An Internet Gopher for the University of Victoria

Michael O'Henly
Library Systems Office
lux@sol.uvic.ca

"A time of turbulence is also a time of great opportunity for those who accept, understand and exploit the new realities." — Peter Drucker

Introduction

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 libraries and that we (the UVic Libraries) should consider implementing it at our university.

This document is directed to anyone with an interest in new library technologies and, in particular, how these may impact on the function of the UVic Libraries.

Part One:

Describes Gopher generically as a *networked application* and as an *information server*.

Part Two:

Proposes ideas about how an internet Gopher could be implemented at the University of Victoria.

Appendix One:

Bob Alberti and others, The internet Gopher protocol: a distributed document search and retrieval protocol, draft 11, University of Minnesota Microcomputer and Workstation Networks Center, Spring 1992.

Appendix Two:

Debbie Wong, Campus Wide Information Systems Evaluation, University of Alberta, January 1992.

Part One:

The Internet Gopher

What's an Internet Gopher?

The University of Minnesota's Microcomputer and Workstation Networks Center describes it 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…

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 that the server can connect to. (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 their 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…

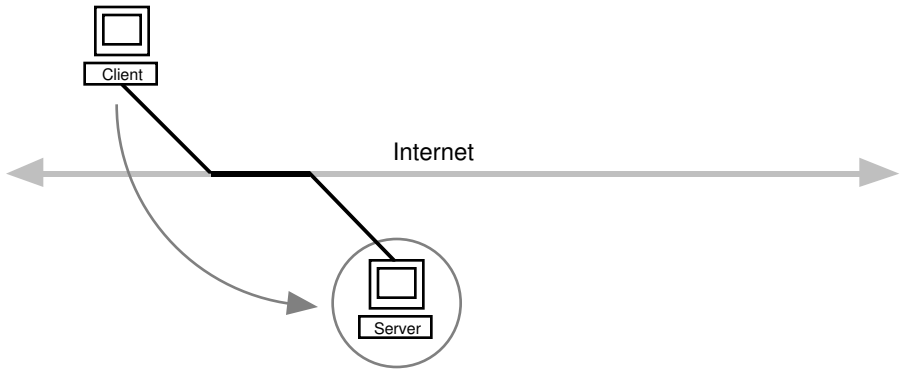

A Gopher server can provide links to other Gopher servers. In this configuration, the server that the user's client connects to is referred to as the *top-level* server. The diagram below shows the client connecting to a top-level server[1] and accessing a resource — via an intermediate machine — which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

