Volume 2
Number 1

ISSN 0730-9295

Information Technology and Libraries

March 1983

ITLBDC 2(1) 1–128 (1983); Inf Technol Libr

Information Technology and Libraries

Volume 2, Number 1: March 1983

CONTENTS

ITLBDC 2(1) 1–128 (1983); Inf Technol Libr

Information Technology and Libraries

Information Technology and Libraries is the official publication of the Library and Information Technology Association, a division of the American Library Association, 50 E. Huron St., Chicago, IL 60611; *Executive Director:* Donald P. Hammer. The journal is issued quarterly in March, June, September, and December.

Information Technology and Libraries publishes material related to all aspects of library and information technology. Some specific topics of interest are: Automated Bibliographic Control, AV Techniques, Communications Technology, Cable Systems, Computerized Information Processing, Data Management, Facsimile Applications, File Organization, Legal and Regulatory Matters, Library Networks, Storage and Retrieval Systems, Systems Analysis, and Video Technologies. *ITAL* welcomes unsolicited manuscripts. Submissions should follow the guidelines stated under "Instructions to Authors" on page 80 of the March 1982 issue.

Manuscripts of articles, communications, and news items should be addressed to: Brian Aveney, Editor, *Information Technology and Libraries*, Blackwell North America, 6024 S.W. Jean Rd., Bldg. G, Lake Oswego, OR 97034. Copies of materials submitted for review should be addressed to: David L. Weisbrod, *ITAL Book Reviews*, Systems Office, Yale University Library, Box 1603A, Yale Station, New Haven, CT 06520. Advertising arrangements should be made with William Z. Schenck, University of Oregon Library, Eugene, OR 97403.

Information Technology and Libraries is a perquisite of membership in the Library and Information Technology Association. Subscription price, $10, is included in membership dues. Nonmembers may subscribe for $20 per year. Single copies, $5.50.

Circulation and Production: American Library Association, 50 E. Huron St., Chicago, IL 60611. Please allow six weeks for change of address.

Publication of material in *Information Technology and Libraries* does not constitute official endorsement by the Library and Information Technology Association or the American Library Association.

Abstracted in *Computer & Information Systems, Computing Reviews, Information Science Abstracts, Library & Information Science Abstracts, Referativnyi Zhurnal, Nauchnaya i Tekhnicheskaya Informatsiya, Otdyelnyi Vypusk*, and *Science Abstracts Publications.* Indexed in *Computer Literature Index, Current Contents, Current Index to Journals in Education, Education, Library Literature, Magazine Index, and NewSearch.* Microfilm copies available to subscribers from University Microfilms, Ann Arbor, Michigan.

Second-class postage paid at Chicago, Illinois, and at additional mailing offices. *Postmaster: Send address changes to* Information Technology and Libraries, *50 E. Huron St., Chicago, IL 60611.*

A Paradigm Shift in Library Science

In other places I have gone on record concerning the unheeded research findings from catalog use studies. I suggested that these studies should be reviewed again, and the most constant recommendations should be incorporated into the design of the new form of library catalogs. What I failed to observe at that time was that some of the findings had been heeded by the library profession but that other findings were not heeded because they were not strong enough to cause a paradigm shift, in Kuhn's terms. In *The Structure of Scientific Revolutions,* Thomas Kuhn describes how research findings, which are substantiated and verified, can challenge scientific "truths" and can cause a scientific community to shift away from common wisdom to a new consensus about what is truth for them.

The paradigm shift in librarianship that did not occur, even though some research showed a certain tendency, has to do with the primary activity or activities at the library catalog. Common wisdom since Cutter's time has been that most users of the library want a catalog where they can find a particular item, a known item. Because we believed that to be true, we have invested a great deal of professional effort in providing catalogs that can help users perform that task par excellence. Generations of catalog specialists have produced descriptive catalog codes, standard unit records in card and MARC form, and have invested greatly in name authority projects. The strength of our belief in this truth is shown in our national bibliographical tool, the NUC, which only provides access to known items, as does the first national online bibliographic utility, OCLC. In recent times we have held international conferences, created international standards, and organized a Universal Bibliographical Control project—all devoted to bibliographic description of known items. Belief in the paradigm that the catalog was for known-item searching is very strong indeed.

Most catalog use studies, until recently, seemed to confirm that this belief was a correct one. But now we have findings from several *online* catalog use studies that dispute this "truth." The CLR/OPAC survey, recently completed (reported in this journal and elsewhere), online catalog transaction log analyses, focused group interviews, and data and statistical reports from several libraries (including the Library of Congress and the University of California Library System) attest to the "fact" that the great majority of library users are performing *topical subject* searches, not author/title or known-item searches. That is the overwhelming finding from these studies. It was also the finding from some earlier catalog use studies, but this time the weight of the evidence cannot be ignored or unheeded.

The important question is: What impact will these findings have on the professional work of librarians, the catalog codes, the efforts to standardize and internationalize our bibliographic records? If the findings are accepted as the "new truth" about catalog use, we may see a perceptible shift in professional work, with greater attention being paid to subject analysis and subject access in library catalogs. I don't expect a pendulum swing away from descriptive cataloging, but there may be just enough of a shift for balance so that the profession, the international scene, and the national libraries and bibliographic utilities will devote more resources to the improvement of this aspect of library catalog design. If this does not occur, we may find the *technology* of online library catalogs dictating the *science* of online library catalogs. That would be unfortunate for subject analysis and subject access are greatly aided by computers and other technologies, but the underpinnings of OPAC developments should be a firm understanding and professional agreement about what subject analysis and subject access *is* or *should be*.

We really don't know if augmented MARC records and keyword access are good enough to replace LCSH. Such an extremist view (which I personally do not advocate) needs to be justified and verified. We really don't know if LCSH online or a new library thesaurus will be the answer to improved subject access. We need to arrive at some professional consensus about it. As library scientists, we need to concentrate on these issues before we confirm which OPAC feature and what online authority and bibliographic records will be universally adopted. We cannot assume, as before, that subject authority formats will match name authority formats. The functions of known-item access and subject access are different and may require different codes, formats, and links to bibliographic records.

The challenge for the library profession now is *not* how to convert card catalogs to online catalogs, but how to convert unifunctional catalogs into multifunctional catalogs where known-item and topical-subject searches can both be successfully completed. Cutter will not roll over in his grave, but he may sit up and take notice! As will Ranganathan, Haykin, Metcalf, J. C. M. Hanson, Bliss, and C. Martel!

PAULINE A. COCHRANE
Professor, School of Information Studies
Syracuse University

EDITOR'S NOTES

Oldest Living Library Automation System?

Walt Crawford of RLIN and I were recently chewing the fat about systems we'd worked on. I noted that the batch acquisitions system I installed at Harvard in 1970 is still running. Walt said he had put a circulation system into Berkeley in 1968 that is still going. We'd be interested in hearing of other "old-timers." Drop a note with the library name, type of system, hardware and software details, and date of installation to *ITAL*. We'll publish results in the journal.

Born-Ag'in Automation

Jim Haas of CLR recently noted that "retrospective conversion is becoming a religious experience."

Doctors

While ALA journals do not usually note doctorates for authors, we were amused by the crop this issue. Nancy Carter's doctorate is in mathematics and higher education, Bennett Price's in classics, and Ed Brownrigg's in Spanish literature. Well-rounded types, we automation folk.


A-F
F-H
H-M
M-R
V-Z

The Concept of a MARC Format

John C. Attig

A MARC format is a distinct set of content designators prescribed for the identification and characterization of data in a particular category of MARC records. Of the various criteria that might be used to categorize records, certain ones are coded in the Legend of each record. The format of the record is deduced from these codes. Problems have arisen because formats have not been defined for exclusive categories of records and because consistency of content designation across formats has not been maintained. This paper proposes that (1) separate formats should be defined for functionally distinct types of records—bibliographic and authority records; (2) the present bibliographic formats should be merged into a single format; (3) inconsistencies should be removed; and (4) the criteria for identifying categories of records in the Legend should be re-examined.

In 1980, the Library of Congress Network Development Office conducted a review of the MARC formats. Recognizing that "analysis of the family of formats was needed in the light of advances in technology and interlibrary cooperation"[1] since the introduction of the MARC II format in 1968, the Library desired to examine the principles underlying the formats, to inventory the problems that had arisen, and to initiate revisions. As a result of this review project and subsequent work done by the Library, representatives from the national libraries and bibliographic networks, and ALA's MARBI Committee, considerable progress has been made in "the laying of a general foundation on which revised MARC formats can firmly stand."[2] A statement of the principles underlying the US/MARC formats has been approved.[3] A number of long-standing problems—such as institutional responsibility for data in a MARC record and the linking of related records—have been resolved and new principles have been formulated.[4] Remaining problems have been listed, and priorities and procedures for their solution are being discussed.[5]

The most significant remaining issue concerns the meaning and significance of the term *format*. In an appendix to the report of the format review project, its author states, "An effort should be made to provide a coherent rationale for dividing up the universe of bibliographic description into multiple MARC formats. This rationale would provide answers to such questions as: (1) Why do we need the various MARC formats for bibliographic data? (2) What are the criteria on the basis of which we distinguish one format from another?"[6] Behind these questions lies another: What is the meaning of the term *format*?

THE CATEGORIZATION OF MARC RECORDS

As an initial working definition, it will be assumed that a MARC format is a dis-

John Attig is a senior assistant librarian at the Pennsylvania State University Libraries, University Park, where he is a cataloger in the Bibliographic Services Department. He represents the Reference and Adult Services Division on the MARBI Committee. Manuscript received December 1982; accepted December 1982.

tinct set of content designators prescribed for the identification and characterization of data in a particular category of machine-readable record. Thus format is one of the criteria by which records may be sorted into categories—in this case, the set of content designators that are valid for the record. Unlike other criteria for sorting records, however, format is artificial and derivative. Each format is the result of a decision that a particular category of MARC records deserved a distinct set of content designators.

Before examining the categories for which the present formats have been defined, it would be useful to pursue the question of record categorization in general terms. By analyzing the universe of MARC records and the systems in which they interact, it should be possible to identify various criteria that might be used to sort records into categories. Although a rigorous analysis is beyond the scope of this paper, the following outline is a beginning. Each criterion suggested should be considered an independent dimension in an *n*-dimensional classification matrix.[7]

General

Type of Record

MARC formats have been defined for two fundamentally different types of records: bibliographic and authority records. Although the two types are related and indeed must interact in most automated systems, they are yet distinct both in function and content. Bibliographic records contain data that may be used to identify, control, and retrieve the material for which the record is a surrogate. Authority records contain data that may be used to identify, control, and retrieve the individual names and titles associated with various bibliographic items. The distinction between types of record is of primary importance, and in the outline below separate sections are devoted to criteria that apply only to bibliographic or to authority records.[8]

Level of Generality

Most MARC records are general in nature; the information is intended to be universally applicable—to all copies of an item, to any occurrence of a name or title. However, it is also possible for records to be specific to a network, an institution, or a single copy. In complex bibliographic systems designed to support multilevel data structures, clustering of records, local authority files, and copy-specific bibliographic descriptions, it may be necessary to categorize records in this way.

Conventions Governing Content

The MARC formats are standards for content designation. They do not specify which data elements are required for any record nor (except for such elements as coded data) the form in which the data is recorded. Content is governed by standards such as cataloging rules or classification schedules. Records need to be categorized according to the content standards that apply to a given record and the extent to which they have been followed.[9]

Autonomy

MARC records have tended to be self-sufficient. However, with the introduction of the technique for linking related records, with the advent of linked files in many systems, and with the development of interactive record exchange, this is no longer a valid assumption. Data may be split out of a record and maintained in separate, linked records. A complete record may need to be reassembled from various files before it can be used. Such records need to be identified as a distinct category in the communications format.[10]

Bibliographic Records

Intellectual Nature of Content

Bibliographic records are usually sorted into categories based on characteristics of the material described in the record. Categories based on intellectual characteristics include natural-language text, cartography, music, performance, numeric data/software, etc. These categories are not mutually exclusive, and it is doubtful whether a list of mutually exclusive categories is possible.

Physical Presentation

Physical presentation is logically independent of the intellectual nature of the material described. For example, the same language material may be printed, handwritten (manuscript), (micro-) filmed and projected, recorded or encoded for computer processing. Again the categories are not mutually exclusive.

Hierarchical Level

Bibliographic items often exist within a hierarchy of containing items—a monograph in a series, for example, or an article in a journal. A bibliographic record may contain information about several levels in the hierarchy, but one level (called the "target item" in US/MARC) is the primary focus of the description. If the hierarchical level of the target item is identified in the record, it is possible (1) to determine the relationship of any two hierarchically related records, or (2) to sort related records into a hierarchical arrangement.

Publication Pattern

Because material is published (or *not* published) in certain ways, it must be acquired, processed, and described in different ways—which in turn means that the records are structured and used in different ways. Monographic material is generally published once. Serials are continuously being published. Collections, although treated as single units, were published (or created) separately. Analytics (now called "component parts") are published as inseparable parts of larger units. Manuscripts might be defined as material that is *not* published. Another possible category that has distinct features is dynamically updated material, such as loose-leaf services or dynamic data files, where the content of the material is constantly changing.

Stability

Another way of looking at differences in publication pattern is to consider the stability of the bibliographic record itself. The records for completed monographs or dead serials are stable; no changes in the basic description should ever be required. Records for live serials, incomplete multi-piece monographs, dynamic data files, or archival record series, on the other hand, are always subject to amendment. All these types of unstable records present similar problems of control and thus constitute a significant category within a bibliographic database. Note that this distinction is similar to that described under *Publication Pattern*, but would draw the lines in slightly different places.

Originality

Bibliographic items are often related historically. They may be originals or reproductions of those originals. Reproductions may be *bibliographic* reproductions (such as a reprint or even a new edition of an original work) or *physical* reproductions made from an original master (either in the same form or in a different form). Physical reproductions also vary in exactness, the term *facsimile* indicating a particularly faithful re-creation of all the physical characteristics of the original.[11]

Uniqueness

Most bibliographic records are the surrogates for items that may be held in substantially the same form by many institutions. Certain types of material, however, are unique—either by their nature (manuscripts, art originals, handmade objects) or because their existence and identity are the result of actions subsequent to their creation (e.g., a group of pamphlets bound together or an archival collection). In such cases, the bibliographic record itself becomes the description of the holdings of a single institution.

Authority Records

Type of Record

In some systems, authorized headings are stored in authority records, while references to these headings are stored in separate records. Reference records may also contain general information notes about a name or title.

Type of Heading

The name or title controlled by an authority record may be characterized according to its type—personal, corporate or conference name, title, topical or geographical term.

Use of Heading

The heading controlled by an authority record may be related to bibliographic material in a number of ways—as a descriptive entry (author, editor, etc.), as a subject entry, or as a series entry. These categories are not exclusive, and the same heading may be used in various ways on different bibliographic records.

All of these criteria (and probably others equally significant) define categories of MARC records, which it might be useful to distinguish. Most of these categories may be identified more or less clearly from data elements presently defined in the MARC formats. Certain categories, indeed, have been selected as particularly significant and are identified by codes in the Leader of each MARC record, in a series of data elements called the Legend. It is from the Legend codes that the *format* of the record may be deduced.

THE IDENTIFICATION OF MARC RECORDS

Every MARC record carries in its Leader four character positions called the Legend. According to the *MARC Formats for Bibliographic Data* [*MFBD*] "the function of the Legend is to identify and describe the record. This identification and description implies (1) the content of the record, e.g., printed language material cataloged as a serial or manuscripts cataloged as a collection; and (2) the meaning of the content designators, i.e., in which MARC format the record is transcribed."[12] "The presence of these codes in the Leader was intended to assist users in separating records for various kinds of materials, e.g., only records for musical sound recordings are desired, because of different processing requirements, different collection policies, etc."[13] Thus the Legend is designed to provide the means for sorting records into categories by automated processing systems. The codes (1) provide *identification* of the type of record or the type of material described in the record, and (2) support *processing* of the record by defining the format (i.e., the set of valid content designators) and triggering special processing routines required for special categories of records.

The Legend is a fixed-length data element made up of four character positions. In the Authorities format, only one position is defined ("Type of Record") and only one code is valid ("Authority"). In the bibliographic formats, however, the situation is more complex and more problematic. Two character positions have been defined—"Type of Record" and "Bibliographic Level." In terms of the dimensions of record categorization described above, both are ambiguous.

"Type of Record" is a misnomer. Codes in this position identify the type of material described in the record. The data element is also ambiguous because it characterizes the material simultaneously by two distinct criteria—the intellectual nature of the material and its physical presentation. As a result, a large number of codes are required to identify combinations (e.g., printed music, manuscript music, musical sound recording, nonmusical sound recording). One significant category of physical presentation—microform—was originally included as a separate category. However, the code was removed when it was recognized that items of quite different intellectual natures and physical presentations could be reproduced in microform and that a large number of codes would be needed to characterize such material.

"Bibliographic Level" might also be considered a misnomer. In *MFBD*, no meaningful definition is given. The two codes most commonly used are those for monograph and serial—which suggest an attempt to categorize material according to publication pattern. However, the remaining categories—collection and analytic[14]—do not fit this interpretation. Indeed, this data element is intended to convey more than publication pattern. It has an interesting history.[15]

In the original *MARC II Format*, the definition of the Legend included a third function—to convey "the hierarchical level, e.g., a book which is part of a series, etc."[16] In an appendix, this function was explained: "It is frequently necessary to describe not only the item cataloged, but also all other bibliographic items to which the item is related. For example, in order to describe a journal article, it is also necessary to describe the journal. One way to accomplish this is to provide a separate subrecord for each bibliographic item described in the record. Each subrecord would be equivalent to a bibliographic level. The legend for a given record would indicate which levels were present."[17] The Legend was made up of five character positions, three of which were called Bibliographic Level and expressed "the relationships between the work cataloged and any other bibliographic entity of which it is a part"[18] by including a code for each level described in the record. Thus the original intent of the Bibliographic Level data element was to define hierarchical relationships.

