Internet gateway" for US government document archives. Rather than reinventing the wheel, a server could simply have a directory item point to a TELNET session that connects the user directly to the USF archive application.

Figure 1 illustrates how resources on a remote Gopher system can be linked in such a way that the user sees them as extensions of her "local" system.

[2] Archie is a system to query Internet software archives.

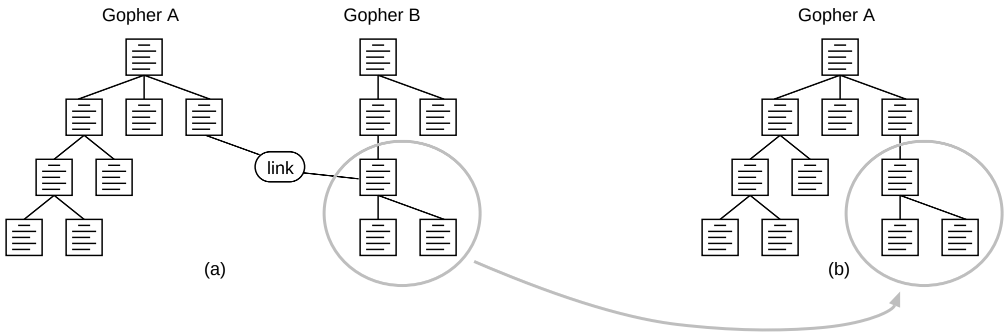

Figure 1: What Gopher sees (a). What the user sees (b).

1.4 Where Are Internet Gophers Being Used?

The University of Minnesota's Gopher has a subdirectory called "Other Gophers" that lists dozens of servers around the world. When Farhad Anklesaria, from the University of Minnesota's Microcomputer Research Group, was asked how many production Gopher servers there are, he said:

> *Glancing at our directory of links to "Other Gophers," a quick count of the running Gophers that want to be accessible is around 100. This only counts primary (top-level) Gopher servers, and does not count secondary or index servers. If you take a broader view of the Gopher network to include all the anonymous FTP servers that are accessible via our virtual Gopher → Archie → Gopher-FTP gateways, then the number jumps into the thousands. If you include all the WAIS sources, then things get higher still. All this without leaving Gopher's little paws. If you consider all the "TELNET based" servers that can be reached by launching a TELNET session based on a Gopher descriptor, then, well, you get the picture.*

Aside from the question of numbers, Anklesaria is making an interesting point: from the start, Gopher has been developed with as many "hooks" to other information services as possible. In this quote, he mentioned production gateways to Archie, FTP, and WAIS servers and then went on to describe development gateways to Worldwide Web and MIT's TechInfo.

2 Why Should UF Be Interested?

Gopher offers a pragmatic solution to many of the publishing and access challenges our university faces as well as an opportunity to provide new, value-added services for our users.

As a vehicle for publishing, Gopher is available 24 hours a day to users with network or dial-up access. Anything that can be viewed onscreen can be downloaded or printed locally.

As an access tool, Gopher offers the user an environment in which she may do productive things with virtually no instruction. For the more sophisticated user, it offers speed, convenience, and full-text searching.

If you imagine the heirarchy of Gopher directory lists mirroring the heirarchy of university departments, it becomes evident why this would be an appropriate interface for a campus-wide information system (CWIS). "Departments," in the form of subdirectories and servers, can be added to the system as campus resources become available and can actually be located and maintained in the departments they represent.

For instance, if the athletics department produces a document outlining its annual schedule, then the machine on which the document is maintained could be designated a Gopher server[3] and added to the network. As changes are made to the document, those updates would be immediately available to Gopher users.

[3] It should be noted that not all microcomputers are powerful enough to support "background" operation as a Gopher server. In this case, the document would have to be posted on a more appropriate machine.

Gopher was recently chosen for the University of Alberta's CWIS[4] based upon the following selection criteria:

1. The CWIS must be accessible from all our campus platforms (i.e., micros/workstations with Ethernet, coaxial, or dial-in network connections, terminals and supported mainframes).
2. The CWIS must have an easy-to-use program interface.
3. The system should support delegation of information administration.
4. The support and general implementation of the product should not be too onerous.
5. The system should provide extensions such as file transferring and printing.
6. The CWIS should have compatibility with current and future data communications standards.