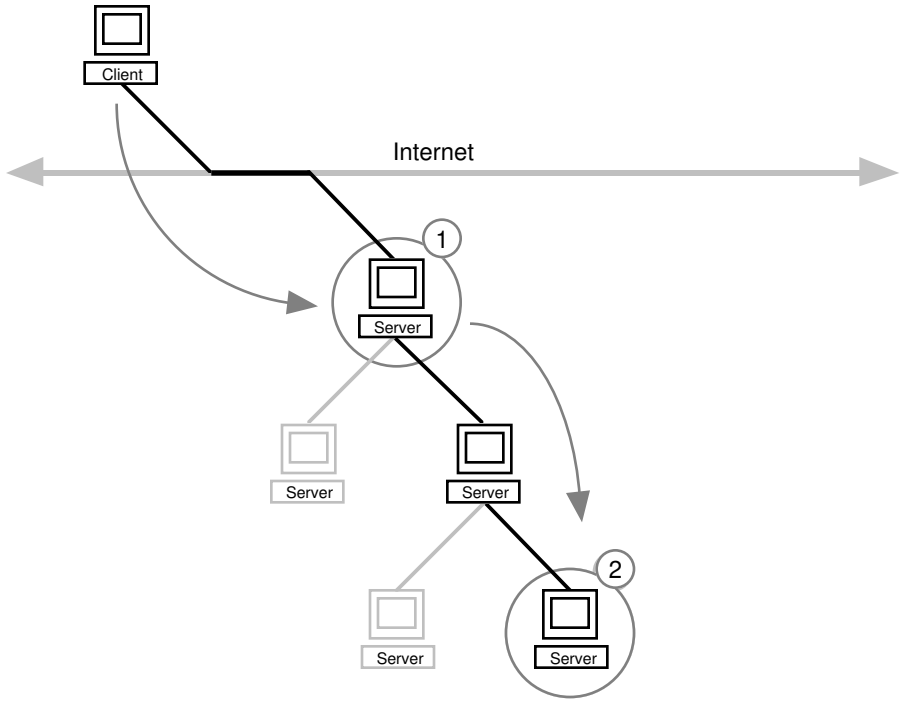

Other Gopher servers may be located within a particular top-level server's *domain* (for example, distributed across a university campus)…

…Or they may be in other domains[3].

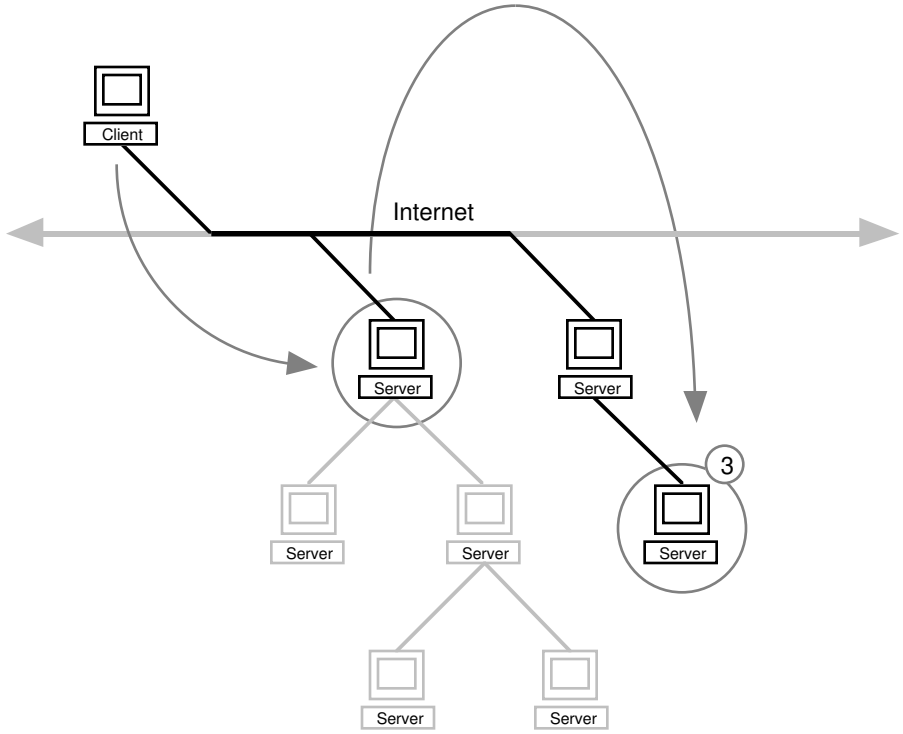

A Gopher server can also include links to other non-Gopher hosts — for instance, an internet-accessible Library OPAC[4].