This definition of Bibliographic Level was never implemented. Before the format was published, the Library of Congress in cooperation with the British National Bibliography made a number of changes that were included in a supplement issued with the MARC II document. One major change involved the Legend. It was reduced to four character positions, and Bibliographic Level was reduced to a single character indicating "the aspect of the work cataloged"—in other words, identifying the target item. Two positions were left undefined. "Although the term 'Bibliographic level' remained, the concept of relationship disappeared from the legend, and the legend became a description of the physical item, and remains so today in the United States MARC formats."[19]

Bibliographic Level is thus an attempt to identify the level of the target item of a record using categories such as monograph, serial, component part, and collection. Using these categories, it is often possible to determine the hierarchical relationship that exists between two related records. For example, the relationship between a monograph in a series and the serial of which it is a part or between a journal article (component part) and the serial that contains it will be clearly identified. However, the relationship between a subseries and its parent serial cannot be determined from the Legend (both are serials) but must be determined by examining the tag of the entry for the parent serial in both records. Thus Bibliographic Level does not explicitly identify hierarchical relationships in every case. Nor is this the sole purpose of this data element. Despite its name, it is most often used as an identification of publication pattern. This criterion is so significant in defining processing categories—material processed once (monographic), material processed continually (serial), material processed in aggregate (collections), and material processed as part of something else (component parts)—that it has become entrenched in Bibliographic Level. The result is a tension between two different uses for the same data element.[20]

Thus both elements in the Legend are ambiguous. Proposals to eliminate these ambiguities have been made. The most substantial effort to reorganize the Legend was inspired by "a UNESCO-sponsored project to design a common communications format at the international level to encompass the needs of the library and information communities and the abstracting and indexing services," as well as by a desire "for the orderly and logical expansion of the MARC formats to cover different types of materials issued in different forms and different levels."[21] Although the proposal was tabled, the issues remain and work continues as part of development of the common communications format. This work provides an opportunity to reexamine the universe of bibliographic and authority records, to identify the most significant characteristics that divide them into categories, and to try to resolve some of the ambiguities in the current definition of the Legend in the US/MARC formats.

THE DEFINITION OF A FORMAT

The format of a MARC record is deduced from the codes given in the Legend. An authority record is simply identified as such. Bibliographic records are categorized by two criteria—Type of Record, which characterizes the material described in terms of the intellectual nature of its content and its physical presentation, and Bibliographic Level, which identifies the publication pattern or other processing category of the record as an indication of its hierarchical level.

A MARC format is a distinct set of content designators prescribed for the identification and characterization of data in a particular category of records. The US/MARC formats comprise a family of such formats. At present, there are seven formats for bibliographic records—for Books, Films, Manuscripts, Maps, Music and Sound Recordings, Serials and Machine-Readable Data Files—and a format for authority records.[22]

These eight formats are the result of a historical process. MARC II was intended to provide a single format structure "capable of containing bibliographic information for all forms of material."[23] This information was to be broken down into individual data elements and identified and characterized by explicit content designators (tags, indicators, and subfield codes). "It was recognized that under ideal conditions the universe of material would be studied at one time for a more coordinated approach to the assignment of content designators. However, those responsible knew the magnitude of such a task, the time required, and the need for specialists to be involved. Consequently, to make progress in the near term, it was decided to handle one form of material at a time, beginning with books."[24] Additional formats have been developed whenever a need was perceived for a type of record not adequately covered by existing formats.

The criteria defining each format vary. Authorities is a type of record. Books is defined by the intellectual nature of its content (language material), its physical presentation (printed), and its bibliographic level (monograph or component part or collection). The Films format includes a variety of categories based on intellectual nature of content and physical presentation and excludes serial-level material. Maps is based on a category of intellectual content and includes all physical presentations and bibliographic levels of cartographic material (except atlases, which are Books). Manuscripts are a type of physical presentation (or a level—collections). Music includes two types of material, that which is musical in nature and that which is presented as a sound recording, and excludes serial-level material. Serials includes all serial-level material, except maps and data files. Data Files are a category of physical presentation (machine-readable). The situation is thus complex, perhaps even confusing.[25] To say that the Legend *identifies* the format of a record is an oversimplification. It is better to say that the format *can be deduced* from the Legend.

As formats multiplied and grew in complexity, it became increasingly difficult to maintain complete consistency in the content designation defined for the same data elements in different formats. This was partly due to the fact that format specifications for certain types of material (such as manuscripts, audiovisual materials, technical reports, and machine-readable data files) had been developed in consultation with, and were primarily used by, special communities. The material was kept separate from general library collections, different descriptive conventions were used, and separate bibliographic files were maintained. For these single-format users, consistency across formats had no great significance. In recent years, however, the trend has been toward large integrated databases, toward multimedia collections and catalogs, and toward consistency in bibliographic description for all materials. The International Standard Bibliographic Descriptions and the *Anglo-American Cataloguing Rules*, 2d ed. [*AACR2*] are monuments to this trend toward integration.[26]

Likewise, the advent of multiformat systems makes consistency of content designation across formats a definite advantage. Such systems perform most of their processing activities without regard to the format of the record. All records, regardless of format, must be read, indexed, retrieved, and displayed in the same manner. Any inconsistency in definition or usage of content designators means that exception routines must be invoked or else entirely separate processing routines for each format, with a great deal of redundancy, must be maintained.

An extreme example of existing inconsistency is provided in field 240. The first indicator in this field has been defined as a yes/no switch to enable/suppress display of the field and/or generation of an explicit access point (added entry). In the Books, Films, Manuscripts, Maps, and Serials formats, the indicator controls display. In the Data Files format, it controls access. In the Music format, it controls both. There is similar lack of consistency in the other additional title fields (210–240), both across formats and between fields. Such inconsistency in an element controlling such fundamental processes as display and indexing cannot promote efficient processing. Different types of records should not be treated differently in the MARC formats if they do not need to be treated differently by their users—either the catalogers who create the records or the systems that process them. Both types of users are moving toward integration.

Another factor points in the same direction. Since the format of a record is deduced from codes in the Legend, that record must belong to one—and only one—format. Yet the formats do not represent mutually exclusive categories of material. An item may be both music and manuscript, for example, or both map and serial. For such mixed-format material, a choice must be made, and this choice means that certain data cannot be included. The record for a serially published map can include all the cartographic fields (such as the mathematical data fields) or all the serial fields (such as the control information in the 008 coded-data field for serials)—but not both. The separate formats have erected barriers that are not always accurate reflections of the material they represent. Each record must fit completely within a single format, and each format must contain all the data elements needed for any record coded in that format. To an extent, it is possible to deal with mixed-format items by duplicating fields from other formats—by including serial fields in the Maps format, for example. At best, this practice simply expands the already extensive list of data elements common to several formats. At worst, it can add to inconsistency across formats if the same data element is treated in different ways.[27] Furthermore, this solution cannot be complete, because coded-data elements are defined separately in each format in a nonrepeatable field with the same tag (008).

Thus mixed-format material will continue to present difficulties as long as a record can belong to only one format and as long as formats are defined for separate but overlapping categories. Either the categories that distinguish formats must be redefined to eliminate any overlaps or a multiformat record must somehow be designed—or the separate formats must be merged into a single format.

To what extent should the formats be merged? The purpose of format integration is to promote processing efficiency through consistency and to break down artificial barriers between formats. All records that are created and processed in much the same way, which are basically similar in structure and function, should be consistently coded. Only one of the criteria listed above distinguishes categories of records that are fundamentally different in these ways—Type of Record. Authority and bibliographic records are routinely created, stored, and processed in different ways. They are so distinct in structure and function that their integration into a single format seems unwarranted.[28]

The term *format*, then, should be redefined as a distinct set of content designators prescribed for the identification and characterization of data in a type of MARC record that is distinct in function from other types of MARC records. Under this definition, there presently exist two such formats—for authority records and for bibliographic records. The seven bibliographic formats should be merged into a single format.

A SINGLE FORMAT FOR BIBLIOGRAPHIC RECORDS

The concept of a single format for bibliographic records is not a new one. It is embodied in several of the national implementations of MARC and particularly in UNIMARC. In the US/MARC formats, a significant step in this direction has already been taken—the preparation of the integrated *MARC Formats for Bibliographic Data*. This single-format document shows the large number of content designators that are consistently defined in all formats. In preparing this document, many needless variations in terminology and usage were eliminated. Remaining inconsistencies are easily identifiable from the format tables.

Although much has been achieved, *MFBD* does not yet constitute a single format. Such a format would contain a single set of content designators, each of which could be used in any bibliographic record where it was appropriate. Put another way, any data element required for any category of bibliographic material would be unambiguously identified and characterized by an appropriate content designator. This unambiguous relationship between data elements and content designation, independent of the category of material being described, has not yet been achieved.

Several sorts of inconsistencies remain. First, there are inconsistencies in *terminology*. For example, field 555 is defined as a Cumulative Index Note for serials and as a Finding Aids Note for manuscripts.[29] Second, there are inconsistencies in *usage*. For example, a note about languages is tagged 546 for serials and archives; for other categories of material, language information is given in a general note (field 500). Although these discrepancies are primarily historical accidents of format development, they often do reflect the different descriptive conventions that are commonly applied to certain categories of material.[30] Third, there are inconsistencies in the *definition* of the content designators themselves. The indicator in the 240 field, described above, is an instance of this.

The organization of *MFBD* provides a context for dealing with these inconsistencies. The documentation of a MARC field has five components: (1) the content designation proper—tags, indicators, subfield codes; (2) definitions—the name of the data element that corresponds to each content designator; (3) a table indicating validity in each of the present formats; (4) descriptive notes on meaning and usage of data elements; and (5) examples. The first step toward a single format is the elimination of the format table. If a single name can be given to each content designator, variations in terminology and usage appropriate to different material or descriptive conventions can be pointed out in the descriptive notes. Such notes can also describe inconsistencies in definition—at least until they can be eliminated. Work to remove the remaining inconsistencies from the format has already begun and must continue.

The result of these changes would be a single format for bibliographic records. It would be made up of two types of data elements—a core of common elements appropriate to all material and additional elements appropriate to particular categories of material or particular descriptive conventions. In general, the core elements would be those concerning the integrity of the record (Leader, Directory, Control Number), access (fields 100–130, 400–490, 600–655, 700–730, 800–830), classification (fields 050–082), and general description, much as it is specified in *ISBD(G)* and chapter 1 of *AACR2* (fields 245, 250, 260, 300, 500). The category-specific elements would be those required for special categories of material for description (e.g., fields 255, 362, 511) or access (e.g., ISBN, ISSN). The concept of content designators being *valid* only for certain categories of material would no longer apply. However, some elements would (by their nature or by an explicit practice note) be *appropriate* only for certain categories. Artificial barriers within the format not imposed by the nature of the material or by descriptive conventions would largely have disappeared.

There would also be subtle changes in the way the formats are maintained and developed. Any change in content designation would have to be applicable for all records in which the data element was appropriate. On the other hand, the format could be extended to new categories of material without going through the cumbersome process of drafting an entire "format" from scratch. A core of common data elements would already be present in the format. It would only be necessary to add descriptive notes as needed to explain the application of data elements to the new category of material and to include appropriate examples. The only noneditorial changes to the format would be the definition of content designation for any new data elements required.

The definition of the term *format* proposed here also has implications for the Legend. The identification of the format of a record, the set of valid content designators, becomes much simpler. Authority records are explicitly coded as such; anything else is a bibliographic record.[31] The principal function of the Legend then becomes the identification to the system of categories of records that have significance for processing. It is not at all clear that the categories presently coded in the Legend are in fact the most significant ones. As part of the review of the US/MARC formats and of the standardization of the Legend internationally, a rigorous analysis of the universe of

MARC records and the ways in which they might be categorized should be undertaken. The results must then be evaluated. Revisions to the Legend would have a profound effect on existing systems and records. On the other hand, such a revision may be the only way to provide an orderly and logical schema for categorizing records, which can apply both to present applications and to future expansions of the formats.

SUMMARY OF RECOMMENDATIONS

The most significant issue facing the authors of the US/MARC formats is the significance of the term *format*. The historical development of the formats has imposed artificial barriers that are not always reflected in the material described in the records. Formats have been built as separate boxes, with a great deal of redundant lumber, but also with a great deal of inconsistency in construction. An alternative concept of a format should be formulated. Such a concept will structure our understanding of the existing situation and provide a framework for analysis and decisions affecting future developments.

This paper has proposed such a concept. The term *format* should be defined as a distinct set of content designators prescribed for the identification and characterization of data in a type of MARC record. Separate formats should be defined for types of records that are mutually exclusive and distinct in function—such as authority and bibliographic records. As a corollary of this definition, *MARC Formats for Bibliographic Data* should be considered a single format. Any content designator should be valid in any bibliographic record in which the corresponding data element is appropriate.

In order to implement this concept, the following steps are recommended:

1. The definition of *format* given above should be included in the statement of the principles underlying the US/MARC formats.
2. *MARC Formats for Bibliographic Data* should be transformed into a single format. The format table should be eliminated and replaced by descriptive notes on variant practices.
3. Remaining ambiguities and inconsistencies in the definition of content designators should be eliminated.
4. The definition of the Legend should be revised to make clear that, within a format, validity of content designators is not a consequence of Legend codes. The function of the Legend thus becomes the identification of type of record (format) and of other significant categories.
5. The categorization of MARC records should be investigated and revision of the Legend considered.

ACKNOWLEDGMENTS

This paper was originally inspired by two proposals (since tabled) that came before the RTSD/LITA/RASD Committee on the Representation in Machine-Readable Form of Bibliographic Information (MARBI) and representatives of the national libraries and bibliographic networks in October 1980. An earlier version was discussed at the January 1982 MARBI meeting. I would like to thank my colleagues at these meetings, as well as Kaye Gapen, Dennis Reynolds, and David Weisbrod, for their comments and encouragement.

REFERENCES AND NOTES

1. "Principles of MARC Format Content Designation" (draft, Washington, D.C.: Library of Congress, 1981), p.1. This report of the review project was prepared by the Library from material submitted under contract by David Weisbrod of Yale University. The quotation is from the foreword by Henriette D. Avram.
2. Avram, in "Principles of MARC Format Content Designation," p.2.
3. A draft has been published: John Attig, "The US/MARC Formats: Underlying Principles," *Library of Congress Information Bulletin* April 23, 1982, 120–24; reprinted in *Information Technology and Libraries* 1:161–74 (June 1982). A revised version was approved on Oct. 29, 1982, and will eventually be published in the introduction to *MARC Formats for Bibliographic Data* [*MFBD*] (Washington, D.C.: Library of Congress, Automated Systems Office, 1980–   ), loose-leaf, updated quarterly.

4. Institutional responsibility is the subject of the Responsible Parties Rule, section 2.6 of the statement of principles. For the linking of related records, see Sally H. McCallum, "MARC Record-Linking Technique," *Information Technology and Libraries* 1:281–91 (Sept. 1982). The format specifications described will appear in Update no.8 to *MFBD*.
5. "Discussion Paper No. 5: Topics for Future Consideration by MARBI," Library of Congress, Automation Planning and Liaison Office, March 20, 1982, revised Oct. 29, 1982.
6. Weisbrod, in "Principles of MARC Format Content Designation," p.65.
7. Some of these criteria and some of the terminology used were suggested by Weisbrod, ibid.
8. A third type of record—a local holdings record—is currently being defined: "Proposed MARC Format for Holdings and Locations: Working Draft," Southeastern ARL Libraries Cooperative Serials Project and Library of Congress, 9/15/82. A holdings record might be thought of as a bibliographic record viewed from the perspective of a single holding institution.
9. This is particularly important if it is recognized that the MARC formats are not confined to any single set of descriptive conventions, but should rather serve as a vehicle for bibliographic and authority information of all kinds, for all users. Many data elements in the MARC record try to convey information about descriptive conventions and level of completeness—Encoding Level and Descriptive Cataloging Form, for example. For users of *AACR2* contributing to a national database, various National Level Record standards for full and minimal records have been published by the Library of Congress, and field 039 has been defined to carry information about the level of the record with regard to these standards. As most bibliographic networks have their own content standards, this attempt to create a national standard and records defined in terms of that standard has not been an outstanding success. Outside the circle of library conventions, the situation is considerably less well defined; see note 30 below.
10. A code has been defined in the Leader (byte 19) in the bibliographic formats, to indicate that a linking-entry field (760–787) contains only the control number of the related record and that the relevant data must be retrieved from that record. McCallum, "MARC Record-Linking Technique," p.282–83.
11. The ambiguity of bibliographic and physical reproduction makes it difficult to categorize material. In some cases (e.g., films and sound recordings), the desire is to distinguish master negatives or recordings from the copies produced from them. Using this same criterion, however, all books are reproductions of the printing plates from which they were produced. Until agreement on a useful concept of reproduction has been achieved, *MFBD* will advise the use of a fill character in field 007/2 (Original versus Reproduction Aspect).
12. *MFBD*, Leader/06(*), p.1.
13. "Proposal No. 80-2: Revision of MARC Legend," Library of Congress, Automation Planning and Liaison Office, 10/7/80, tabled 2/2/81, p.1.
14. The code for "analytic" has now been broken up into two codes—"component part, monographic" and "component part, serial." McCallum, "MARC Record-Linking Technique," p.287. Also, as part of the revision of the Manuscripts format, a new value will be added for a subunit of a collection (which may itself be a collection).
15. Most of the following discussion is based on Henriette D. Avram, "Recommendations towards the Establishment of a Common Exchange Format for Use by All Agencies in the Information Community," in H. Dierickx and A. Hopkinson, eds., *Towards a Common Bibliographic Exchange Format?: International Symposium on Bibliographic Exchange Formats, Taormina, Sicily, 27–29 April 1978* (Budapest: OMKDK-Technoinform, 1978), p.135–46. The substance of the following section was previously presented in John Attig, "Discussion Paper No. 6: Hierarchical Relationships in the Legend," Library of Congress, Automation Planning and Liaison Office, August 29, 1982.
16. Henriette D. Avram, John F. Knapp, and Lucia J. Rather, *The MARC II Format: A Communications Format for Bibliographic Data* (Washington, D.C.: Library of Congress, 1968), p.14.
17. Ibid., p.147.
18. Ibid., p.42.
19. Avram, "Recommendations," p.139. Another casualty of this redefinition of the Legend was the multilevel record consisting of subrecords describing each level involved. US/MARC has never implemented a subrecord technique, a decision recently confirmed after a long and difficult struggle with the problem of analytic cataloging. It was decided that full descriptions at each level should be maintained as separate records. These records may be linked either implicitly through entries (as is done for series in the

case of a monograph in a series) or explicitly through a linking field that contains a *citation* of the related item (an entry, a brief description, or the record control number of the full description). See McCallum, "MARC Record-Linking Technique."

