Library 2000

June 15, 1992

Academic Staff

Jerome H. Saltzer, Group Leader

Research Staff

Mitchell N. Charity

Undergraduate Students

Jeremy A. Hylton
Robert C. Miller
Arthur W. Min
Manish D. Muzumdar
Ronald Weiss
(Brandeis University)

Reading Room Staff

Paula Mickevich
Carol A. Nicolora
Maria Sensale
Rebecca J. Soble

Support Staff

Lisa M. Kelly

Library 2000

Library 2000 is a new research project, exploring the system implications of large on-line storage. The method of the project is pragmatic, to develop a prototype of an on-line electronic library using the technology and approaches expected to be feasible in the year 2000. Initial support for Library 2000 has come from grants from the Digital Equipment Corporation and the IBM Corporation, and uncommitted funds of the Laboratory for Computer Science.

The basic hypothesis of the project is that the technology of on-line storage, display, and communications will, by the year 2000, make it economically possible to place the entire contents of a library on-line and accessible from computer workstations located anywhere. The goal of Library 2000 is to understand and try out the system engineering required to fabricate such a future library system. The project's vision is that one will be able to browse any book, journal, paper, thesis, or report in the library using a standard office desktop computer, and follow citations by pointing–the thing selected should pop up immediately in an adjacent window.

The key technology required to realize this vision is low-cost disk storage. In the last decade, mass-produced disk storage has fallen in cost/bit by a factor of 1000. Projections of current research and development activity suggest that disks in volume production by the year 2000 will fall in cost/bit by at least another factor of 100. This change of five orders of magnitude calls for a complete rethinking both of what is economically feasible, and also of engineering for effective use. A back-of-the-envelope estimate suggests that the cost of the magnetic disk storage needed to hold scanned images of all the books and serials in a large library will, in the year 2000, be about equal to the annual budget for that library. When floor space is considered, the cost of holding images on magnetic disk will be substantially less than the cost of holding the paper form.

Three other technologies are also advancing at a pace that is likely to provide effective support to an on-line electronic library system:

- Data communications. Medium bandwidth networks already in place, in the form of campus networks and the NREN, provide the bandwidth needed to allow access to a library from a distance. Planned higher bandwidth networks should change transmission of scanned images from the occasional to the routine.

- Display technology. Multi-plane megapel displays are becoming common on medium-cost engineering workstations. By the year 2000, these displays will probably be standard on the lowest-cost personal computers. (The grey-scale capability that comes with a multi-plane display is critical to make reading of scanned images acceptable.)

- System organization. The client/server model of organizing distributed computation has matured sufficiently that it appears to be the method of choice in designing an electronic library system.

Our research problem is not to invent or develop any of these four technologies, but rather to work out the system engineering required to harness them in a usable form. The engineering and deployment of large-scale systems is always accompanied by traps and surprises beyond those apparent in the component technologies. Finding things, linking things together (especially across the network), keeping the whole system reliable, and making the system last for a period of decades will all require new ideas.

Architecture

The overall system architecture of Library 2000 consists of a workstation client that is responsible for matters of presentation, user interaction, and usage coordination, together with a multiplicity of storage servers and of index servers. Storage servers hold the raw information, in at least two forms: scanned bitmap and ASCII text. Index servers provide indices of the ASCII text to allow searching. The overall paradigm of use is deceptively simple: a user expresses interest in some item to the workstation client, the client dispatches a query to one or more index servers, and if the query is successful, uses the information returned by an index server to request items from one or more of the storage servers.

This architecture is appealing because it allows modular, independent, competitive design and replacement of the user interface, the index services, and the storage services. It also permits a uniform user interface to a wide variety of different library collections as well as to personal files, mail, or other non-library databases. The architecture is simple and the functional boundaries match both physical and administrative boundaries. Finally, traditional library back office operations such as cataloguing, circulation, acquisition, journal control, etc., fit in gracefully as additional clients and servers.

Two technology observations interact to produce an interesting architectural consequence. The cost of magnetic disk memory has for many years been between 10 and 100 times cheaper than that of random-access memory. Similarly, the amount of space required for (compressed) scanned images is between 10 and 100 times larger than that required for the corresponding ASCII text. The architectural consequence is that if one spends about the same number of dollars for each, there will be space in RAM for a complete index of the words stored in page-image form on the disk. This observation leads to an interest in index-preparation and searching algorithms that are optimized to operate directly in large random-access memory, even for very large databases.

The network protocol that connects the presentation client with the index and storage servers is stateless, to achieve robustness in the face of network and server failures. It uses unique identifiers, to allow separation of the index and the storage services. Storage servers are replicated with wide geographical diversity, to counter threats to persistence over time periods measured in decades.

The Prototype

The general research strategy of Library 2000 is to build a small but extendable prototype system, stock it with live data, see how it works, and then iterate the design at successively larger scales. The first prototype was placed in service during the summer of 1991. It involves a very simple presentation client accessible by telnet from anywhere in the Internet, and a combined index/storage server that contains the catalog card records of the joint reading room of the M.I.T. Laboratory for Computer Science and the M.I.T. Artificial Intelligence Laboratory. This card catalog has about 15,000 records. The prototype also contained a second collection of 16,000 library card records, on the subject of computer science, extracted from the M.I.T. library catalog. During the fall, Mitchell Charity added a third collection, consisting of all available abstracts of about 2000 LCS and AI technical reports and memos.