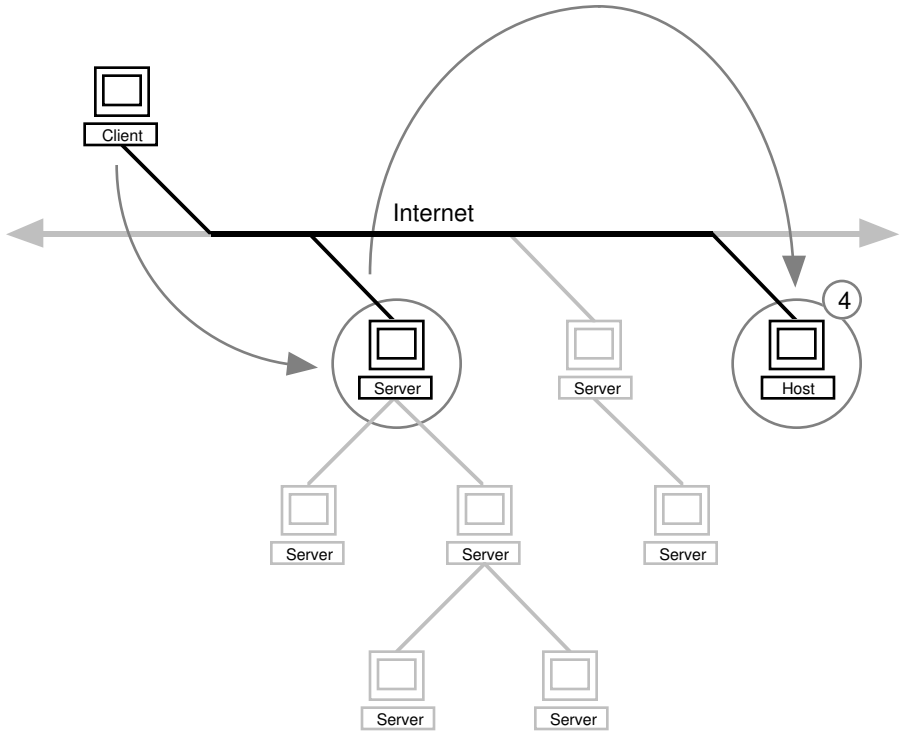

Gopher as an information server…

When the user connects to a Gopher server, the client software displays a directory listing like this…

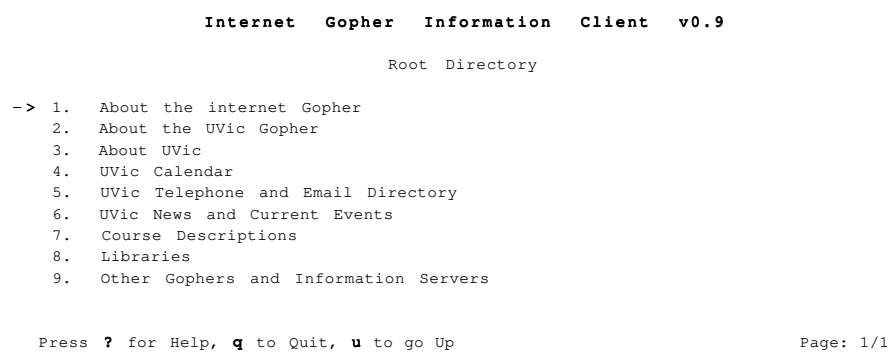

In "The Internet Gopher Protocol" (Appendix 1) eleven types of directory list items are defined. (Only those items printed in bold are described here.)

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. [Use of this type is discouraged.]
5 Item is DOS binary archive of some sort. [Use of this type is discouraged.]
6 Item is a UNIX uuencoded file. [Use of this type is discouraged.]
7 Item is an Index-Search server.
8 Item points to a text-based telnet session.
9 Item is a binary file. Client must read until the connection closes. Beware.
+ Item is a redundant server (same information as the previous server).

Type "0" — Item is 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 and Gopher displays the full text on their screen. A simple text-browsing facility allows them to page backward and forward through the file, perform string searches, etc. When they quit the text browser, they are returned to the directory list from which they made their selection.

Type "1" — Item is 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 "Libraries". Selecting this leads to a new subdirectory listing…

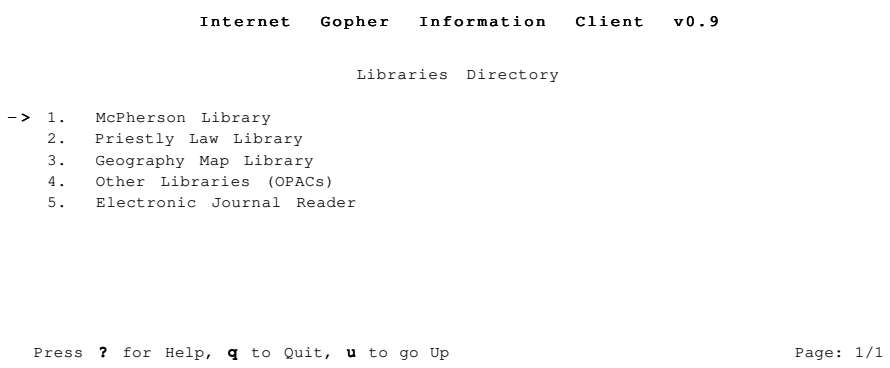

Selecting "Priestly Law Library" then brings up a further subdirectory that lists items specific to that library. And so on...