20. In the UNIMARC format, Bibliographic Level and Hierarchical Level are separate data elements: *UNIMARC: Universal MARC Format* (2d ed. rev.; London: International Federation of Library Associations and Institutions, 1980), p.12.
21. "Proposal No. 80-2," p.1. This proposal recommended separation of Type of Record into two data elements—Type of Entity and Type of Physical Presentation—and a clearer definition of Bibliographic Level and its categories.
22. These formats are published in two format documents: *MARC Formats for Bibliographic Data* (1980) and *Authorities: A MARC Format* (Washington, D.C.: Processing Services, Library of Congress, 1981), 116p. The format for machine-readable data files will be published in Update No. 9 to *MFBD*. Two of the formats are undergoing revision. The Manuscripts format is being revised to accommodate archival collections, and specifications for two- and three-dimensional graphics are being added to the Films format. In addition, a format for holdings and location information is being developed (see note 8 above).
23. Henriette D. Avram, *MARC, Its History and Implications* (Washington, D.C.: Library of Congress, 1975), p.7.
24. Ibid.
25. This same confusion of criteria can be seen in the chapters for description of particular types of material in *AACR2*, and in the lists of General Material Designations (rule 1.1C1). *Anglo-American Cataloguing Rules* (2d ed.; Chicago: American Library Assn., 1978), 620p.
26. Although it comes from a different bibliographic tradition, the *American National Standard for Bibliographic References: ANSI Z39.29* (New York: ANSI, 1977), 92p., also aims at consistently formatted references for print and nonprint material of all sorts.
27. This actually occurred when the element for Frequency was included in the Maps format. Whereas the Serials format provides three fields—310 (Current Frequency), 320 (Current Frequency Control Information), and 321 (Former Frequency)—the Maps format defines a single field (315). This single field was also included in the Data Files format. This inconsistency makes it difficult to deal with these records together in the context of a serials control system or union list—which is presumably why the serial data was included in the first place.
28. Other possible types of records, such as address records or acquisitions fund records, are also clearly distinct. The local holdings record, on the other hand, is less distinct and seems to be considered a separate format as a matter of convenience. Holdings records are related to a single bibliographic record. The bibliographic record is a union record containing information about all holdings; the local holdings record contains data elements for information specific to a particular institution's copy or copies; and several data elements are included in both records as a link.
29. Actually this inconsistency has already been resolved in *MFBD* by combining the names: Cumulative Index/Finding Aids Note. The fact remains, however, that different terms are used for this data element when applied to the different categories of material.
30. For example, field 740 is not defined in the Serials format because serials conventions call for variant titles to be characterized explicitly as provided by the indicators in field 246. Variations in usage become even more pronounced outside the scope of *AACR2* and related content standards. It has become apparent in recent years that not all MARC users fall within the scope of library practice. Descriptive conventions for audiovisual materials, technical reports, and archival collections have all been accommodated (more or less happily) within MARC. The result is that some data elements are appropriate, not simply to particular categories of material, but to particular descriptive conventions, and that discrepancies in the usage of data elements abound. This underlines the fact that MARC is *not* a content standard and that content standards are in fact needed in order to maintain a usable file of records from diverse sources.
31. The situation will become somewhat more complex with the introduction of a format for holdings. This might simply be another type of record, like Authorities, which is not further categorized. On the other hand, one recurring comment on the draft format was that it should be possible to communicate a merged record, i.e., a bibliographic record that includes data elements for institutional location and holdings. Which type of record this would be and how it should be identified in the Legend have not yet been decided. ■■

Integrating Library and Book Trade Automation

Sharon C. Bonk

Librarians rely on booksellers to provide the link between publishers and libraries. Increasing automation in the book trade and in library acquisitions systems will lead to a new type of bookseller-library interface and relationship. Electronic order transmission has begun on a limited scale in both the United States and Great Britain. Librarians must reexamine their requirements of booksellers in the context of the new possibilities automation will provide.

Librarians rely on the services of booksellers in purchasing materials for their collections, and the bookseller is an important link between the publisher and the library. The nature of that link is likely to change significantly with the advent of integrated library and bookseller computer systems. Up to now, in both Britain and the United States, specialized library booksellers have formed a distinct part of the book trade. Library booksellers compete with publishers and retail booksellers for sales of publications to libraries. This competition is based on discounted prices, extra services, or both. In Britain, prices are fixed by the publisher under the Net Book Agreement, which does not permit any discount other than the 10 percent allowed to public and other Publishers Association–approved libraries which permit use of their books by people outside the institution.[1] In the United States, discount is a major factor in attracting a library's business. In both countries it is the ability to provide the books quickly and provide extra services efficiently that keep library customers. Specific booksellers are selected by each acquisitions librarian for empirically demonstrable reasons and perceptions of quality service.

The textbook reason for use of library jobbers, other than externally imposed directives by funding authorities, is rationalization of effort.[2] By placing numerous orders with a limited number of reliable booksellers, librarians eliminate or reduce a measurable amount of paperwork which booksellers take on. Librarians have fewer sources to deal with in order, follow-up, and payment. The jobber takes on the work of prepayment, claiming, and combined shipments and invoices, and presents the library with books, invoices (to library specifications), and reports on titles not currently available from the publishers. Booksellers have also added services and functions that are related but not directly part of the book purchase transaction: supply of catalog cards; supply of bar-coded labels for circulation systems; physical preparation of books for immediate use upon delivery.

ELECTRONIC ENVIRONMENT

The library-bookseller relationship described above is the familiar one. In the future, libraries should be able to receive new bookseller services in electronic ordering and the related confirmation of orders, status reports, and invoices. These services

Sharon C. Bonk is head of acquisitions at the University Libraries, State University of New York at Albany. Manuscript received July 1982; accepted November 1982.

will result from the use of automated order processing and inventory control within the book trade and the response to the use of computers for order generation and follow-up by libraries. Work has begun on exploring the gap between the library and book trade systems. Meetings have been held between librarians and booksellers on the desirability of interfacing systems.[3] Sandra Paul has reported on the activities of the American book trade that have implications for library acquisitions systems.[4] Julian Blackwell has brought a similar message to British librarians.[5] The Book Industry Standards Advisory Committee (BISAC) and the American National Standards Institute Z39 subcommittees have developed and are reviewing and revising format standards for electronic transmission of purchase orders, title updates, invoices, and other transactions. These formats are designed to fit the needs of publishers, booksellers, wholesalers, and librarians. Each group has representatives on the committees, and public review of their proposed standards is invited.[6]

Thus, the present activities and conditions in the book trade and in libraries are creating the environment out of which progress will be made toward electronic telecommunication between libraries and their suppliers.

1. Libraries are increasingly automating their traditional book and serials ordering, receiving, and fund accounting functions and linking them interactively with other internal library systems which are already automated or in the planning stage for automation.

2. Library booksellers are also in the process of installing, if not already in possession of, interactive order, receipt, and accounting systems which are designed to support customer service as well as meet bookseller internal needs.

3. Publishers in the mid-to-large range are installing online systems to link their production, order, distribution, and accounts information.

4. Through cooperative industry efforts, publishers and booksellers in Britain and America are beginning to provide for book and serial ordering by electronic means.

5. There exist in machine-readable form the components of a bibliographic chain of information useful in book order and purchase transactions which may be transferred between the groups by telecommunications methods.

6. The state of technology and software development is such that it is possible to provide interfaces between systems which would make this interchange of bibliographic and other data possible.

Progress toward shared data in electronic form can be made only if there are mutually beneficial reasons for librarians and booksellers to implement the links between data handling systems. Those mutually beneficial areas are:

- Identification of publication availability
- Access to databases of current book information
- Ease of handling data entry
- Decrease in preorder search effort
- Receipt of "clean" unambiguous orders
- Reduction or elimination of duplicate effort
- Speed of processing
- Speed of delivery to libraries
- Increased stock turnover
- Increased service
- Increased sales
- Staff savings

Since most book trade–library interfaces are in test phases or are operating at the minimal level, the empirical data necessary to prove the mutual benefits to be derived do not yet exist. Therefore, these mutual benefits are posited from examples and observations of systems in operation, plans for enhancements, and discussion with those in the process of planning and developing.

LIBRARY REQUIREMENTS

Direct telecommunication with booksellers has been listed or implied as desirable by librarians in both countries.[7] In a summary of his interviews with librarians, Boss notes that online ordering was mentioned frequently as one of the most important characteristics of automated acquisitions systems. Public librarians, especially, wanted to speed the receipt of titles. Boss notes the contradiction that exists. The high volume of ordering that makes online ordering attractive also presents the problem of requiring that a wide variety of vendors be able to receive and transmit information. This is the limitation of systems in which a library is linked electronically to one bookseller. The three large booksellers supporting these systems, Baker & Taylor, Brodart, and Blackwell's (Oxford), provide for generation of paper orders to other suppliers. The smaller booksellers who are currently working on electronic ordering systems with libraries provide only for direct order transactions to that bookseller. Obviously, the bookseller is interested in providing for increased orders to his firm at the expense of competitors. However, most librarians do not wish to limit their options by making contractual agreements and investing in equipment with one bookseller. Librarians who find that one vendor can supply the majority of its acquisitions would welcome this type of interface. In general, librarians do not wish to acquire different hardware for each system and use several different sign-on procedures and search keys for order placement to multiple vendors. Therefore, most librarians will be cautious in their approach to bookseller systems currently available.

The existence of large library networks with acquisitions subsystems makes the centralization of electronic order generation possible. Recent developments by the WLN and UTLAS networks provide acquisitions subsystem users electronic transmission of orders to vendors who wish to receive the orders in this form.[8] OCLC has announced this feature but not a systemwide implementation date. It is likely that RLIN will also consider the provision of this component in its acquisitions system.

This concentration of source of orders will surely be beneficial to vendors by reducing the member of possible interfaces that many individual library systems would require. However, local conditions and costs may dictate that a library not acquire its materials through a network. Yet, it may wish to establish electronic communication with vendors.

What is required of book-trade interfaces?

1. Integration with other library systems is imperative. Any transfer of data in machine-readable form should integrate with other parts of the bibliographic and financial control systems of the library and its funding body. The same information should not have to be processed twice by the library in order to record information for internal use and again for bookseller use.

2. Libraries, including those using network acquisition modules, should benefit from use of book-trade databases without the penalty of duplicate or excessive effort. As large and as comprehensive as bibliographic utility databases are, they do not contain purchase information: current price; in-print status; information on sales of multivolume sets; bindings; special prepublication offers. Obtaining this information is still a manual process completed before use of the automated acquisitions system. It is here that a book-trade interface through the network could contribute much to the integration of preorder work into the library's automated acquisitions system.

For some libraries, whose acquisitions are generally current imprints from standard publishers, bookseller databases could provide one source for bibliographic information, identification of availability, price, and an indication of probable delivery time by stock or printing status. This would reduce and, ideally, eliminate duplicate searches and multiple transcriptions of data and, thereby, be even more attractive than a network database.

3. An online file of all orders placed through the system should be accessible to the library and bookseller for updating or other modification. Reporting on status of orders and claims should be accomplished by the bookseller initiating transmission of information. This would eliminate the printing, batching, and mailing of reports by booksellers and the posting of such reports by libraries.

4. Librarians would like copy-specific inventory information from the bookseller. Use of this data would not necessarily be an essential part of daily library ordering but would have specific application for rush orders or stock editing. In general, knowledge of availability at time of order would free

funds encumbered for "unavailable" publications.

5. Online communication may not be desirable for all aspects of library bookseller systems, especially while library systems are not fully online. Libraries will accept tape and disc records of their orders and invoices until computer links are established and costing is done to indicate where savings will occur in the use of telecommunications to exchange data with vendors.

BOOK-TRADE RESPONSE TO LIBRARY REQUIREMENTS

Booksellers as a group do not have common spokesmen, so it is difficult to find published or public views that speak collectively for the businesses that are in competition with one another. However, it is observable from services and products offered and bookseller internal order processing systems that their concerns are for decreased order turnaround time, lower per-item costs of processing, and increased volume of business with static or decreased staff. From presentations and in individual discussion and communications, one can conclude that library booksellers are aware of the library's increasingly sophisticated service needs and are preparing for or are already in position to offer new services related to the book purchase transaction in electronic form. These services are or will be tailored with the particular bookseller's needs very much in mind.

Supply of traditional printed catalog cards has been updated by provision of cataloging data in machine-readable form. Computer-printed invoices to customer-specified requirements now may include magnetic tape options. Many booksellers would gladly drop paper reports in favor of a method of more efficiently updating information for customer use. Many of these services stem from the booksellers' internal order processing systems in place. Certainly, customized billing, reporting, and statistics lend themselves to production from a computer file. The ability to manipulate data and deliver related services is an area for competition and one which has great potential.

Several library booksellers have pondered their role in a telecommunications age. Others wish to know the ways libraries view use of bookseller databases.[9] In general, booksellers seek this information so that they may add to or redesign systems to offer these services and keep competitive. The larger library suppliers will respond to librarians' requests to provide interactive computer links if the specifics of these requests are not unique to an institution and do not interrupt the internal operations upon which a bookseller's system rests. The middle-to-small-sized library booksellers may be slower to react to library needs since they do not typically have the resources to formulate large systems and market them as services or products for purchase.

Booksellers have begun to respond to library requirements, as outlined above, in the following ways:

1. Booksellers recognize librarians' concerns for integration of systems. At present, libraries must provide integrating software to link the bookseller data with the library's bibliographic and financial systems. Libraries might agree to purchase order processing and fund accounting software developed by the bookseller to use on a library's own equipment. Payment might be by outright purchase, per record use, or by increased number of orders. How far booksellers are able or willing to pursue this varies considerably within the library bookseller group.

Baker & Taylor offers a system (LIBRIS II) that provides for order placement, fund accounting, and receipt of approvals and standing orders. Those orders and approval plan volumes supplied by Baker & Taylor may be processed online without need for paper orders, follow-up, and invoice entry online. Baker & Taylor also can supply cataloging data and tape services. Brodart's order system (OLAS) provides order placement directly through the terminal and fund accounting. Brodart can also supply cataloging data in card or machine-readable form.

Other booksellers are working on this order-purchase function. The Blackwell's (Oxford) database supports the library acquisitions system of one of the English county library systems. Blackwell's and the library have built the interfaces between two main-frame computers, and as such, the system in operation is not designed for exact replication. However, Blackwell's has developed software for library computers for order files and performance of fund accounting.[10] Blackwell North America, long a source of cataloging data as well as a supplier of books, is planning to develop interfaces with the systems of its customers and the networks.[11]

Askew and Sons and Woolston and Blunt, two large British library suppliers, have online order receiving systems that allow the libraries to dial into the bookseller's database from a standard terminal and, after searching the database, to place orders for books. At present, the messages are one-way and an open order file is not held for or by the library. Both companies are looking to their customers to guide them in developing useful services, expanding this order transmittal system into a two-way message system which might eliminate the need for traditional paper orders and paper files.

A source of integrating software may be the acquisition modules under development by circulation system vendors. Currently, no circulation system acquisitions module is far enough along in development and implementation to interface with the book trade in electronic form, although CLSI and DataPhase have agreed to incorporate BISAC order formats in their acquisitions systems. A different approach is offered by the Ringgold Nonesuch acquisitions system, which will incorporate the book industry standards. This system leaves the catalog and/or circulation interface to the library to develop.

2. Little work has been done on the feasibility of networks utilizing bookseller databases. Networks and their members have not seriously considered bookseller system capabilities. Since two of the American library booksellers who have large databases of current book information also offer acquisitions systems in competition with the networks, they will probably not be in the lead of integrating the two systems for the benefit of the library community. However, booksellers of varying sizes are represented on industry standards committees and do work with the library community toward resolution of problems.

3. The technical capabilities for booksellers to meet library requirements exist. Increased capabilities of terminals will allow libraries and booksellers to acquire equipment that can be programmed for data transmission using different protocols. Booksellers with computer systems should and will be able to accommodate either direct or dial-up access using industry/library standards. However, since each bookseller's goal is to provide services that will result in increased orders, booksellers would prefer to have an exclusive link to a library. What should be incorporated into their planning is the realization that it is to the booksellers' collective advantage that a standardized approach to their systems be followed. When the ability exists for a library to communicate directly or through a network with multiple booksellers, then librarians' collective interest in these systems will increase. If the library has the software and capacity to receive machine-readable transmittals from the booksellers, booksellers who use automated methods to generate orders to publishers should be able to transmit order confirmations, reports, and updates to the library in book-trade standard formats.

4. No library supplier has true stock level online for customer access. Machine-readable codes in books will permit this to be done more easily and more comprehensively, but booksellers are reluctant to provide this information and doubt that it will be of general use. Brodart has a speeded order fulfillment system which is linked to its inventory, but it has not incorporated a link to its acquisition system. Many booksellers believe that the cost of making this information available for the few who would regularly use it cannot be balanced by use of the information or increased orders as a result of the information.

5. Economic considerations are obviously of concern to booksellers as well as libraries. Some booksellers see declining library budgets as a reason not to get involved beyond the basic level dictated by their competitors. Others see savings to the libraries as the most important reason for these systems. Booksellers with operating systems see financial benefits in the form of savings from direct order systems and are

looking to produce statistics for their customers to compensate for the time-consuming and costly statistical analyses done in some libraries or not compiled at all. Booksellers will offer a variety of modes of access and levels of use of their systems to satisfy large and small customers. Some booksellers may develop a turnkey package that will be cost-effective for some types of libraries.

6. When bookseller and library systems are designed for data exchange in machine-readable form, much paperwork and most files can be eliminated for each user. Use of records already created by booksellers saves the library and bookseller time and eliminates transcription errors. This will eliminate duplicate efforts, provide information on orders and expenditures, and reduce personnel and paper-processing costs. These savings will, in turn, need to be applied to the cost of implementing machine-based systems, and, in the long term, will mean true savings for libraries and booksellers.