The University of Florida could likewise adopt these criteria in its efforts to develop a CWIS. If so, then Gopher would serve as a good implementation. In particular, the software is already available for a wide variety of machines. And the software is free. Clients and servers exist for VM/CMS, VAX/VMS, Unix, NeXT, Macintosh, and IBM PC.

There are virtually unlimited applications for a CWIS information system at UF. Some that come to mind are:

- **Facilities Information.** Information about building, department, facility and service locations, parking, security, and emergency procedures, etc.
- **Services Information.** Information about cafeterias, the bookstore, campus computer store, bank machines, bus routes, athletics and recreational programs, etc.
- **Campus Directories.** The conventional printed directory is expensive to produce, goes out of date quickly, does not include email addresses, is only useful for those who have a copy, and is difficult to modify between printings. A Gopher-accessible directory[4] could be published dynamically. For instance, a new faculty member's directory information would be available online the day he or she is appointed, telephone number or email address updates could be recorded immediately, and UF's directory information would be available from anywhere on the Internet[5].
- **Upcoming Events.** Information about public events (theatre, lectures, workshops), conference schedules (UF and elsewhere), service interruptions, etc.
- **Academic Calendar.** An online version of the calendar could include examination timetables, registration information, course descriptions, etc.
- **Public Affairs.** Many of the University's printed publications could be made available with Gopher: articles from the *Alligator*, minutes of Senate and Board of Governors meetings, press releases, statements of University policy, etc.
- **Libraries.** The UF libraries could provide facility and service information, user orientation, reserve reading lists, and LUIS orientation and tutorials all through Gopher. An electronic journal reader[2] might be implemented as well as electronic gateways to other archives and databases for all disciplines.

3 The Current State

Gopher has arrived at the University of Florida, but only a few servers[6] have been established. Furthermore, these servers aren't coordinated conceptually or physically; almost no links exist between the current servers.

3.1 Existing Servers

There are a handful of Gopher servers currently operating on campus. All of these are in the experimental stage and are under construction. Most are being used to test and illustrate basic Gopher capabilities.

1. **NERDC.** This server contains the skeleton for a UF Gopher "root" with a hierarchy for all the UF departments. Resources currently

[4] Such as the "Ph"[3] or "CSO" system previously proposed by ECS-DS.

[5] Gopher supports gateway access to X.500 directory servers. In the future, X.500 will probably be used to provide campus directory information at UF.

[6] Known as "Gopher holes".

available include the Ph/CSO electronic directory and links to other popular information services (like weather). The only departmental information available is from the library. Host: `nerix.nerdc.ufl.edu` (an IBM 3090 running OSF/1 Unix).

2. **CIRCA.** This server is the closest to being a "production" system; a list of undergraduate course offerings complete with availability/seating information is updated at half-hour intervals from the Registrar. During Drop-Add, terminals in the CSE lab run a Gopher client and thus serve as course browsers. A skeleton for CIRCA public lab information also exists. Host: `elm.circa.ufl.edu` (a DECStation 5000 running Ultrix).

3. **CIS.** This server houses a directory for general Gopher information including local UF information and technical information such as the software distribution and protocol documents. A skeleton for CIS information such as the graduate program and technical report archive also exists. A link to an archive of Clinton/Gore position papers and speeches exists. Host: `gopher.cis.ufl.edu` (a Sun SPARC running SunOS).

4. **ECS-DS.** This server contains a skeleton for Unix/networking technical information. Access to the RFC[7] and FTP software archive is available. Host: `gopher.eng.ufl.edu` (a Sun SPARC running SunOS).

5. **Center for Mammalian Genetics.** This server was developed by Andy Wilcox of ECS-DS to mirror the Indiana University Biology Archive. This archive contains access to Genbank sequences, the NSF Gopher, and other biocomputing resources. A local archive for Drosophila genetic/biological data also exists. Host: `cmg.health.ufl.edu` (a Sun SPARC running SunOS).

6. **Maternal & Child Health.** Not much is known about this server. A small archive on family planning and infant health exists. Host: `mchnet.ichp.ufl.edu` (a Macintosh).

[7] Request For Comment, documents used by the Internet community to detail standards and practices.

3.2 Plans for a University-Wide System

In an effort to move forward, David Nessl of NERDC, Mike Wright of CIRCA, and I have been meeting[8] to discuss Gopher in general as well as specific models for a cohesive UF deployment. Our discussions have focused on the following:

- *Establishing a single, well-known server to act as the "root" for UF.* To this end, the network alias `gopher.ufl.edu` has been established. This address will be used for the UF Gopher root. Currently, this address points to `nerix.nerdc.ufl.edu`, the NERDC's Unix partition running on it's 3090 mainframe. David Nessl administers this machine.