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

Type "7" — Item is 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, they are 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 can be selected from or browsed using the same comands).

Type "8" — Item points to a 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 "Queen's University", for example, from the "Other Libraries (OPACs)" directory list, and be immediately connected to the Queen's library services menu. When they select "Quit" from the Queen's system, they would be returned to the "Other Libraries (OPACs)" directory list.
- *Documents and other resources may be located on servers distributed across the internet.* As a user navigates down through a directory heirarchy, they may also 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 the University of Saskatchewan has developed a particularly good "internet gateway" for library OPACs. Rather than UVic reinventing the wheel, we could simply have the "Other Libraries

(OPACs)" directory item point to a telnet session that connects the user directly to the University of Saskatchewan application.

The following diagram illustrates how resources on a remote Gopher system can be linked to in such a way that the user sees them as extensions of their "local" system.

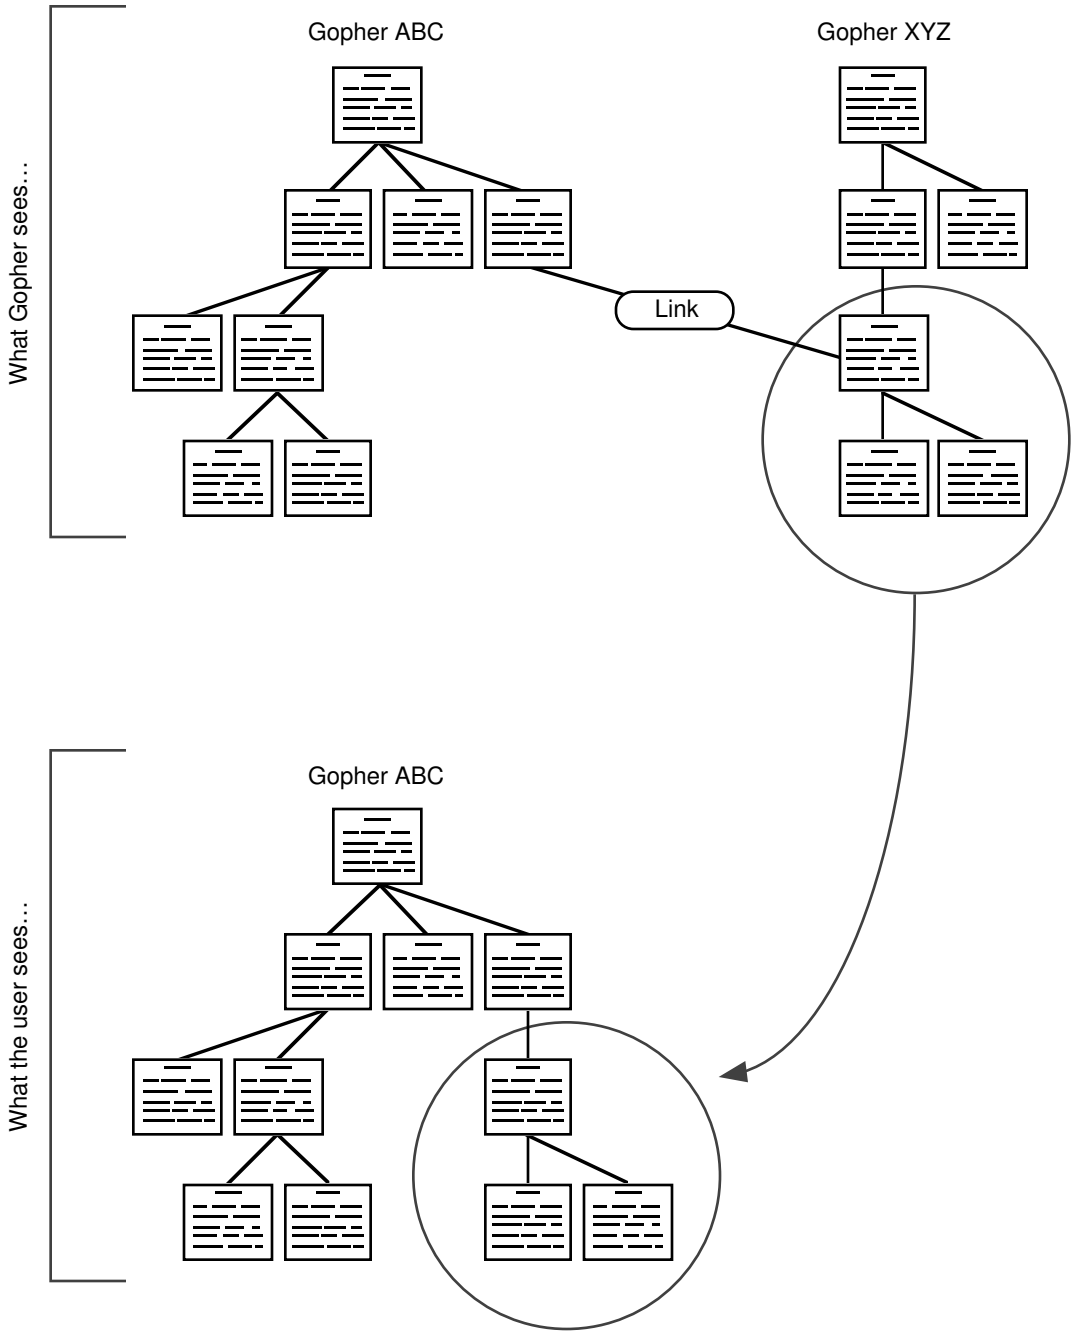

Part Two:

Gopher and UVic

Why should the UVic Libraries be interested?

Gopher offers a pragmatic solution to many of the *publishing* and *access* challenges our Libraries face, as well as an opportunity to provide new, value-added services for our users. In terms of sheer practical benefit, I believe that a Libraries information server could contribute as much toward meeting the needs of the university community as VICTOR.

I would like to see a jointly-developed library information system that would reside at UBC, SFU, and UVic, allowing resources at these universities to be shared — by each other, and by all the member insititutions of the Electronic Library Network. As the Colleges and Institutes are brought onto BCnet, they could (by setting up their own servers) make resources they have developed locally accessible to the rest of the system.