Ironically, many library suppliers have greater technological capabilities and flexibility in software and systems than the libraries they serve. Also, in contrast to librarians, many vendors are able to assess costs of specific operations. They are able to see where profit and service interact or diverge. Since library booksellers are generally reactive in terms of developing and marketing these computer-based services, most will not formulate systems, like those of Baker & Taylor and Brodart, nor market them because of the comparatively small size of the companies and the amount of financial investment needed. Larger library suppliers will not respond to librarians' requests to provide interactive computer links with the library if the specifics of these requests are unique to the requester and not cost-beneficial to the booksellers. Almost all library booksellers are open to adapting or redesigning systems in order to maintain a competitive edge. Therefore, they are interested in what librarians' plans are for use of machine-readable records and their desires for electronic transmission. But as a group, vendors do not see their role as one of leadership but rather as one of support.

Retail and specialist academic booksellers who also supply libraries are in a different category, for most do not have databases or large order-processing systems that can spin off library products such as book lists, customized invoices, book labels, catalog cards, or machine-readable order records. At present, the most automated would probably only be able to accept an order directly into their systems and forward it to the publisher. Herein lies a dilemma. The academic librarians with the most sophisticated needs of bibliographic records, availability information, nontrade publications, integration of serials into their databases, and analysis of orders and accounts, deal with large numbers of booksellers who are least in a position to provide any machine-based services. Therefore, library planners do not look toward them for solutions or services. This in turn reinforces the booksellers' thinking that these services are not required.

Not all library booksellers will be able to compete equally in providing expanded computer-based services. The market is now structured by type of library, and it will continue to be so. Some booksellers will become full-service operators and others will offer a more limited range of services. This is not necessarily a disastrous situation, for all libraries will not require the same level of automated services. Many libraries will be content with a system wherein the library holds its order file on microcomputer and sends and receives messages through it. Likewise, the utilities' systems may develop to the stage of integrating their systems with that of the library suppliers who may be required by customers only to receive and send messages via the acquisitions module. Other libraries will see the benefit of the bookseller providing software for the library or handling some functions at a fair price and saving the library the cost of programming and systems development.

Any mistrust of a commercial motive may need to be reexamined by librarians who are looking for ways to stretch their resources. If booksellers can profitably offer services and software, they should not be automatically suspect, and the services be rejected as commercial enterprises, especially if the library finds the nature of the services appropriate to the library and does not have the staff to provide the software or function itself. Pursuing this approach, two library network staff members have outlined the mutual benefits to be derived from cooperation between commercial system suppliers and library networks.[12] Although they do not discuss book supply systems specifically, their views are that libraries, networks, and commercial firms are all engaged in internal planning based on each other's capabilities and needs. Librarians are planning the use of commercial and/or network services within their own contexts. Networks are as interested as commercial vendors in what is desirable and acceptable for libraries, in determining the libraries' priorities, and in focusing competing perspectives of need and difficulty of design and implementation. Networks and commercial systems are competing to supply services to libraries. Open three-way discussion may lead to networks providing planning, technical, and financial support to member libraries for book-trade interfaces in a network context.

DESIGNING INTERACTIVE SYSTEMS

Librarians should analyze costs of their current systems, whether manual or automated, and see where savings in time, equipment, and personnel would be possible if processes were automated, or upgraded from batch to online, or if links to external data files were made. Librarians must be able to recognize any time, financial, and personnel savings in standardized software or other bookseller services and be able to view telecommunication charges realistically against benefits derived. They must identify the value they place on increased access, control, and information that a fully automated acquisitions system can provide.

Serious consideration should be given to what the present capabilities of booksellers are for receiving orders in electronic form and what specific plans each bookseller has for the future. Although a tradition of good service is valued and commendable, it should be viewed as only one part of the decision as to why a particular bookseller will continue to be used. There must be a change from historic to future considerations, and choice of bookseller should be made from an enhanced definition of services that are made possible by the computerization of the book trade and advances in telecommunications. Will the library be able to develop a two-way communications system with the company? Will machine-readable records be available as an interim step until the library has the capability to receive bookseller transmissions?

Along with consideration of internal needs, there is a need to use the standards for content and format of machine-readable records and work with the book trade on standards for technical interfaces. At the most basic level, librarians must become aware of book industry changes and incorporate standards such as ISBN and SAN.[13] Certainly, the demand by librarians for unique interfaces will be more problematical to library suppliers than the traditional requests for processing of orders, for the financial investments are greater. Independent systems and turnkey systems will continue to be part of the scene. Their external communications in electronic form should be subject to library and book-trade standards so that collective gains will not be subverted by individual libraries' requirements or demands on the booksellers.

Since all librarians would agree that booksellers supply service as well as books, the onus is on librarians to define the services in this area just as they have done in the past, but with greater emphasis on standardization procedures and data. Since most librarians would surely agree that order transmission is not the basis around which a library develops its acquisitions procedures, the ability to dial up a number of library booksellers and enter orders via a terminal is not integral to the system and may even be viewed as only a technical novelty if that is the only use of that technology. Thus, the one-way order transmission systems we see in use today are or should be just the beginning and testing phase to determine problems and areas for expansion.

REFERENCES

1. "Publishers Association: Net Book Agreement (1957)" and "Publishers Association: Library License (1957)," in *A Librarian's Handbook* (London: Library Association, 1976), p.548–50.
2. Stephen Ford, *The Acquisition of Library Materials* (Chicago: American Library Assn., 1973), p.58–60.
3. "Vendors and System Users Form Mutual Interest Group," *American Libraries* 14:196 (Mar. 1982); "Automated Book Ordering: Proceedings of a One Day Seminar" (Hertford, Eng.: Hertfordshire Library Service, Nov. 1980).
4. Sandra K. Paul, "Emerging Standards for Book Acquisition," *Technicalities* 1:8–9, 11 (Mar. 1981); Sandra K. Paul, "Computer-to-Computer Communication in the Acquisition Process," *Journal of Library Automation* 14:299–303 (Dec. 1981); "Utilities Agree to Use BISAC Order Format," *American Libraries* 14:197 (Mar. 1982).
5. Julian Blackwell, "Publishing v. Bookselling—The Shape of Things to Come," in *Nationwide Provision and Use of Information: Aslib/IIA/LA Joint Conference* (London: Library Association, 1982), p.248–56.
6. Charles C. Stewart, "Update on Ordering Standards," *Information Technology and Libraries* 1:341–43 (Dec. 1982).
7. Kenneth J. Bierman, "Vendor Systems and On-line Ordering," *Journal of Library Automation* 13:170–80 (Sept. 1980); Richard Boss and Deanna Marcum, "On-line Acquisitions Systems for Libraries," *Library Technology Reports* 17:122 (Mar./April 1981); Martin Dudley, "Annual Seminar," *MARC Users' Group Newsletter* (Fall 1981).
8. Bruce Ziegman and Brian Aveney, "WLN Online Order Transmission," *Information Technology and Libraries* 1:346–48 (Dec. 1982); Richard Ellis and Leslie Straus, "Electronic Ordering at UTLAS: A Chronicle of Library/Book Vendor/Bibliographic Utility Cooperation," *Information Technology and Libraries* 1:343–45 (Dec. 1982).
9. Dana Alessi, "Book Jobbers—There Yesterday, Here Today, Gone Tomorrow?" *Library Acquisitions: Practice and Theory* 5:21–25 (1981); William Connolly and Philip Hoban, "Annual Seminar," *MARC Users' Group Newsletter* (Fall 1981).
10. Blackwell Technical Services, "Bookline/Perline" (Oxford: Blackwell's, 1981), 9p.
11. "BNA Plans Interfaces," *Library Systems Newsletter* 2:3 (Jan. 1982).
12. Ronald Diener and Carl A. Anderson, "Organizational Issues for Library Automation: A Network Perspective," *Technicalities* 1:3–4, 11 (Mar. 1981).
13. James K. Long, "Electronic Order Transmission," *Journal of Library Automation* 14:295–97 (Dec. 1981). ■■

Access to Periodicals: Search Key versus Keyword

Susan U. Golden and Gary A. Golden

As more libraries replace their card catalogs with online systems, the searching algorithms of those systems become important to patron success in retrieving items. This research compares the retrievability of periodical titles in a fixed-length algorithmic search (Library Computer System) to a keyword search (Washington Library Network). Because serials often have nondistinctive or similar titles, they present a unique retrieval problem. One hundred and fifty-two periodicals were studied to determine which type of search algorithm is most successful with titles of varying lengths.

The development of online catalogs is big news in library automation for the 1980s. As the cost of keeping a large card catalog up-to-date skyrockets, it is not surprising that institutions are turning to online catalogs as a way to plan for the future. Many libraries, having already developed online circulation systems, feel the next logical extension is an online catalog. With the advent of this new breed of library catalog comes the age-old problem of serials control and access. Serials control is an area where several systems are making advances—but what about patron access to serials for searching online? How successful is a patron when searching serials online? In a research library, searching for periodicals can be more common than searching for books. Because of various complicated rules of entry for serials in the cataloging rules of the past, patron success often requires perseverance and luck. The question is, Will online catalogs provide any improvement over the traditional card catalog for access points and/or search strategies?

When a user evaluates a service, he considers cost, time, and quality of the database.[1] A user of a card catalog, be it traditional or online, is likely to base his evaluation on the two latter criteria since direct cost to the patron is usually not a factor. As a patron searches a catalog, he is interested in finding what he wants quickly and with the least effort. If a patron enters the library with a specific citation to a book or a serial, he approaches the terminal and queries the system expecting a quick, direct response that will satisfy his need—is the item owned by the library or not. The file structures of most systems are such that all bibliographic material is treated alike. That is, serials are not stored separately from other types of materials. They may in fact be tagged in some fashion, such as fixed-field MARC tags for bibliographic level, but storage and searching are likely to be in one bibliographic file. This means that unless a patron can limit a search by type of material (e.g., book or serial) when a search is initiated, the computer must search and print out all titles matching that search code. Since periodicals often have short titles like *Time*, *Psychology Today*, *Look*, and *Life*, serial searching can be cumbersome. Some systems do have this ca-

---

Susan U. Golden is assistant director of General Services for Automated Systems and Gary A. Golden is Central Information Services librarian, University Library, University of Illinois at Urbana-Champaign.

pability to limit searches by type of material. However, many do not have this feature.

The purpose of this study was to compare two systems without such a limiter, but which have two very different search algorithms. It is expected that studies such as this one will have an effect on the design of searching algorithms for online catalogs. Will traditional search codes be sufficient, or will more sophisticated methods such as the capability for noting adjacency be needed? Since both the Library Computer System (LCS) and the Washington Library Network (WLN) are, or soon will be, used for public access to bibliographic records, a study of the effectiveness of their title searches is both timely and necessary.*

DATABASES COMPARED

This study, conducted in mid-1980, takes two online systems and compares their searching methods. The first, an online circulation system and known-item short record catalog, is the Library Computer System (LCS) as it exists at the University of Illinois at Urbana-Champaign. The second is a full bibliographic record online cataloging utility network—the Washington Library Network. At the time of this study LCS had public access terminals, but WLN did not. Several organizations are planning to use modified WLN software to provide the basis for online catalogs (because of its sophisticated searching capabilities), so it will be a public access system soon. OCLC was not included in this research for two reasons. First, while the public does use OCLC to verify titles and to borrow via interlibrary loan, it was not believed to be the same use as an online catalog. Secondly, OCLC has the capability to limit searches by material type; WLN and LCS do not.*

*The Library of the University of Illinois at Urbana-Champaign has purchased the WLN software to provide the full bibliographic record portion of its online catalog. The authors searched the WLN database through dial access during the planning stages for UI's online catalog.

*After this research was completed, LCS obtained the capability to limit a search to a serial or a monograph. A follow-up study using these limiters is planned.

The title search in LCS is a code consisting of a command followed by the first four letters of the first significant word of the title followed by the first five letters of the second word. Hence, *Foreign Affairs* becomes "TLS/FOREAFFAI." A list of stopwords applies, and these words should be ignored when constructing searches (see appendix A). A search for the *Journal of Social Issues* becomes "TLS/SOCIISSUE."

The title search in WLN is constructed by a command followed by all significant words in the title in any order. A search for *Foreign Affairs* would be done by "find title *Foreign Affairs*" or "f t *Foreign Affairs*" and produces a list of all items with those two words anywhere in the title or subtitle in any order. A search for the *Journal of Social Issues* in WLN would be done by "find time *Journal Social Issues*." Full words are searched when the exact title is known. As with LCS, a list of stopwords applies (see appendix A).

Because of differences in search strategies, several things need to be noted. First, LCS can distinguish one-word titles. Hence the search "TLS/TIME" searches for matches on "TIME" as a one-word title. The absence of a second five-letter keyword indicates the absence of a second word. WLN, with a similar search "f t *TIME*," will look for the word anywhere in the title or subtitle as a one-word title or in combination with any number of other words. Secondly, LCS cannot distinguish a full word from a truncated word. A search for the *American Archivist* (TLS/AMERARCHI) produced forty matches, among them *America in Archives*, *American Architecture*, and *Amerikanisches Archiv*. Thirdly, LCS searches the first two words in the title in the order submitted, with no other words in between. Hence, an adjacent situation is assumed by the four-to-five combination. WLN has no capability to indicate this type of adjacency.

In May 1980 the LCS and WLN databases had 2,255,073 and 1,676,831 titles, respectively.[2] Neither had the capability at the time of this research for limiting searches by type of material desired. Both display multiple matches in screens of ten titles. The total number of matches are arranged in random order (not alphabetical) so screens must be searched consecutively, ten matches at a time.

One hundred and fifty-two periodical titles were randomly chosen from a pool of 445 serial titles known to be sought by patrons at the University of Illinois–Urbana Library.[3] The same titles were searched in both databases; periodicals were chosen because they had a higher probability of having a one- or two-word title and would therefore be a more rigid test of the search algorithms. Some of the periodicals were scholarly in nature (*Semiotica*, *American Journal of Obstetrics and Gynecology*), while others were more popular in nature (*Time*, *Seventeen*, *Look*). Table 1 shows the composition of the sample. The number of searchable words in the title varied from one to six, with WLN having sixty-one titles (41 percent) searchable using more than three terms. LCS, because of its limited search algorithm, could only be searched on the first two terms. Each title was searched by these researchers in LCS and WLN in an effort to determine which method was more effective for searching serial titles. It should be emphasized that the study was to measure effectiveness in locating the bibliographic record, and that eventual location of individual holdings was not studied.

METHODOLOGY

The 152 unique titles were randomly picked from a larger pool of 445 titles. Only those titles appearing to be periodicals were chosen to narrow the scope of the study. Annual reports, proceedings, and other nonjournal publications were excluded. Each title was searched in LCS and WLN and the following was noted on a coded sheet: number of keywords in the title, whether the title was found on the first screen, total number of matches found, number of screens searched to find the title, second search strategies, number found on first screen of second search, number of matches in second search, number of screens in second search, total search strategies, and a notation for those titles not found (see figure 1).

Table 1. Searchable Terms in Title

Words	LCS	%	WLN	%
One	32	20	23	15.1
Two	120	80	68	44.7
Three	—	—	42	27.6
Four	—	—	14	9.2
Five	—	—	4	2.6
Six	—	—	1	.7
Total	152	100	152	100.

Both databases were searched until the title was found except for those searches that produced more than 800 matches. It was felt that most patrons would not take the time to search through potentially eighty or more screens. No difference was made between a hard or microfilm copy of any title.

The hypothesis to be tested was that more success would be exhibited (less searches and/or screens will be needed) when keyword searching was done than when searching is done via a fixed-length algorithmic search. It was anticipated that by increasing a search from parts of the first two words of the title (LCS) to searching all significant words in that title (WLN), one would find a title more quickly.

RESULTS

One reason for long searches was a large number of matches or hits. Table 2 groups the number of matches in both LCS and WLN. At least seventy-three titles in LCS and eighty-two titles in WLN were found on the first search because they had fewer than ten matches (each database displayed ten titles per screen). WLN had nine titles ranging from 893 (*Forum*) to 17,348 (*America*) matches, while the highest number of matches in LCS was for *Latin American Perspectives* with 579 matches. As expected, the number of matches was directly influenced by the number of searchable terms, but it was also found that the number of matches was related to the uniqueness of the term or terms being searched.

Only three one-word titles produced ten or fewer matches in LCS (*Cognition*, *Mankind*, and *Tax Journal*). There were no one-word titles in WLN that had ten or fewer matches (see table 3). Eight of the highest number of matches from a one-keyword search in WLN (*Time*, *Aging*, *Geology*, *Forum*, *Challenge*, *Nature*, *Perception*) produced a range from a high of 275 matches (*America*) to a low of 19 matches (*Aging*) in LCS. The ability of LCS to dis-

TITLE: Amer j of pol. sci	LCS		WLN	
	104	1-3		
	5	4		
# OF KEYWORDS	02	5-6		24-25
FOUND ON FIRST SEARCH	2	7	1	26
# OF MATCHES FOUND	0143	8-11	0001	27-30
# OF SCREENS USED	03	12-13	01	31-32
2ND SEARCH STRATEGY		14		33
FOUND		15		34
# OF MATCHES		16-19		35-38
# OF SCREENS USED		20-21		39-46
TOTAL STRATEGIES		22		47
NEVER FOUND		23		48

The first count field (1–3) gives the worksheet number, the second (4) indicates the number of words in the title, and the third (5–6, 24–25) indicates the number of searchable words, excluding stopwords. The remaining count fields correspond to the labels on the left.

Fig. 1. Example of Worksheet.

tinguish a one-word search did not appear to reduce the number of matches produced by LCS in a one-term search (27 of 30 one-term searches produced 10 or more matches). The fact that WLN searched whole titles and subtitles did, however, appear to contribute to an increased number of matches for one-item searches (8 of 25 produced 601 or more matches).

Searches conducted with two significant search terms produced more titles with 10 or fewer matches (70 of 120, or 59 percent) for LCS than for WLN (32 of 68, or 47 percent). Titles such as *Scientific American*, *Library Journal*, *Social Inquiry*, *Modern Theatre*, and *American Home* produced fewer matches (1 to 36) in LCS than in WLN (65 to 169). The primary reason for this appeared to be the ability of LCS to do adjacency searches. A search on LCS such as "TLS/AMERHOME" asked the system for a word beginning with "AMER" to be adjacent to a word beginning "HOME" at the start of a title entry. The same search in WLN, "f t *American Home*," asks for those two specific words to appear anywhere in the title or subtitle and would therefore have the potential to produce many more matches.