This online catalog has now become the primary catalog of the LCS/AI reading room; the paper catalog was closed in January, 1991. Several support systems, including a simple emacs-based cataloguing system were also developed, so that operations of the reading room could be transferred entirely to the new online catalog. A shelf-reading project is underway, and as of the end of the reporting year all proceedings, all books received since 1980, and all journals are now fully catalogued and checked. All technical reports received since January 1991 are also catalogued; about half of those received before that date are also represented in the catalog. The staff of the reading room reports that circulation of technical reports has increased markedly since the on-line catalog became available.

During the fall, the entire system migrated from a borrowed VAX server to a set of three replicated IBM RS-6000 servers, with a view towards scaling up to much larger quantities of data. The index was recoded from the PERL (interpretive) to the C language, and it is now in the midst of a complete design revision. By taking advantage of the large RAM available on the RS-6000s, this revision is intended to provide the high performance needed to allow indexing of the 700,000-record M.I.T. library catalog. During the spring, this new index server was tested on a 25% sample of the M.I.T. catalog with good results; the complete index system should be ready to deploy some time this summer. Mitchell Charity did most of the work on the prototype design and implementation.

Three related projects were completed this year. First was a parser for library standard (MARC) records, by undergraduate Art Min. Second was a preliminary prototype of an X-based user interface to the online catalog, initiated by Mitchell Charity with further development by Jeremy Hylton. The third was a completely different user interface based on the Organization Engine, an information organizing system in use at the Digital Equipment Corporation Cambridge Research Laboratory. This interface was developed as an undergraduate thesis at Brandeis University by Ron Weiss.

Currently, Manish Muzumdar is developing a general client interface toolkit, and Robert Miller is working on the indexer overhaul.

Joint Projects

Library 2000 is working closely with two other groups interested in this area: the M.I.T. Distributed Library Initiative, and a group of computer science departments that are collaborating on technical report distribution. These collaborations represent an opportunity to obtain real, production library materials in electronic form, and also an opportunity to influence the architecture of future library systems.

The Distributed Library Initiative is a new, joint activity of the M.I.T. Libraries and M.I.T. Information Systems. Its purpose is to develop and experiment with new technologies for the M.I.T. Library system. The DLI is installing prototypes of possible future library electronic delivery systems, and is working with publishers who are undertaking experiments with electronic distribution systems. Library 2000 is working with DLI at several levels to help develop the DLI plans, to look for opportunities to avoid unnecessary mismatches of data formats, protocols, and programming interfaces, and to make joint use of software where feasible. The acquisition by Library 2000 of the M.I.T. library catalog and its update stream is one tangible result of this collaboration.

The technical report collaboration represents an extension of the prototype system in another interesting direction, that of linking indexed text with page images. The plan of the collaborators is that each university (initially Stanford, the University of California at Berkeley, Cornell, Carnegie-Mellon, and M.I.T.) would place page-images of its computer science technical reports in a local server. Each would also prepare electronic bibliographic records and distribute them to the other participants. Each participant would then take this set of bibliographic records, place it in some local index/access system, and work out ways of providing links from those records to the local and distant page image servers. This plan has progressed to the point that the Corporation for National Research Initiatives, acting as coordinator, has made a proposal that DARPA fund the project. In anticipation of funding, an online discussion group has already developed a proposed format for exchange of the bibliographic records. The M.I.T. portion of the proposal is actually a joint activity of Library 2000 and the M.I.T. Library System, with the library system planning to do scanning and paper delivery, and a potentially expanded version of the Library 2000 prototype as the underlying search, storage, and presentation system.

Plans

The general plan of the Library 2000 project is to continue the development, iteration, and extension of the prototype until it becomes a full-fledged on-line library system. The immediate extensions underway are those already mentioned, to the full M.I.T. library card catalog, and to the page images of the technical report collaborative project.

The addition of page images actually leads to extensions in several different areas, each exposing a variety of new system design problems:

- Development of an image storage service and storage service protocol.

- Development of an X-based client for query and image display.

- Collection of original word-processing text for full-text indexing.

- Investigation of how to relate the ASCII text to the corresponding scanned images.

- Development of a document linking plan.

Two other, related projects are also contemplated:

- Implementation of a simple replication system that works well under geographic separation of the replicants.

- Development of a collection discovery and rendezvous system.

These plans are actually too ambitious when compared with the intended size of the research group; some selection will occur.

Publications

1. Li Gong, T.M.A. Lomas, J.H. Saltzer and R.M. Needham, "Protecting Weak Secrets from Guessing Attacks," submitted to the *IEEE Journal of Selected Areas in Communications.*

2. J.H. Saltzer, "File System Indexing, and Backup," in *Operating Systems for the 90's and Beyond, Lecture Notes in Computer Science 563,* edited by Arthur Karshmer and Juergen Nehmer, Springer-Verlag, New York, 1991, pp. 13-19.

Theses Supervised

1. R. Weiss. Integration of the Organization Engine and Library 2000. Undergraduate honors thesis, Brandeis University Department of Computer Science, May, 1992.

Talks

1. J.H. Saltzer. Technology, Networks and the Library of the Future. Lecture given at MIT EECS Department Colloquium, October 28, 1991.