As a vehicle for publishing, Gopher is available 24 hours a day to users with network or dialup access (see **Implementation Ideas** for details). Anything that can be viewed onscreen can be downloaded or printed locally. It is a multi-user system and could be expected to support 25-30 simultaneous sessions without unacceptable effects on performance.

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

The following are examples of ways we might use Gopher in the UVic Libraries…

Facility Information

Descriptions of our facilities that anticipate questions like…

Are there legal-sized photocopiers in the Curric Lab?
Who do I talk to about booking a study carrel?
Are there any special facilities for the handicapped in the Law Library?

Service Information

Descriptions of our services that anticipate questions like…

Is the Geography Map Library open on Victoria Day?
What date do my books have to be returned by at the end of term?
What is the Library's loan policy on serials?
What is the charge for an interlibrary loan?

User Orientation

Printed handouts that are produced for users could be posted on Gopher, including…

- Guides to reference materials
- Guides to CD-ROM and Datasets collections
- Changes to Library policies and procedures

Reserve Reading Lists

Students and faculty members could use Gopher to access (and print out) up-to-date reading lists without making a trip to the Libraries.

Finding Aids

The University Archives has prepared extensive finding aids for its collections. Publishing them on Gopher would make useful information about Archive holdings available to anyone in the world with internet access.

Finding aids and holdings lists could also be posted for the Film Centre, Music & Audio, and the Geography Map Library — three other areas with collections that are not easily accessed through VICTOR.