As more search terms were available, WLN improved its efficiency (i.e., less matches were found). Of the fifty-seven titles having three or more searchable terms, 86 percent produced ten or fewer matches. Unfortunately, not many periodicals have titles with three or more words.

The other phenomenon related to the

Table 2. Number of Matches Found per Search

Matches	LCS	%	WLN	%
1–10	73	49	82	55
11–100	50	33	41	28
101–200	10	7	7	5
201–300	9	6	5	3
301–400	4	3	3	2
401–500	3	2	0	0
501–600	1	1	0	0
601–700	0	0	1	1
701–800	0	0	0	0
More than 801	0	0	9	6
Total	150	101*	148	100

*Figures add up to more than 100, due to rounding.

Table 3. Searchable Term by Matches

Matches	1–10	11–100	101–200	201–600	601–
One term					
LCS	3	9	5	12	0
WLN	0	8	4	3	8
Two terms					
LCS	70	41	5	4	0
WLN	32	27	3	4	2
Three terms					
WLN	34	6	0	1	0
Four terms					
WLN	13	0	0	0	0
Five terms					
WLN	1	1	0	0	0
Six terms					
WLN	1	0	0	0	0

number of matches was the uniqueness of the terms being searched. The partial correlation of whether a title was found on the first search in WLN when compared to the number of keywords, while controlling for matches and screens used, is −.3717 ($p = .001$). This indicates that only 18 percent of the variance in whether an item was found on the first search in WLN could be explained by the number of keywords in the title. The corresponding partial correlation for LCS is only −.1555 ($p = .03$) and explains only 2 percent of the variation. In both databases, general titles such as *Demography, Accountancy, Political Studies, Economist*, and *National Review* produced more matches than more specific titles such as *Industrial Gas, Farm Journal, Electronic News, Dun's Review*, and *Mother Jones*. Also, in LCS the four-to-five-letter sequence for searching tended in some cases to allow for too many permutations of a word. For example, *Progress in Mathematics*, which was searched by "TLS/PROG-MATHE," produced titles with *progress, programmation, progressive* and *mathematics, mathematical*, and *mathematiques* in them. The fact that "journal" and "of" are stopwords in LCS meant that searches were being conducted on third and fourth words in a title. Specific titles such as *Journal of Education, Journal of Politics*, and *Journal of Marriage and Family* were searched under the more general terms of *education, politics, marriage*, and *family*. This helped contribute to an increased number of matches.

Since both databases display titles in a random order, the number of matches does not preclude the user from finding an entry on the first screen even though there were more than ten matches. Table 4 shows how many titles were verified on the first screen. Ninety titles in LCS and eighty-seven titles in WLN were found on the first screen ($x^2 = .04$, phi = .22). Titles with matches as high as seventy-five (*Journal of Marriage and Family*) in LCS and with fifty-two (*Economic Journal*) and forty-four (*Semiotica*) matches in WLN were found on the first screen displayed. The conclusion is that the difference in finding a title on a first search in LCS when compared to WLN as large or larger than that in table 4 could arise by chance in about four out of one hundred trials. Therefore the difference is probably significant and not due to sampling error.

It should also be noted that fifty-five searches in LCS having 11 to 200 matches and seven searches with greater than 201 matches were successful in fewer than ten screens. In WLN, fifty-one searches having eleven or more matches were conducted in fewer than ten screens and only eleven searches with eleven or more matches took more than ten screens. Whether a patron

Table 4. Found on First Screen

	WLN	LCS
Yes	87	90
No	56	60

$x^2 = .04$ phi = .22

will search up to or even more than ten screens is not at present known.

CONCLUSION

The hypothesis that fewer screens would have to be searched using a keyword search than in a corresponding fixed-length algorithmic search could not be supported. The number of titles found using fewer than ten screens was higher in LCS than WLN. In fact, WLN produced more titles that had greater than 600 matches (10 versus 0 for LCS). The fact that a note appeared on the screen that the user's query produced 1,409, 5,799, or 17,348 hits might indeed influence whether that patron will continue with his search.

Although LCS was a fixed-length algorithmic search, there were some enhancements over the keyword searching in WLN. First, adjacency appeared to help in searches that had nongeneral terms. Second, the stopwords list acted as both a help and a hindrance. It helped when the words being searched were nongeneral (e.g., industrial or experimental), but caused many more matches when the only searchable word became general (e.g., education or politics). Also, the ability to distinguish one-word titles helped LCS outperform WLN when searching for a title having only one word because WLN produced nine one-word titles with more than 800 matches.

The results of this research suggest that libraries considering a keyword online catalog such as WLN or a fixed-length known-item public access system like LCS, need to make certain improvements. Either searching algorithm could be improved by having the ability to limit a search by the type of material (serial or monograph) requested. Two things would be helpful in keyword searching: the ability to specify that a title consists of only one word and the ability to show that words are adjacent. Following are two examples:

Had the search of *America* indicated that the title was a serial and consisted of only one word, many of the 17,348 matches would have been excluded. *American Home* as an adjacent term search would not have retrieved 169 matches in WLN. As the number of searchable words increased, the number of matches decreased. Therefore, LCS would benefit by having the ability to add a third adjacent term to its searching algorithm. If *Latin American Perspectives* was searched on three terms with a serial limiter in LCS, much fewer than 579 matches would have been retrieved. In addition to the above limiters, both systems would benefit from having an alphabetic sort of matches. Currently, both display the total matches retrieved in random order. Having the ability to put matches in one alphabetical order would decrease the number of screens a patron would need to search. A patron would also be able to see similar titles in an alphabetical system. Having all or some of these improved capabilities would greatly enhance the searching power of either database. Since programming trade-offs are a reality, more study into which of these capabilities is the most important when searching for serials is needed.

There are also several additional questions that this study did not address. For instance, do larger catalog databases cause more searching problems because they increase the number of matches? Will a patron search through five, six, or more screens? How many and what stopwords are necessary to ensure the most successful search with the least user effort? These questions need to be addressed in future studies of online catalogs.

As is the case in many studies of this type, more questions were raised than answered. With the current trend toward closing card catalogs and the move toward online catalogs, librarians need to know what are the best search algorithms to build into a system. More studies are necessary to make the future catalogs responsive to differing levels of bibliographic searching.

REFERENCES

1. F. Wilfrid Lancaster, *Information Retrieval Systems: Characteristics, Testing and Evaluation* (2d ed.; New York: Wiley, 1979), p.108–9.

2. "Data Base Statistics—May 6, 1980," *The WLN Participant*, 3:5 (June 1980).

3. Gary A. Golden, Susan U. Golden, and Rebecca T. Lenzini, "Patron Access to Serials: A User Study," *College & Research Libraries* 43:23–24 (Jan. 1982).

APPENDIX A. LIST OF ENGLISH LANGUAGE STOPWORDS

WLN

a	an	at	if	in
of	on	or	and	but
for	the	from	with	to

LCS

a	from	of	to
an	guide	on	u.s.
and	how	proceedings	u. s.
annual	in	report	united states
bulletin	international	reports	united states.
conference	introduction	study	with
department	journal	studies	yearbook
dept.	national	symposium	
for	new	the	

Communications

UTLAS-Japan Communications Link

George Gorsline, Jr., and Wyley L. Powell

On February 11, 1982, the first high-speed data communications link was established between Canada and Japan using a 9600-baud dedicated submarine cable circuit. This allowed the Maruzen Company Limited and International Christian University, both in Tokyo, to connect online to UTLAS in Toronto. The 9600-baud circuit uses a conventional voice grade circuit and allows UTLAS to multiplex up to thirty-two simultaneous channels at 300 baud each. Maruzen had had previous experience with these arrangements, having used an identical system to connect to Dialog and to interface to MARUNET, their domestic network which gives Maruzen customers throughout Japan (e.g., Osaka, Fukuoka) access to their MASIS Center in Tokyo.

Maruzen, founded in 1869, is one of Japan's largest companies in the book industry. In addition to its functions as a bookdealer, publisher, importer of foreign publications, and exporter of Japanese materials, it distributes information retrieval services and supplies stationery and office equipment. Maruzen distributes UTLAS' services and products in Japan.

International Christian University joined the UTLAS user network, through Maruzen, in April 1982. It was founded in 1949 by a group of Japanese and North American educators and currently has a student body of about 2,100. Among the unique titles which ICU contributes to the UTLAS database are older materials on religious studies, materials on Japanese culture and language, and English-language books published in Japan and other Asian countries.

UTLAS is currently investigating International Communication Access Service (packet service) to provide international service on a pay-as-you-go basis and to satisfy the Japanese and other countries' regulations regarding electronic mail ("message switching").* Packet switching was not originally considered because the Japanese overseas telecommunications company, Kokusai Denshin Denwa Company Limited (KDD), did not then support ICAS protocol, nor had UTLAS yet announced its packet network support. A satellite link was rejected because signals would have required over one-half second to travel up to the satellite and return to Earth. These delays are unacceptably long for dedicated data service using the full-duplex mode of operation because of their impact on terminal response time.

To establish the Toronto-Tokyo link, it was necessary to install multiplexing equipment and order the 9,000-mile-long circuit. Since seven different companies were involved in providing segments of the circuit,† the integration and testing end-to-end became quite challenging. Pieces do not always fit together as expected, and, with data moving in both directions on a

George Gorsline, Jr., is director of technical support and Wyley L. Powell is in Public Relations, UTLAS, Toronto, Ontario.

*Outside North America, communications facilities (telephone, telex, data, etc.) are operated by government agencies and regulated on a content basis. As computer message switching could displace telex revenues, many countries require use of ICAS, payment per message, instead of flat rate for unlimited use lines.

†UTLAS, Bell Canada, CNCP Telecommunications, Teleglobe Canada, KDD, Nippon Telephone and Telegraph (NTT), and Maruzen were the companies involved. The circuit runs from Toronto to Vancouver to Hawaii to Guam to Okinawa to Tokyo and Osaka.

full-duplex circuit, it is possible for the connection to work well one way but not the other. The fourteen-hour time difference and different technical standards (again North American vis-à-vis the rest of the world) further complicated the testing. However, problems are to be expected in setting up any complex circuit, and the dedication and goodwill of many people working during their normal sleeping hours made the circuit work.

More recently, UTLAS installed a TYMNET host node, thus providing local-call dial access in virtually all U.S. centers and giving UTLAS an immediate presence throughout North America. ■■

The WLN/RLG/LC Linked Systems Project

Wayne E. Davison

INTRODUCTION

This communication draws heavily on documents which are the work of many of the Linked Systems Project team members. A list of documents produced by the project is to be found in appendix A.

The WLN/RLG/LC Linked Systems Project (LSP) began early in 1980. With funding from the Council on Library Resources (CLR), the Washington Library Network (WLN) and the Research Libraries Group, Inc. (RLG) began work, along with the Library of Congress (LC), on an online communications link and intersystem data retrieval and maintenance facility to support a shared authority file. The communications link will enable heterogeneous computer systems to exchange data, and is extensible to additional systems beyond the initial three. This link will be based on existing and emerging international standards to the fullest extent possible. The shared authority file will support the Name Authority File Service (NAFS), and help to overcome a barrier to effective exchange of bibliographic data caused by the frequent lack of heading consistency across systems. The NAFS and the retrieval facilities will serve as the basis for future sharing of bibliographic data. Thus, the LSP will provide the foundations for a national bibliographic network.

While the computer-to-computer communications link will first be used for the exchange of authority data, the existence of the link will allow later sharing of full catalog records, location and holdings data, the transmittal of interlibrary loan requests, communication with book vendors, et al. At the completion of the Linked Systems Project, there will be a network in place upon which many applications can be implemented, using the cooperative creation and maintenance of consistent databases, search and retrieval between systems, and exchange of free-form messages. Truly a national network will have come into being, based upon a diversity of systems linked by a standard interface.

ORGANIZATION OF THE LINKED SYSTEMS PROJECT

WLN, RLG, and LC participate in the LSP as peers, and all decisions are made by consensus. C. Lee Jones is the program officer who oversees the project for CLR. The senior officers responsible for the project at each of the member institutions currently are Raymond DeBuse for WLN, Tina Kass for RLG, and Henriette Avram for LC. There are two working groups, each composed of representatives from each of the three institutions. The Telecommunications Group is responsible for all aspects of the communications link between the applications programs in each system. The Authorities Group is responsible for the applications support of authority record retrieval and maintenance between the systems and between each system and its own terminal users. The LSP is divided into three segments, named Project 1, Standard Network Interconnection (SNI), and Authorities Implementation (AI) (see figure 1). Project 1, the requirements definition segment, contained two parallel components: the authorities component and the

Wayne E. Davison is manager, Communications Development, for The Research Libraries Group, Inc., Stanford, California.

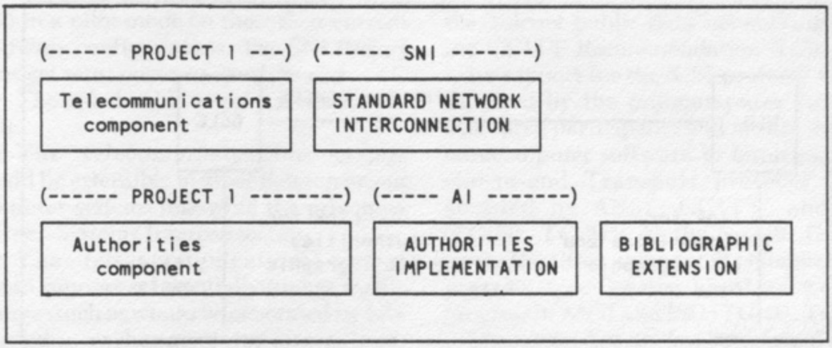

Fig. 1.

telecommunications component. A fourth segment is currently being defined to extend the functions of the LSP to include exchange of bibliographic records as well as authority records.

THE NEED TO LINK SYSTEMS

In order to meet the service needs of their users and to deal with economic realities, libraries must share information resources to the fullest extent possible. While each of the bibliographic utilities has facilitated this sharing among its own membership, each utility has been isolated from the others. A major reason for this isolation is the lack of common data communications procedures for interconnecting dissimilar computer systems. As the cost of computer hardware continues to fall, more individual libraries are installing their own local systems to supply their users with improved services in the areas of circulation control, online catalogs, etc. Unless there is a mechanism to link these systems, the problem of isolation will be compounded.

The linking of computerized systems to permit one system to exchange data with another system is a basic and essential task to any effective cooperative intersystem resource sharing. By developing standardized techniques for linking systems and implementing these on three major library automation systems, a significant contribution toward meeting the goals of the Council on Library Resources' Bibliographic Service Development Program (BSDP) will be made. For a brief description of BSDP, see *CLR Recent Developments*, V.9, no.1 (April 1981).

The rapid growth and success of the online shared cataloging systems would not have occurred without the development of a standard format for the communication of machine-readable cataloging data (MARC). The adoption of the MARC formats as standards allowed the development of different systems, all of which could process the same data as it was received from the Library of Congress or from other sources adhering to those standards. MARC will continue to be the standard as online data exchange between systems is developed. Similarly, the creation of standards for online links between bibliographic systems will facilitate both the expansion of resource sharing among libraries and the development of entirely new services, such as the Name Authority File Service (NAFS).

The potential for use of the standard linking mechanisms between other systems and for other applications is significant (see figure 2). Cooperative projects such as CONSER could use the developed techniques and protocols for intersystem file maintenance between LC and OCLC and for record contribution from other utilities. Members of utilities that have their own computer system, such as Northwestern University, could use the standard linking mechanisms in their cataloging and authority work.

THE LSP NETWORK LINK

During Project 1, the Telecommunica-

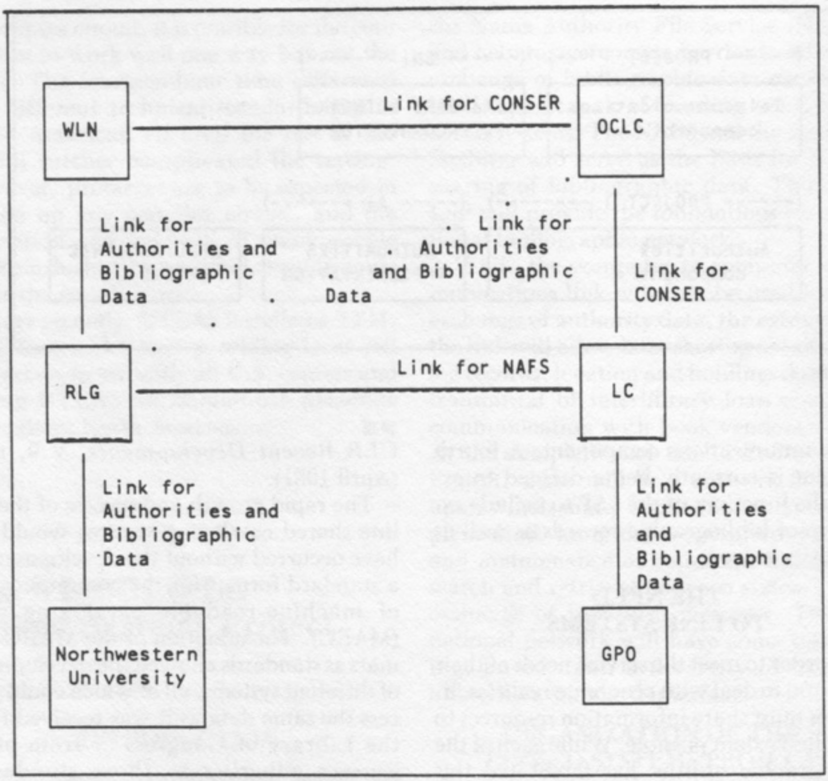

Fig. 2. Potential Configuration and Applications.

tions Group surveyed the current state of computer networking, including offerings from computer hardware and software vendors, data communications vendors and the standards work being done in the International Standards Organization (ISO), the American National Standards Institute (ANSI), the Consultative Committee for International Telegraphy and Telephony (CCITT), and the U.S. Department of Commerce, National Bureau of Standards (NBS). The possible network solutions were measured against the following criteria:

- The reference model of Open Systems Interconnection [TG.4]* should be followed in the design of the telecommunications link.
- The link must support multiple simultaneous application-to-application interactions among host computers (i.e., it is not a terminal emulation protocol).
- The software involved in the link must have available a sufficient level of support to ensure its operational reliability. The link should not require excessive operational personnel. Network management must not be dependent on any single node.
- The software must belong to a clear evolutionary software architecture with an expected lifetime of at least five years. In the dynamic world of networking, we must not base the link on software that will be shortly obsolete.
- The link must be implemented by LC, WLN, and RLG by no later than October

*The TG accession numbers which appear in brackets throughout the text refer to items listed in the bibliography.