- *Establishing guidelines for UF servers.* This would ensure a consistent "look and feel" to the servers at UF. The NERDC and CIS servers reflect our ideas. The main thrust has been for clarity, flexibility, and accountability. This work is continuing.

- *Developing documentation for those interested in running their own server.* A localized "how-to" manual could describe how best to structure and present information through Gopher. This topic needs to be addressed further.

- *Providing servers for those who wish to publish information via Gopher but don't have the resources to run their own servers.* It has been suggested that the NERDC server be used for this purpose. In this scenario, NERDC would provide the computational facilities for Gopher (disk space, access) as per its current mission. The details of this service need to be worked out.

- *Establishing a single point of contact for Gopher help.* This resource would have knowledge of the Gopher system in general, possess the software distribution, and otherwise distribute the help resources discussed above. Currently, the NERDC support desk will field Gopher questions by phone at 904/392-2061. Questions by email may be directed to `gopher@ufl.edu`.

- *Establishing a forum through which local Gopher development can continue and include those who wish to participate.* Accordingly, we have created a mailing list

[8] Electronically, of course.

`uf-gopher@ufl.edu` for this purpose. Requests for additions, etc., should be directed to `uf-gopher-request@ufl.edu`.

Our intentions are to provide a prototype system that illustrates the basic capabilities of Gopher. We hope that we can then demonstrate the system and motivate its further development in an official capacity. We believe that Gopher has tremendous potential as a CWIS.

3.3 UPS Project

To illustrate how departmental information might actually be distributed and used at UF, I contacted University Personnel Services (UPS) and proposed that it participate in a project to make some employee information available via the NERDC Gopher.

After speaking with Jodi Gentry, communications coordinator for UPS, we decided to put the *A&P and USPS Employee Handbook* and *Training Connection* schedule into Gopher. Using the CIS server for development, we decomposed the handbook and training schedule and presented the data internally to UPS. The response was positive so the data are being announced formally in the next issue of *UF Perspective*, the UPS quarterly that will reach over 11,000 employees by the second week of September.

The NERDC server has been configured to log all of its transactions, so we should be able to evaluate quantitatively the usefulness of the data. At that point, we will probably analyze the requests by topic, host, and time of day.

3.4 What's Needed for the Future

The Gopher system has been discussed by the Networking and Telecommunications Subcommittee (NTSC) of the Council for Information Technology Services (CITS) and by CITS itself. Thus, Gopher is known to the official policy-setting bodies for UF computing.

The input I have received has been mostly positive and has driven our current situation considerably. It is hoped that CITS considers Gopher more closely now that a prototype system is nearly in place. This document may serve to widen the audience of those interested in Gopher at UF.

Those experimenting with Gopher should continue to do so, and those interested in getting involved should step forward. The next few months should be an interesting time for Gopher at UF.

We'll see you there.

A Addendum

For users with access to the network (Internet, including UFNET) or NERDC's VTAM, but who do not have a Gopher client, courtesy Gopher access accounts have been established. These accounts allow users to try out the Gopher system and learn where to secure the client software.

Internet/UFNET users can TELNET to the address `gopher.ufl.edu` and log in as the userid `gopher`. NERDC users can request `UFINFO` at the `NERDC VTAM IS ACTIVE` screen.

No passwords will be required. The accounts illustrate the Unix and CMS Gopher clients and the UF Gopher "root" server.

References

[1] Bob Alberti, et. al. *The Internet Gopher Protocol: A Distributed Document Search and Retrieval Protocol*, Draft 11, University of Minnesota Microcomputer and Workstation Networks Center, Spring 1992. Available via the UF Gopher at `gopher.ufl.edu` in the "Computing and Networks" directory under the "UF Gopher Information/Technical" menu.

[2] Michael O'Henly. *An Internet Gopher for the University of Victoria*, Library Systems Office, April 27, 1992.

[3] Bradley C. Spatz. *Ph: An Electronic Campus Directory System*, Distributed Systems, Engineering Computing Services, College of Engineering, University of Florida, February 4, 1992. Available in Postscript form via anoymous FTP from `eng.ufl.edu` in the `pub/ph` directory as `manual-draft.ps`.

[4] Debbie Wong. *Campus Wide Information Systems Evaluation*, University of Alberta, January 1992.