Library Directory

Most students and faculty at UVic probably have little idea of "who does what" in the Library. One way of narrowing the gap between Library professionals and users would be to publish a function-oriented directory with names, phone numbers, and email addresses.

VICTOR Orientation and Tutorials

Many microcomputers and workstations support a "windowing" user interface. One interesting application of Gopher would be an interactive VICTOR tutorial. The user would run the Gopher client in one window, selecting the tutorial from a directory list. In a second window, they would run another application that would connect them directly to VICTOR. Then they would work through the tutorial instructions displayed in the Gopher window while actually performing them in "real time" in the VICTOR window.

Applications like this could supplement the traditional Library Lessons.

Internet Gateway

There is an ever-increasing number and variety of resources becoming available across the internet. Gopher provides a simple, extensible framework for the delivery of these services. No particular expertise is required on the part of the user and the workload involved in adding and deleting pointers to internet resources would be minimal.

These resources include…

- numeric, bibliographic and full-text databases for all disciplines
- document archives for all disciplines

- Library OPACs
- other (non-Gopher) information servers
- archives of public domain software
- commercial bibliographic database and document delivery services

Electronic Journal Reader

No one disputes that electronic publishing is going to have a critical impact on libraries, but there has been little discussion about how ejournals might be offered to the library user. Currently, users subscribe to the titles that interest them and receive copies of each issue as email. If fifty people at a university subscribe to one title, then fifty copies wind up in fifty mailboxes…

The Libraries could offer a valuable service in this area by configuring a Gopher server to act as an ejournal "reader". This would mean that only one subscription would be required for each title, only one copy would have to be stored, and the user could browse spontaneously among a broad range of titles. Most importantly, we could offer our users full-text indexing across an entire ejournal collection.

There are many unanswered questions about the future of electronic publishing. It's hard to predict, for instance, how cost and copyright issues will bear on the way users ultimately access ejournals. From the Libraries' point of view, however, putting up an ejournal reader on Gopher could give us some practical experience in both developing an access method and sharing resources with other institutions. Whether or not this will prove to be the best method in the long run is perhaps less important than the notion that we should start experimenting now with access tools and inter-institutional relationships that will be commonplace in the library of the future.

Why should other UVic departments be interested?

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 (see **Implementation Ideas**).

For instance, if Athletics and Recreational Services produces a document outlining their annual schedule, then the machine on which the document is maintained could be designated a Gopher server and added to the network. As changes are made to the document, those updates would be immediately available to Gopher users. (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 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.

[Note: The complete University of Alberta evaluation is reprinted in Appendix Two.]

There are virtually unlimited applications for a campus-wide information system at UVic. 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 could be published dynamically. For instance, a new faculty member's directory information would be available online the day they are appointed, telephone number or email address updates could be recorded immediately, and UVic's directory information would be available from anywhere on the internet. *[Note: Gopher supports gateway access to .X500 directory servers. In future, .X500 will probably be used to provide campus directory information at UVic.]*

Upcoming Events

Information about public events (theatre, lectures, workshops), conference schedules (UVic 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 Ring, minutes of Senate and Board of Governors meetings, press releases, statements of University policy, etc.

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 I asked Farhad Anklesaria, from the University of Minnesota's Microcomputer Research Group, 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 networ 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 base 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 email message, 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*.

For Canada, the following sites are listed (and can be connected to by selecting them from the directory)…

- Lakehead University
- Nova Scotia Technology Network
- Queen's University
- University of Guelph
- University of Manitoba
- University of Waterloo

The University of Alberta plans to have a Gopher-based CWIS (campus-wide information system) in production by May 1, 1992 (see Appendix 2). Simon Fraser University and other member institutions of COPPUL (Council of Prairie and Pacific University Libraries) are also reported to be investigating Gopher.

Implementation Ideas

These are my own thoughts about how Gopher might be implemented at UVic. *[Note: This section is unfinished as of Apr 27. I'm still thinking about it...! M.O'H.]*

First, I think it's important to establish what function a UVic Gopher would be intended to serve. I can imagine a system based on three different models…

1. **A Libraries information server**

   A Library information server would offer only Library-related services. It would be administered by Library staff and would require minimal support from Computing Services. Future growth would be in the direction of resource-sharing with other ELN and COPPUL institutions.