1983 in a pilot mode on their then-current hardware configurations. The SNI Project 2 budget must cover its development.

• The link should be cost-effective to operate.

• The telecommunications system should be extensible to other heterogeneous computer systems following the principles of Open Systems Interconnection.

• The telecommunications system should support substantially higher traffic volume (such as would be generated by bibliographic exchange, interlibrary loan, etc.) without major redesign.

The preferred solution was one based on the ISO Open System Interconnection—Basic Reference Model (see figure 3 and the section "Background: Computer-to-Computer Links" below). Each institution will use a minicomputer to interface with the Telenet public data networking using the CCITT Recommendation X.25. Software support for the X.25 protocol will be supplied by the minicomputer vendors. The three participants will jointly develop minicomputer software to implement the end-to-end Transport protocol being adopted by ANSI, CCITT, and ISO [TG.56, TG.57]. At the Session Control layer, the three parties will collaborate on program specification based on work in progress in ANSI and ISO [TG.60, TG.61].

No standards work has been found at the Presentation and Application Layers that is applicable to bibliographic systems. Therefore, CLR contracted with James Aagaard of the Northwestern University Libraries to head a project to draft specifications for these areas. This work began in July 1981 and was completed in October 1982

```
                         WLN

                         |7|
                         |-|
                         |6|
                         |-|
                         |5|
                         |-|
                         |4|
                         |-|
                         |3| .....
                         |-|   ↑
                         |2|   :
                         |-|   :
                         |1|   :
                         |↑|   :
                          :   X.25
                         |v|   :
                          1    :
                          -    :
                          2    :
                          -    :
                          3    v
                         |T| .....
                          e
R ____________  _______   l   _______  ____________ L
L 7|6|5|4|3|2|1:<>:1|2|3 Telenet 3|2|1:<>:1|2|3|4|5|6|7 C
G ____________  _______   n   _______  ____________
          :          :    e    :           :
          :<---X.25--->:  t    :<---X.25--->:
```

Fig. 3.

[TG.1.5, TG.50, TG.63, TG.64]. CLR has provided additional support for this team to coordinate protocol development with the Bibliographic Interest Group (BIG) in Canada. BIG is currently working on an experiment to transfer batch MARC search results from the National Library of Canada (NLC) to five Canadian universities using the iNET [TG.29a, TG.29b]. In addition to being used by the LSP, Aagaard's work will serve as the basis for a standard to be proposed by ANSI Committee Z39 Subcommittee D for national and international adoption.

THE STANDARD NETWORK INTERCONNECTION (SNI)

The Project for a Standard Network Implementation (SNI) covers the design, development, and implementation of a standardized telecommunications link between the systems at WLN, RLG, and LC to allow library applications on one system to exchange data with library applications on another system.

The SNI portion of the LSP began in January 1982 and will extend into late 1983. During the course of the SNI project, WLN, RLG, and LC will acquire and install the minicomputer hardware, and the vendor-supplied X.25 software to connect to Telenet, and test that connection. Using a Digital Equipment Corporation (DEC) PDP11/24 minicomputer, RLG will take primary responsibility for developing software in the PASCAL programming language to implement the Transport Layer. WLN and LC will adapt this software to operate on their Data General Eclipse minicomputers. The three participants will collaborate on program specifications for Session Control including interfaces to the application and presentation protocol processes that will be developed later. LC will take primary responsibility for development of Session software in PL/1 for the CICS environment which it shares in common with WLN. At the completion of the project, the linking mechanism will be in place to support a wide variety of library resource sharing applications.

THE NEED FOR SHARED AUTHORITY FILES

Shared authority files take on a double importance in a linked system environment. First, it is important to share authority information because authority work is one of the most costly and labor-intensive aspects of cataloging. Second, bibliographic data cannot be shared effectively among linked systems unless there is some mechanism to overcome the lack of heading consistency across those systems.

THE AUTHORITY APPLICATION IMPLEMENTATION (AI)

The AI, which began in late 1982, covers the detailed design, implementation, and evaluation of the capability for online authority record exchange between the WLN, RLG, and LC. The three participants will develop applications to support intersystem search and response and intersystem file maintenance facilities. There are two general types of sessions, interactive and batch; both will use the computer-to-computer link being developed by the SNI. Only two systems will be involved in a given session. The interaction is a computer-to-computer interaction. The

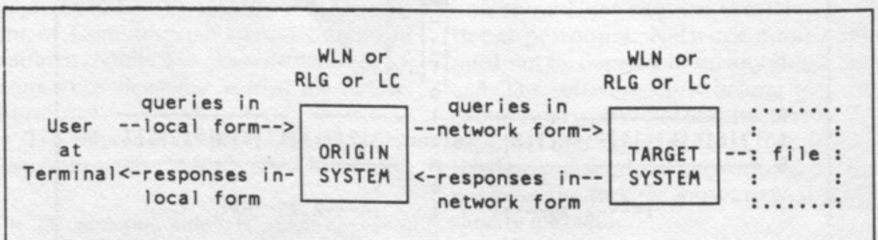

Fig. 4. Authority Searching.

terminal user communicates with the local system, which communicates with the target system on behalf of the user. The target system communicates its results to the originating system, which, in turn, communicates the results back to the user. Search/response is a general interactive facility which will permit searching of authority files on the various systems. The batch file maintenance facility, while general in concept, will be implemented specifically to support the building of the Name Authority File Service. These two facilities will support the cooperative creation and maintenance of the Name Authority File Service.

The master name authority file will be mounted on the system at LC. The systems at WLN and RLG will have their own authority files to support their users. Any one of the systems will be able to search any of the authority files on any of the three systems. Basically, a searcher on one system formulates a search using the search language of that system (see figure 4). The search is translated into an intersystem syntax by the originating system and sent to the target system where it is then translated by the target system into the query syntax of that system. After the request is processed, either the number of retrieved records or the actual records are sent to the originating system. The records are transmitted in the US MARC format, but the originating system displays responses to the terminal user in the screen format of the user's system.

The file maintenance facility will consist of three components: record contribution, reporting, and record distribution. For record contribution, when a user on an LSP system completes a new or updated record, the system at LC can request a copy for contribution to the master file (see figure 5).

If the record fails the edits, the master file system produces a notification for the system from which the record came, indicating the reason for failure. The systems at WLN and RLG can request these reports about the addition of these records to the master file at LC (see figure 6). The master file system does not maintain a copy of the failed record.

If the record passes all edits of the master file system, it is added to the master file, the originating system is notified of the approval, and the record is marked for distribution to all subscribers to that file, either online or on tape. The systems at WLN and RLG can request copies of records that have been added to the master file (see figure 7).

Record contribution and record distribu-

```
              LC                                WLN or RLG
                          requests in                          ...............
.........   +--------+  --network form->  +--------+           :   records   :
: master :<-| ORIGIN |                    | TARGET |<-:          to be     :
: file   :  | SYSTEM |  <-records in----  | SYSTEM |           : contributed :
:........:  +--------+    network form    +--------+           :.............:
```

Fig. 5. Authority Record Contribution.

```
WLN or RLG                           LC
               requests in                     ..............
+--------+   --network form->    +--------+    :            :
| ORIGIN |                       | TARGET | <-: record edit :
| SYSTEM |   <-reports in----    | SYSTEM |    :   reports   :
+--------+     network form      +--------+    :............:
```

Fig. 6. Reports on Record Contribution.

```
WLN or RLG                                  LC
               requests in                  ..............
  ORIGIN     --network form->   TARGET      :   records   :
  SYSTEM                        SYSTEM      :    to be    :
             <-records in----               : distributed :
               network form                 :.............:
```

Fig. 7. Record Distribution.

tion use the same record transmission mechanism: a sequential file or queue is generated on the system generating the transactions and the receiving system asks for records from this queue. The records are then transferred one at a time. The receiving system keeps track of where it is in the queue so that if restart is needed, it knows where to ask for restart. Each record is checked on receipt for integrity at a very low level. If the record is a contribution to the master file, it is also passed through the master file edits.

Record contribution and record distribution differ in that responses, e.g., edit reports, are returned for record transmission that takes place because of record contribution, whereas responses are not returned for distribution records. The response transmission is handled like record transmission in that the system generating the responses puts them into a queue and the receiving system asks for items from that queue. Thus record contribution and record distribution use the same general mechanism for transmission of both types of items, records and responses.

BACKGROUND: COMPUTER-TO-COMPUTER LINKS

Research into linking computer systems began in earnest in the early 1970s when the U.S. Department of Defense Advance Projects Research Administration developed the ARPANET. This work was based on packet-switching technology and connected a wide variety of computer systems. Although development of the ARPANET ceased for all practical purposes in the mid 1970s, this effort remains a fountainhead in the area of computer networking.

In the mid-1970s the European community took a major step in the area of computer networking when the Consultative Committee for International Telegraphy and Telephony (CCITT) adopted its Recommendation X.25 specifying a common interface between computer equipment and public packet-switching data networks.

At the same time, U.S. computer manufacturers were introducing their own proprietary computer networking solutions. These systems facilitated interconnection of each manufacturer's computers with others from the same manufacturer, but created barriers against connection with computers from another manufacturer.

During the mid-1970s, the library community also became involved in computer networking. Early efforts in this area were made by the Telecommunications Committee of ALA in 1975 [TG.68]. This work was then taken over by the NCLIS/NBS Task Force in 1976 [TG.19]. The task force produced a set of procedures for program-to-program communication, but stopped short of addressing the issues of bibliographic record retrieval and maintenance. This work was used by a group of four Norwegian universities in a research project during 1981 to link their automated bibliographic systems [TG.53, TG.54].

In 1977 the Research Libraries Group, Inc., established a computer-to-computer link between its host system at the New York Public Library and LC to permit RLG catalogers to search the LC database through the NYPL system and bring records back from LC to be incorporated into the RLG database [TG.70]. Because no computer-to-computer communications protocols were supported by either system, a terminal emulation mode was selected for easy development for the pilot. In this experiment, the NYPL computer appeared to the LC computer as four terminals rather than as another computer. Although this technique did permit the exchange of data, it was realized that the link was both unreliable and made inefficient use of the communications resources. This experiment illustrated the problems inherent in using a master-slave terminal emulation approach in place of a peer level program-to-program approach to computer networking. It was also felt that a more standard protocol designed especially for computer-to-computer communications would be more extensible than the IBM-specific terminal emulation technique used in the pilot link.

THE OPEN SYSTEMS INTERCONNECTION— BASIC REFERENCE MODEL

The national and international standards organizations have a major activity under way in the area of computer communications through the development of the International Standards Organization (ISO) Reference Model for Open Systems Interconnection (OSI). The result of these efforts was a basic reference model for open systems interconnection. The model has now attained the status of an international draft standard and has served as the touchstone for discussions of network architecture during the last two years. See draft proposal ISO/DP 7498 for a full description [TG.4].

The purpose of the OSI (Open Systems Interconnection) Basic Reference Model is to provide a common basis for the coordination of standards development for systems interconnection, while allowing existing standards to be placed into perspective. Open Systems Interconnection refers to standards for the exchange of information among processes, computers, people, networks, etc., that are "open" to one another for this purpose by virtue of their mutual use of the applicable standards.

This model describes a layered approach to defining and implementing communications protocols. In this way the functions and services of each layer can be specified and implemented independently from the other layers. Layering is a technique that imposes an order and structure on the functions required to support computer-to-computer communication. All of these functions must be performed; layering clarifies the problem and simplifies the interactions between these functions. The ARPA Network protocols introduced layering into data communications. Since that time layering has become the accepted software engineering technique for organizing network architectures. It has been used as the basis for virtually all networks developed in recent years. IBM's System Network Architecture (SNA) and Digital Equipment Corporation's DECnet are two notable examples.

Some of the benefits of layering are:

- the ability to "divide and conquer" a complex technical problem
- reduction of complexity by imposing a discipline on the interactions between system components by defining and limiting those interactions
- definition of a standard nomenclature for discussion and design of network systems
- definition of a standard architecture that facilitates and encourages the development of standard procedures for all levels
- the ability to replace any single building block with a better one without disturbing any of the others
- simplified error detection and program debugging

In the OSI model, each system is logically composed of a hierarchically ordered set of seven processes (see figure 8). Adjacent processes within a single system communicate through an interface at their common boundary, the vertical dimension. Processes of the same rank in separate systems collectively form a layer, the horizontal dimension. Except for the highest layer, each layer provides processes in the next higher layer with services. The cooperation between processes within a layer is governed by a set of protocols specific to that layer.

In order to cooperate, processes in any layer, other than the lowest, communicate by means of the set of services provided by the next lower layer. In order for informa-

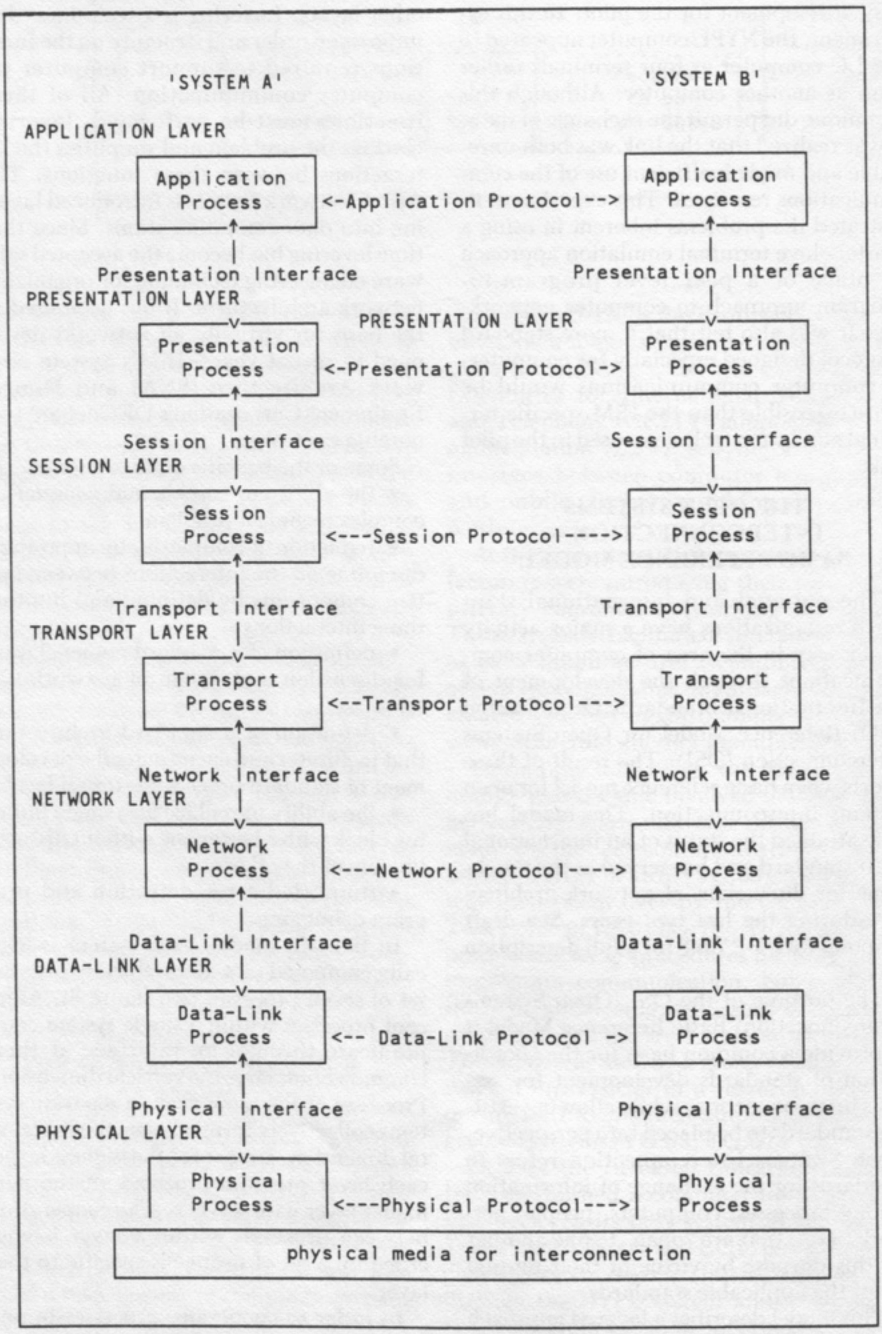

Fig. 8. Open Systems Reference Model.

tion to be exchanged between two peer processes, a connection must be established in the next lower layer using a protocol of that next lower layer. This pattern applies downwards until an available connection, the physical medium or a common process, is encountered. Protocol control information is exchanged between two peer processes to coordinate their joint operation. Interface control information is exchanged between a higher level process and its next lower level process to coordinate their joint operation.

The Application Layer is the highest layer in the architecture. This layer exists to carry out the application specific aspects of interprocess communication between application processes. The Application Layer is where the actual information processing is done. This is the highest layer in the OSI architecture. Protocols of this layer directly serve the end-user application process by providing the distributed information service appropriate to an application and to system management. All functions directly used by application processes, including those functions which are handled by a request for a service from a lower layer, are part of this layer. The other layers exist only to support this layer.

Presentation processes exist to provide services to application processes. The primary service provided by a presentation process is to perform transformations of data and operations between local forms and network canonical forms. The other service is to provide a pass-through connection to a session process. The Presentation Layer provides general functions for the representation and manipulation of structured data. The purpose of the Presentation Layer is to provide the set of services for the Application Layer to enable it to interpret the meaning of the data exchanged between application processes. The Presentation Layer is concerned with the management of formats and the management and performance of translations for the users in the Application Layer. The Presentation Layer is on top of the Session Layer, which serves to link a pair of processes in the Presentation Layer.

Session processes exist to provide services to presentation processes. The primary services provided by a session process are:

- session establishment
- context management
- session identification
- session recovery
- session release
- data exchange
- data delimiting
- dialogue management

The Session Layer establishes, controls, and clears the cooperative relationships between different application processes. These relationships are called sessions. The Session Layer serves to connect distributed activities into a logical relationship for the control of data exchange, and to disconnect them. The Session Layer specifically supports interactions between cooperating presentation processes.