2. **A campus-wide information system co-ordinated by the Libraries**

   A Libraries-coordinated CWIS would offer a full selection of information services representing the requirements of all areas of the university. The Libraries' role would be essentially editorial: establishing standards of presentation, liaising with participating departments, and ensuring the overall effectiveness of the system. Technical support and networking issues would be the responsibility of Computing Services. *[Note: University of Waterloo has selected this model.]*

3. **A campus-wide information system co-ordinated by Computing Services**

A CWIS coordinated by the Computing Centre would offer the same selection of information services, but would be developed and supported by this department in the same way as other major administrative applications (i.e., MAIL/BOOK, SABS, CHRIS, etc.). *[Note: This is how the majority of public Gophers are implemented.]*

I think the first model — a Libraries information server — is the one most consistent with the resources available to us. The system I am about to describe would fulfill all foreseeable requirements for the Libraries, but has the potential to become part of a fully-fledged CWIS if other departments show interest and support.

Libraries Information Server

The Libraries' Gopher server would run on a Library-owned NeXT workstation (see **System Costs** below) connected via Ethernet to the internet. It would be assigned the domain name, *gopher.library.uvic.ca*, and a UVic network identity.

This system would be available to the university community, to the internet, and to member institutions of the Electronic Library Network (see **Questions** below).

The Libraries and Computing User Services would make microcomputer and workstation Gopher clients available at no cost (i.e., documentation and a copy of the client software in exchange for a blank diskette). Mainframe clients would also be available to all account holders on systems operated by Computing Services (i.e. *UVVM*, *sol*, *ra*, etc.) For the ELN dialup and university network user, a VT100 client would be provided on the Gopher server.

Resources provided by this server could include all the things listed in "**Why should the UVic Libraries be interested?**" above. Making these available is not so much a technical issue as one of organizing the production of source material. For example, adding a VICTOR tutorial to the Gopher server would be a trivial task — but first the tutorial document must be written, and that will take someone a chunk of time.

As a starting point, I would suggest that…

- Libraries user orientation documents
- an internet gateway to regional OPACs (UBC , SFU, other ELN institutions)
- an electronic journal reader and index server

…would be resources that we could expect to make available at an early stage.

System Requirements

The NeXT is a Unix workstation with additional features that make it attractive for this application. In particular, the NeXTstep operating system includes a full-text search engine and the NeXT version of the Gopher server has been written to take advantage of this. Through an arrangement with SFU and UVic's Faculty of Engineering, we are able to purchase NeXT products at significant discounts. In fact, the NeXTstation Turbo that I recommend would cost less than a high-end Macintosh.

Storage would be provided by a 2 gigabyte external hard disk (2 Gb. = 2000 Mb. = about 1,024,000 pages of text). Data would be backed up onto a 2 Gb. DAT tape drive. At the moment, prices are falling dramatically on large volume drives and tape backup systems. Performance and reliability are excellent (five year warranties are common for third-party products).

System Costs

NeXTstation Turbo (400 Mb. HD, 16 Mb. RAM)	.....$7360.	(Cdn., PST/GST incl.)
2 Gb. external HD	.....$2500.	(U.S., approx.)
2 Gb. DAT Backup	.....$2200.	(U.S., approx.)
Miscellaneous (cabling, software, supplies, etc.)	.....$1500.	(Cdn., approx.)

Questions

This description of a Libraries information server is, of course, quite incomplete and leaves several important implementation and policy questions unanswered.

- How do we provide operational and user support for this system?
- How do we provide access for ELN institutions that are not yet connected to BCnet?
- How do we gain access to Library OPACs without a “local” login? At the moment, both SFU and UBC require userids and passwords.

Appendix One

Bob Alberti and others, **The internet Gopher protocol: a distributed document search and retrieval protocol**, draft 11, University of Minnesota Microcomputer and Workstation Networks Center, Spring 1992.

Appendix Two

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