The Transport Layer performs all functions necessary to bridge the gap between the services provided by the Network Layer and the services needed by the Session Layer. Therefore the functions to be performed by the Transport Layer are dependent on the services provided by the underlying Network Layer and the services required by the Session Layer. Potential Transport Layer functions include:

- address mapping
- multiplexing end-to-end connections onto network connections
- end-to-end sequencing
- end-to-end error detection and recovery
- end-to-end flow control
- end-to-end notification of delivery
- end-to-end segmenting and blocking
- expedited data transfer

The Transport Layer exists to provide interprocess data transfer in a network environment. The Transport Layer itself controls data transportation from source host to destination host, thereby providing transparent transfer of data between session processes. It provides the end-to-end control necessary to guarantee error-free and properly sequenced messages. The complexity of this layer is inversely related to the reliability of the layers below it. This layer relieves the higher layers of all concern with the transportation of data between them.

Network processes exist to provide services to transport processes. The primary services provided by a network process are:

- routing
- segmenting and blocking
- error detection and recovery
- sequencing
- flow control
- notification of delivery

The Network Layer provides a connection path between a pair of transport processes. It provides transport processes with independence from routing and switching considerations. In the Open Systems architecture, some systems are final destinations for messages, while others may act as intermediate nodes that forward data to other systems. The required routing procedures are grouped in the Network Layer. The standard for the Network Layer is the CCITT Recommendation X.25: "Interface between DTE and DCE for Terminals Operating in the Packet Mode on Public Data Networks." In this context, the terms "DTE" (Data Terminal Equipment) and "Terminals" refer to all data processing equipment, including main-frame computers. "DCE" (Data Communication Equipment) refers to a packet-switched network, such as Telenet or Tymnet.

Data-Link processes exist to provide services to network processes. The purpose of the Data-Link Layer is to mask the characteristics of the Physical Layer from the network processes. The primary services provided by a Data-Link process are:

- address mapping
- delimiting and synchronization
- node-to-node sequence control
- node-to-node error detection and recovery
- node-to-node flow control

The Data-Link Layer provides the functional and procedural means to establish, maintain, and release one or more data links between network processes. Some physical connection media require rigorous checking procedures to overcome the inherently high error rate. Conventional telephone lines are a case in point. These procedures have been codified in data-link control procedures, which have been studied and standardized for a number of years. Past usage indicates that these procedures should be placed in their own layer. X.25 specifies a data-link protocol that is identical with one of the modes specified in the ISO HDLC (High-level Data Link Control) protocol standard. The American National Standard version of HDLC is ADCCP (Advanced Data Communication Control Procedures, X3.66-1979).

The primary function of the Physical Layer is to provide transmission of transparent bit streams between data-link processes. The Physical Layer provides the mechanical, electrical, functional, and procedural characteristics to establish, maintain, and release physical connections. It is essential that the architecture permit the use of a realistic variety of physical connection media; e.g., satellite, fiber optics, conventional telephone lines. Because different control procedures are needed for different media, these were isolated in this layer. There are several standards for the Physical Layer. By far the most widespread in the United States is the Electronics Industries Association (EIA) RS-232-C.

BIBLIOGRAPHY

The items listed in this bibliography are not produced by the Telecommunications Group: the TG accession numbers are given for convenience in citation only.

TG.1.5 *Application Level Protocol Development for Library and Information Science Applications—Volume 1, Service Definition.* Prepared by James S. Aagaard, Kenton E. Andersen, and Wayne E. Davison. October 29, 1982. 48p.

TG.4 *Data Processing—Open Systems Interconnection—Basic Reference Model.* Draft Proposal ISO/DP 7498. December 3, 1980. 100p.

TG.19 *A Computer Network Protocol for Library and Information Science Applications, by the NCLIS/NBS Task Force on Computer Network Protocol.* December 1977. 90p.

TG.29a *A file transfer service for the bibliographic interest group,* by R. J. A. Buhr, D. A. MacKinnon; G. V. Bochmann (consultant). Computer Gateways Incorporated. April 30, 1982. 185p.

TG.29b *A file transfer protocol for the bibliographic interest group,* by R. J. A. Buhr, D. A. MacKinnon; G. V. Bochmann (consultant). n.d. 213p.

TG.50 *Application Level Protocol Development for Library and Information Science Applications, Specification of the Protocol.* Prepared by: James S. Aagaard, Kenton E.

Andersen, and Wayne E. Davison. April 28, 1982. 143p.
TG.53 *The BIBNETT Project: system to system communication using a computer network.* By Liv A. Holm. 1982-06-01. 49p.
TG.54 *Application protocol for library and information systems: the network protocol from the BIBNETT project.* By Liv A. Holm. 1982-03-31. 41p.
TG.56 *Information Processing Systems—Open Systems Interconnections—Transport Protocol Specification.* ISO/TC97/SC16 N1169. June 1982. 85p.
TG.57 *Information Processing Systems—Open Systems Interconnection—Transport Service Definition.* ISO/TC97/SC16 N1162. June 1982. 39p.
TG.60 Draft, *Basic Connection-Oriented Session Service Definition.* ISO/TC97/SC16 N1166. June 1982. 36p.
TG.61 Current Status of Draft, *Basic Connection-Oriented Session Protocol Specification.* ISO/TC97/SC16 N1167. June 1982. Var. paging.
TG.63 *Presentation Level Protocol Development for Library and Information Science Applications—Volume 1, Service Definition.* Council on Library Resources Grant 2033. Prepared by James S. Aagaard, Kenton E. Andersen, and Wayne E. Davison. October 29, 1982. 22p.
TG.64 *Presentation Level Protocol Development for Library and Information Science Applications—Volume 2, Protocol Specification.* Council on Library Resources Grant 2033. Prepared by James S. Aagaard, Kenton E. Andersen, and Wayne E. Davison. October 29, 1982. 43p.
TG.68 *Protocol for Computer-to-Computer Communications.* Technical Communications—*Journal of Library Automation* 9:167–71 (June 1976).
TG.70 *Final Report on the RLG/LC Pilot Cataloging Project: "Establishment of Library Network On-Line Access to the Library of Congress MARC Data Base."* By Peter S. Graham. December 31, 1978. 68p. ■■

APPENDIX A. LSP DOCUMENTS LIST

The following are current project working documents, which are subject to ongoing revisions, and intended for limited distribution. Individuals involved in related developments are invited to address inquiries to the author.

Telecommunications Group

P1-t1 *Telecommunications Coordination Plan.* July 1980.
P1-t2 *Facilities Report.* August 1980.
P1-t3 *Protocol Evaluation, Survey of Network Offerings, Traffic Estimates, and Assessment of Investment in Linkage.* March 1981.
P1-t4 *System Process Model.* July 1981.
P1-t5 *Alternative Link Solutions: Description, Evaluation, and Selection.* December 1981.
SNI.0 *Proposal to the Council on Library Resources for Standard Network Interconnection.* December 1981.
SNI.1 *Development Coordination Plan.* February 1982.
SNI.2 *Network Topology Document.* March 1982.
SNI.3 *Network Layer Specification.* September 1982.
SNI.4 *Connection-Oriented Transport Layer Specification.* September 1982.
SNI.5 *Network Front End Acquisition, Hardware and Software Procurement Specification.* June 1982.
SNI.6 *Connection-Oriented Session Layer Specification.* November 1982.

Authorities Group

LASP-6 *MARC authority record format review* / LASP. — (July 9, 1980)
LASP-7 *General description of the project : authorities* / LASP. — (July 11, 1980)
LASP-8 *Organization and communication coordination plan* / LASP. — (July 8, 1980)
LASP-12 *MARC authority record format review : background material* / LASP. — (October 31, 1980)
LASP-13 *Characteristics of current authority data files* / LASP. — (October 31, 1980)
LASP-14 *Comments on the National Level Authority Record (NLAR)* / submitted by the Washington Library Network and the Research Libraries Group. — (March 1980)
LASP-19 *MARC authority record format review : $w control subfield* / LASP. — (December 4, 1980; revised February 26, 1981)
LASP-20 *Requirements for record content : NLAR and LASP* / LASP. — (February 27, 1981)
LASP-21 *Library of Congress automated authority control system : overview* / LC. — (February 1981)
LASP-22 *Search and response facilities: WLN, RLG, and LC* / LASP. — (February 1981)
LASP-23 *Authorities subsystem external design* / RLG. — (February 3, 1981)
LASP-25 *LASP Joint functional specifications: Intersystem component* / LASP. — (March 31, 1981)
LASP-25a *LASP Joint functional specifications: Intersystem component* / LASP. — (August 31, 1981)
LASP-27 *Organization and communication coordination plan for Phase II* / LASP. — (March 25, 1981)

LASP-29 *Overview of LASP Phase I: Authorities group* / LASP. — (March 31, 1981)
LASP-30 *WLN authority control system: overview* / WLN. — (March 31, 1981)
LASP-33 *Comparison of components and hierarchical relationships of authority headings in WLN, RLIN, and LC* / LASP. — (August 31, 1981)
LASP-34 *Uniqueness of headings: Comparison of WLN/RLIN/LC normalization rules* / LASP. — (August 31, 1981)
LASP-36 *LSP functional specifications: WLN system component* / WLN.— (November 1981)
LASP-37 *LSP functional specifications: LC system component* / LC. — (November 1981)
LASP-38 *LSP functional specifications: RLG system component* / RLG. — (November 1981)
LASP-39 *LSP external design: Intersystem component* / LSP. — (February 26, 1982)
LASP-39a *LSP external design: Intersystem component* / LSP. — (Revised August 31, 1982)
LASP-41 *Proposal to the Council on Library Resources for an authorities implementation and bibliographic analysis* / LSP. — (March 15, 1982)
LASP-41a *Proposal to the Council on Library Resources for an authorities implementation* / LSP. — (Revised September 30, 1982)
LASP-43 *Addendum to the WLN/RLG/LC Linked System Project proposal for Authorities Implementation and Bibliographic Analysis: Further Discussion of the Bibliographic Analysis* / LSP. — (April 2, 1982)
LASP-45 *RLIN Authority File configuration to support the Name Authority File (NAF)* / RLG. — (April 2, 1982)
LASP-47 *LSP external design: WLN system component* / WLN. — (August 31, 1982)
LASP-48 *LSP external design: RLG system component* / RLG.— (August 31, 1982)
LASP-49 *LSP external design: LC system component* / LC. — (August 31, 1982)
LASP-50 *WLN/RLG/LC Linked Systems Project final report* / LSP. — (August 31, 1982)

Microcomputers for the Public in the Public Library

Bonnie S. Fowler
and Duncan Smith

Libraries are moving with the times, seeking better ways to serve their intended patrons. In recent years one can see the movement in the public library by its increased attention to technology and networking. Forsyth County Public Library, located in Winston-Salem, North Carolina, has become one such library through an exploratory program in its Adult Continuing Education Department (ACE). The program, originally funded through a Library Services and Construction Act (LSCA)

Bonnie S. Fowler is a part-time librarian in the Adult Continuing Education Department and Duncan Smith is Kernersville Branch librarian at the Forsyth County Public Library, Winston-Salem, North Carolina.

grant, offers for public access four Apple II Plus microcomputers. Each microcomputer has a tutorial disk and booklet for the uninitiated microcomputer user, and the library also provides a wide range of other more sophisticated software packages and print materials for public use. Data gathered in a recent survey and daily statistics of the computers' usage indicate that the program is achieving its service objectives.

When the idea for such a program was first discussed at the library, it seemed appropriate that computers should be part of ACE's services. ACE itself was originally an experimentally funded LSCA project that was undertaken by county funding when its federal money was exhausted. ACE's goal was then, and still is, to aid adults in making informed decisions in the areas of continuing education, job hunting, and career development. To attain this goal, ACE is set up as an area separate from the adult reference department. In the area are housed all materials concerning colleges and schools; community classes/opportunities; financial aid; job-hunting skills; career choice, development, and change; grantsmanship; various career and school tests' preparation; adult basic education and General Educational Development (GED) test materials. In addition to books, the department has extensive career, vertical, and grantsmanship files, current course information, college catalogs, an automated Guidance Information Service, typewriters for public use, and the local Employment Security Commission's job-listing fiche. ACE's patrons are both regular public library users and referrals from other local agencies. More often than not, these referrals are new users, surprised at what the public library can offer them. ACE therefore views itself, and is viewed by the community, as a vital part of Winston-Salem's adult educational delivery system, and microcomputers were easily integrated into the department's services.

HISTORY

Much has been said and written about the usefulness and appropriateness of microcomputers in the public library; however, most recent action has been directed at how the library itself can use the computer or how patrons can benefit through better services provided by a library's utilization of microcomputers. It has been primarily in schools and as catalogs that patrons have actually interacted with computers in the library. Neither of these fully develops the microcomputer's potential for aiding the adult learner. Considering the growing importance and abilities of computers and the goal of the ACE department, a proposal was developed to explore the use of microcomputers in the delivery of both traditional and nontraditional public library reference/informational services.

After preliminary discussions with other community agencies working in adult continuing education, three primary goals were defined. Ultimately, the services and emphases of the project would rest upon these three first-year goals, as follows:

1. To develop computer literacy and to enhance the ability of adults to interact with computers.

2. To provide adults with opportunities to assess their career interests and/or develop new career-related skills.

3. To improve awareness of and access to educational opportunities for adults.

These goals and how they are being met will be discussed as they relate to their respective services—computer literacy, information/instruction, and development of a local educational opportunity database. It was felt that these goals would offer patrons a variety of opportunities to interact with microcomputers, and possibly even introduce library users and nonusers to a different facet of the public library.

The proposal requested an initial sum for the purchase of various microcomputer hardware, software, and print materials. It was decided that the library would fund the staff for the project. No new staff was hired; existing staff was trained in the use of microcomputers and existing software. Since Apple had the most educationally appropriate software and had been very supportive during initial inquiries, it was selected as the vendor. The project then purchased two Apple II Plus microcomputers with 48K RAM, two 9-inch color monitors, one Epson printer, and a variety of software and print materials.

SERVICES

To meet the stated goals, four areas of services were developed, three of which were immediately available. The first goal of acquainting the public with micros was accomplished through the use of the *Applesoft Tutorial* (a self-instructional manual), other more advanced guides, and various programming books and magazines. A list of books and magazines used in the program is listed in the bibliography. (Included in the list are a number of non-Apple materials, bought to serve patrons who are working with other computers in addition to the library's Apples.) The majority of users were the uninitiated, those who had never before interacted with computers. For them, the *Tutorial* book and disk explained a few of the various abilities of the computers (e.g., graphics, sounds, text), introduced the particular BASIC terminology and principles (e.g., print, precedence, loops), and gave a comfortable familiarity with the machine and its possibilities. As new users read the manual, first they are shown programs already stored on the disk such as games and elementary graphics. Then, users are given elementary programs to type and run, they are encouraged to experiment with various statements in the program and they are taught shortcuts and alternate ways to write statements. Advanced users immediately begin using more-advanced materials—the printer and the double disk drive—as well as using the library's microcomputers for tasks both of a professional and personal nature.

The second goal involved providing adults with the opportunity to develop, refine, and assess various career/educational–related skills. To attain this goal, software packages were purchased to develop both computer and noncomputer skills. The computer-related skills included word processing (*Apple Writer*) and learning different computer languages (*Apple Pascal*, *Apple Fortran*). The noncomputer-related skills included packages to help users prepare for the GED (*Elementary, My Dear Apple*) and develop their typing skills (*Typing Tutor*). Interestingly, many patrons in this group were not concerned with developing computer skills. Rather they viewed the computer merely as a means, similar to books or other educational materials, for learning or developing a necessary skill.

Furthering both the first and second goals of the proposal was the library's decision to utilize the microcomputers as a purely information-dispensing tool. A number of software programs that could be used for primarily information purposes were purchased. Similar information could be found in other materials at the library, but due to the nature of the interaction between machine and user, microcomputer-based information was considered to be quite appropriate and more effective. Unlike previously mentioned users, interested solely in skills development, patrons who planned to use these software programs must have a fairly good understanding of micros and of how they operate. Most of these users have already completed the *Tutorial* program or have had previous computer experience. An example of such a program is an energy- evaluation/saver program (*Energy Miser*), into which the user loads his own personal housing situation and features; the computer allows the user to explore energy-saving options by evaluating various proposed changes the user suggests. Similar programs are available concerning income-tax filing (*Tax Manager*), general money management (*Home Accountant*), and personal record keeping (*Personal Filing System*).

To accomplish the third goal of improving awareness of and access to local educational opportunities, the library staff had to first establish a database of information that it had previously been collecting and printing in the format of a bimonthly calendar. For several years, ACE had collected information from more than 200 schools, agencies, businesses, and individuals concerning current or ongoing courses, classes, seminars, lectures, exhibits—anything of an educational/informational nature for adults. Every two months a calendar describing these upcoming opportunities was published, complete with dates, times, costs, and names of contact persons. The calendar reached a mailing of 1,500 and was available at various locations around the county. It was decided to incorporate this knowledge into a database that could be much more current than the calendar, would allow for easy retrieval of particular course information, and would offer each patron information for his particular needs.

Patrons would be served by the database in three ways: (1) they could peruse a printout of all information currently in the database, (2) they could complete an interest profile sheet and be notified when a pertinent course was listed, and (3) they could request information either by phone or in person concerning a particular subject, field of interest, or organization and receive current information about it. In the most popular method, patrons would complete an interest profile sheet which would inform the librarian of particular subjects or courses of interest to the patron. On a regular basis, library staff would search the database for the indicated areas and notify the patron when appropriate courses appeared. Patrons could change or add to their interest file at any time. In addition to currency and ease of retrieval, it was hoped that the database could record and identify unmet educational needs in the community, needs that could then be referred to and addressed by the appropriate community agencies. Establishing a database such as this was an obvious extension of the already popular community calendar, although it did require extra staff time to organize, input, and search information and special computer materials (*DB Master*, a database management package and Corvus 5 Megabyte, a hard disk drive).

DAILY LIFE

The first months after their arrival, the microcomputers were housed in administrative offices in order for ACE and administrative staff to become familiar with both their operation and their accompanying software. It was important for all ACE staff members to be able to operate the micros so that they could help patrons as the needs arose. After this initial staff training, microcomputers were moved to the ACE department into an area easily supervised from the ACE reference desk.

As soon as the public saw or heard about them, the calls of inquiry began. (Very little real promotion was done; most patrons have heard about them by word of mouth or from a few promotional articles and flyers.) In fact, the microcomputers have increased the department's daily statistics of patrons served by approximately 40 percent. Scheduling procedure for using the micros is not difficult, but it is rigorously enforced. Patrons may reserve a computer for one or two hours each day, and reservations for the computer may only be made one day in advance. Requests are taken either by phone or in person. Hence, there is no long waiting list; everyone has an equal chance to use the computer. Since there is only one double disk drive and one printer, patrons are encouraged to reserve that particular machine if they plan to need those special features. It was quickly obvious that more than two machines were needed. With approval from the LSCA committee, money was diverted from software and print materials into hardware materials and two more Apple II's were purchased. The ACE department then had four microcomputers available for a total of 288 hours of computer time per week. Now, occasionally there are hours when a patron can walk in and use a micro without a reservation, but for the most part reservations are still required.

Because the project is targeted for adult users, children are encouraged either to use the computer in the children's department or to limit themselves to the *Applesoft Tutorial* and other educational and informational programming. A "no games" rule is strictly followed. Actually the staff members have found that some adults are intimidated by the ease with which children understand and use the computers. They also realize that it is very difficult to concentrate upon learning how to use the computer when computer games are being played on an adjacent machine. Thus the staff is very protective of its adult users and does not encourage parents to leave unaccompanied young children at the computers. This is not to say that some children do not use the micros quite appropriately and well. Several children have utilized the computers for school projects and they have even taught the staff some new language and graphics programs. However, it is necessary to consider the role of children in a project like this and to set limits in order to protect somewhat hesitant adults.

In addition to protection, staff members must also offer orientation and troubleshooting for microcomputer users. At their first visit to the microcomputers, patrons are encouraged to seek staff help if they need it, and staff members try to keep aware of users and of their projects. Although the *Applesoft Tutorial* does assume its user has had no previous experience with or knowledge of computers, patrons will usually not understand all that the *Tutorial* explains. Daily, staff members are called upon for further explanation or demonstration of both tutorial and software packages. Although staff members are very willing to help, the difficulty here is twofold. First, it may have been months since the staff member previewed this particular software package, and now called upon, he must quickly review and suggest a possible solution or explanation. Secondly, for patrons doing their own programming, the librarian must first be told and understand the background and logic of the program before any evaluation or suggestions can be made. Thus it is important that each staff member periodically reacquaint himself with both the available software programs and general programming terminology so that he may effectively aid his patrons. Staff members also rely heavily upon books and other library materials to aid them in answering microcomputer questions and in solving programming "bugs."

USER DATA

All users must complete a special registration card upon their first visit to the library's computer center. The information card is designed to help the librarian learn how the computer will be used by the patron and how the librarian can best help the patron to become acquainted and comfortable with the microcomputers. The following is a summary of the first six months' users:

Total number of registered users	571
Total number of individual sessions	1,242
Previous experiences with computers	
No previous experiences	238
College-level course	98
On-the-job experience	91
Own or have access to one	82
Used micro in school	29
Other or no explanation	32
Nature of use	
Personal programming	158
Computer skills	377
Instruction	114
Information	16

A tally of the first six-month service period shows that over 40 percent of users had no previous experience and that most, at least initially, simply wanted computer skills (i.e., to know how micros work).

SURVEY OF USERS

At the end of the first six months of service, a telephone survey of registered users was conducted. Names were randomly chosen and library staff called 130 people reaching 90 of them to complete the questionnaire. Questions on the survey gathered information concerning how patrons had been using and planned to use the computers, the frequency of their use, topics for book and software selection, interest in future workshops, staff support, and general library usage. The survey yielded both interesting and informative data, including the following:

1. Users appreciated the service and planned to continue using it. The top three areas of interest were developing microcomputer skills, personal programming, and instruction:

2. The frequency of patron use depended heavily upon personal time schedules, seasonal changes (during spring and summer, many stated a preference for outdoor or other seasonal activities), and transportation availability (children who were bused to nearby schools could not get to the library during the summer).

3. Users asked for more books (especially introductory and advanced BASIC books) and software, including games (suggested mostly by children) and business languages such as COBOL. Most users seemed pleased with the initial purchases, but noted that books were difficult to get without placing them on reserve.

4. Almost all users requested workshops on some aspects of microcomputers and asked to be informed if the library planned these workshops. Users were especially interested in workshops on general programming skills using BASIC and on home com-

puters and their applications.

5. Staff was viewed as being quite supportive of and available to users, though occasionally lacking in needed expertise.

6. A majority of the users had previously used the library only for books (they said that they had never used library-owned films, magazines, records, or attended any library programs).

ACE is now in the process of incorporating these findings into future plans for the microcomputers. In October, a first workshop about advanced programming was held for microcomputer users and more are planned for the future. (In September, the library cosponsored a statewide workshop with more than 100 participants, aimed at helping librarians make better decisions regarding microcomputers.) Software and print purchases continue to be made on a wider variety of topics, although there are still no plans for buying packaged games. In a recent staff change, a librarian with computer experience was sought and found for a vacated ACE staff position. Plans for the future include beginning a statewide microcomputer users group, cosponsoring a library-based computer camp for children, and possibly offering some computers in the branch libraries.

SOFTWARE

A few words need to be said about the difficulties involved in software purchase for micros. These difficulties include the need to preview software before purchase (one cannot just read a review and buy) and the importance of purchasing software with comprehensive, clearly written documentation. Regarding preview, ACE has found that it is best to view a package in its entirety rather than in the abridged demo form often available from vendors. Only by viewing a complete package can an accurate understanding be gained. If available, librarians should occasionally visit an examination center such as the North Carolina Department of Education's Evaluation Center in order to completely preview the software before purchase. However, this method does hinder the speed and regularity of software purchase. Good documentation is also important to answer any future problems or questions as well as to give a clear understanding of the intended uses and limits of the product. Before purchasing any software, it would be wise to ask the following questions:

1. What does the product actually do? What will a patron obtain from its use?

2. How sophisticated will the user need to be to use the product?

3. How good is the documentation?

4. Is this a product that can be equated with a reference book (i.e., can the product fulfill a patron's needs within the library microcomputer center setting)?

5. What will staff need to know in order to support patrons using the package?

Although even the most cautious buyer will make purchases that do not work as well as he had hoped or expected, careful consideration of these questions will keep such purchases to a minimum.

Although somewhat harried, daily life at the computers is gratifying. In a new sense of the phrase, the library is full of success stories. One patron gained enough confidence and expertise from the computers to obtain a job selling them for a local computer store. Several other patrons have now stopped using the library's center because they learned computers to be valuable enough to invest in one of their own. A GED student, who had once been turned off by learning in high school, was heard chuckling when the computer printed, "That's right! Good for you, Janet," after she correctly finished a division problem. There are also parents glad for the opportunity to keep a step ahead of their children and young adults eager to experiment with the future. One of the best comments about the project came from the telephone survey. One satisfied user said, "I'm glad to see the government finally spending money on something sensible, something I can benefit from. It's a great thing—computers in the library." ■■

APPENDIX A. MATERIALS USED IN THE COMPUTER CENTER

Books

Albrecht, Bob; Finkel, Le Roy; and Brown, Jerald. *Atari BASIC*. New York: Wiley, 1979.

Applesoft Tutorial. Cupertino, Calif.: Apple Computer, Inc., 1981.

Campbell, John L. *Programming the Apple*. Waco, Tex.: Masa Research, 1981.

Castlewitz, David M. *VisiCalc: Home and Office Companion.* Berkeley, Calif.: Osborne/McGraw-Hill, 1982.
Diver, Gerald Van. *Vanloves Apple II/III Software Directory.* Overland Park, Kans.: Vital Information, Inc., 1981.
Dwyer, Thomas. *A Bit of BASIC.* Reading, Mass.: Addison-Wesley, 1980.
Glatzer, Hal. *Introduction to Word Processing.* Berkeley, Calif.: Sybex, 1981.
Inman, Don. *Apple Machine Language.* Reston, Va.: Reston, 1981.
Lewis, Theodore Gyle. *Pascal Programming for the Apple.* Reston, Va.: Reston, 1981.
Moore, Herb; Lower, Judy; and Albrecht, Bob. *Atari Sound and Graphics.* New York: Wiley, 1982.
Nevison, John M. *Little Book of BASIC Style.* Reading, Mass.: Addison-Wesley, 1978.
Poole, Lon. *Apple II User's Guide.* Berkeley, Calif.: Osborne/McGraw-Hill, 1981.

Magazines

Byte. Peterborough, N.H.: Byte Publications.
Compute! The Journal for Progressive Computing. Greensboro, N.C.: Small System Services.
Creative Computing. Morristown, N.J.: Ahl Computing, Inc.
80 Micro: The Magazine for TRS-80 Users.* Peterborough, N.H.: 1001001 Inc.
Softalk. North Hollywood, Calif.: Softalk Publications, Inc.
Personal Computing. Rochelle Park, N.J.: Hayden Publishing Company.

Sheet Music Index on a Microcomputer

Nancy F. Carter

The Music Library of the University of Colorado has acquired, over the years and mostly through gifts, a collection of vocal sheet music consisting of approximately two thousand pieces of music. The sheets date from the middle of the nineteenth century to the present. Some are in good condition, some very poor. It is a typical library collection, not especially valuable or large, but large enough to begin to become useful.

Nancy F. Carter is music cataloger, Music Library, University of Colorado, Boulder.

GROWING NEED

Indexing this material had been discussed many times, but the usual culprits—lack of time and personnel—had prevented any significant progress. The need for such an index was becoming more and more apparent. The study of American music reaches to vernacular as well as classical. Interest in elusive materials such as sheet music is growing and will continue to grow. Professors teaching American music are interested in having the old sheet music available as a valuable tool for studying trends and types of music. Popular music of an era will reflect fads, politics, prejudices, and accepted (or unacceptable) mores of that era. Nostalgia buffs are often asking for music of the thirties or forties or fifties to study or to perform. The index would also provide one more place to look for those elusive song titles. In addition, we hope that knowledge of our index and collection will inspire additional gifts to add to our collection.

PRELIMINARY DECISIONS

The first step in making our index was a general informal gathering of information. Reference personnel in the Music Library—both past and present—were questioned to find out what sort of information patrons needed to help them find a song. Did they often come to the library knowing only a first line or title? Were they interested in composers? Was music of a certain era often asked for?

Next I spoke with a professor of American music to learn what he considered to be the most useful housing arrangement for the music as well as the most useful points of access, considering his assignments and student projects involving the sheet music.

We decided between us that the music would be arranged on the shelves by year, alphabetical by title within each year. The arrangement itself serves as one type of index for anyone who wants to study music of certain years or decades. Secondly, we decided to make indexes for at least first line, title, and composer. Any further decisions would have to wait until we knew more about the computer we would use.

CRITERIA FOR INCLUSION

The sheets of music were scattered

around in various places in the library—envelopes, holding shelves, and unsorted gifts. No previous attempt to organize and index the music proved to be useful in this project. Sifting through the stacks of sheet music, looking at every piece, was the most enjoyable part of the project. That was one reason I kept that entire job for myself; the other reason was to keep the decisions as consistent as possible. Each piece of music was inspected to determine whether it would be included in the indexed collection. Only vocal music was retained. Words had to be in English or have an English translation. A few pieces were in such bad shape that they were pieced together, photocopied, and then discarded. Any incomplete pieces were discarded as well. After I decided that a piece would be kept, I wrote the earliest year of copyright or publication in pencil at the top of the cover. If a piece had no date, ND was written at the top and all those grouped together. At the same time, decisions were made concerning ambiguous titles and composers' names to eliminate confusion for student workers later on. For example, songwriting teams were very common in the twenties and thirties. Information like "words and music by Ned Miller and Chester Cohn" made it difficult to determine who should be designated as composer and who as lyricist. We decided to designate the first one composer and the second one lyricist unless they were names we knew to be in opposite order. (We try to tell our patrons to always search both composer and lyricist indexes.) Also, should Hans von Holtz be listed under H or V? If there is no Name Authority to help, each case must be decided individually.

Often in the old music, a title can be ambiguous. For example, the title page and beginning of score both look like: "He's Got Those Big Blue Eyes Like You, DADDY MINE." Should the title be indexed under "Hes" or under "Daddy"? In most cases we chose the large type, but again it is difficult to make hard-and-fast rules. The best answer would be to have numerous cross-references, but when space is limited or time is short this may not be advantageous. However, some few titles that are better known by some phrase other than the real title were cross-referenced; for example, "Jeanine, I Dream of Lilac Time" was also listed under "Lilac Time." We decided to always use English translations of titles unless it was a phrase common to English usage, like "Auf Wiedershen."

Another decision that had to be made was whether we would use the first line of the verse of a song or the first line of the chorus. We decided to use the first line of the chorus unless it was identical to the title, in which case we would use the verse.

After all these decisions were made, the piece of music was placed on the proper stack, one for each decade. We were fortunate to have an empty office for the summer, which served admirably for a place to stack and sort. There was no need to clean up and put things away between opportunities to work on the music.

COMPUTER DECISIONS

Once we had settled on what was wanted and needed in an index, we could discuss computer usage in detail rather than in generalities. After some discussion of using the university's large computer, we had decided to go with the library's microcomputer, which is a Radio Shack, Model II. The file maintenance program called "Profile," which is an off-the-shelf program by Radio Shack, seemed to be just the vehicle to do our job for us. Thus, our index required no programming, only a few hours spent adjusting the system to our data. The bulk of time, therefore, was spent on data entry.

The decision to use the microcomputer brought both advantages and disadvantages. The advantages were: convenience; personal attention from the computer's operator; his desire to show that the computer could be used advantageously for in-house projects such as this one; and cost. The drawbacks were those consequences of limited space. Considering that each record would contain six fields (four of them indexed), there was little room left for comments. Also, alphabetizing can be done only within each disk. Thus, if two disks are used, two alphabetical lists of, say, titles will be generated.

Working with the operator, we came up with decisions we felt both of us could live with. We had 132 printing spaces to divide among at least six fields. We decided to include year, title, composer's last name, composer's first name, lyricist's last name, and first line. After allotting spaces for vital information, we had only ten spaces left for comments. We used these comments to indicate that the text was given in an additional language besides English (+ in Fre); a third person shared responsibility for composition (+ Evans); the song is a Christmas song (*Christmas); the song originated from a musical show (*Showboat); the library owns copies in two keys (+ High/Low)

Indexes would be generated using four fields: title, composer's last name, lyricist's last name, and first line. We considered using publisher as well, but the other fields were more important to us. An index would not be made using year since our shelf arrangement would cover that need. Figure 1, a portion of a page from the title index, shows how the finished product looks.

Each computer disk holds approximately 800–1,000 records. Since we would have to use at least two to three disks for our collection of 2,000 anyway, we decided to use enough to leave room for future expansion. We used one disk for each of six time divisions:

1800s and ND
1900–1909
1910–1919
1920–1929
1930–1949
1950–

Each division contains 300–400 records. Thus, using six disks leaves enough room to double our collection and still keep our indexes divided by decades. This division works as an advantage if a patron wishes to peruse music of the twenties or thirties.

WORK SHEET

To reduce errors and to prevent the necessity of carrying the actual sheet music to the computer room, a work sheet was designed to record information for input. Each line on the work sheet was marked to indicate the maximum number of characters allowed for that field (see figure 2).

Two students went through each "decade" stack transferring all information from each sheet of music to a work sheet. Any questions arising concerning the information could be resolved here in the library rather than wasting time at the terminal.

INPUT

My student assistant, John Wilkens, was trained in a very short time to input the information into the terminal. The job was completed in about four weeks on a part-time basis. As each decade was finished, a trial printing, alphabetized by title, was run. This enabled us to catch duplicates and correct errors. The corrected programs were then ready for the final run.

PRINTING

After all information was input and corrected, we were ready for final printing. This was done by the programmer and re-

MUSIC LIBRARY
SHEET MUSIC INDEX
1930-1949

PAGE 1
Wed Dec 8 1982

TITLE	COMPOSER		LYRICIST	FIRST LINE	YEAR	COMMENTS
HAPPY TALK	RODGERS	RICHARD	HAMMERSTEI	HAPPY TALK KEEP TALKING TALK A	1949	*SOUTH PAC
HAVE A LITTLE FAITH IN ME	WARREN	HARRY	LEWIS	I WISH I COULD SAY	1930	+YOUNG
HE SHALL JUDGE THY PEOPLE	LAFORGE	FRANK	BIBLE	HE SHALL JUDGE THY PEOPLE WITH	1942	
HEBRIDEAN SONG	TEMPLETON	ALEC	WORDSWORTH	NO NIGHTINGALE DID EVER CHANT	1939	
HILLS OF OLD WYOMIN'	RAINGER	RALPH	ROBIN	LET ME RIDE ON A TRAIL	1936	
HOLD BACK THE DAY	CEPPARO	LEON	COLONNA	MAN AND A MAID IN THE MOONLIGH	1932	
HOLIDAY FOR STRINGS	ROSE	DAVID	GALLOP	WHEN I SEE YOU SMILE AT ME I H	1943	
HOLY BOY	IRELAND	JOHN	BROWN	LOWLY LAD IN A MANGER	1938	*CHRISTMAS
HOME	VAN STEEDEN	PETER	CLARKSON	EVENING MARKS THE CLOSE OF DAY	1931	
HOSANNA TO THE SON OF DAVID	CARR	ARTHUR	BIBLE	AND HIS DECIPLES WORSHIPPED HI	1947	+HIGH/LOW
HOW ARE THINGS IN GLOCCA MORRA	LANE	BURTON	HARBURG	I HEAR A BIRD LONDONDERRY BIRD	1946	*FINIAN'S
HOW DO I LOVE THEE	LIPPE	EDOUARD	BROWNING	HOW DO I LOVE THEE LET ME COUN	1941	
HUMMING BLUES	TEMPLETON	ALEC	TEMPLETON	CAN'T YOU HEAR THE HUMMING BLU	1939	
I CAN'T BEGIN TO TELL YOU	MONACO	JAMES	GORDON	I NEVER HAVE A DREAM THAT I DO	1945	
I COVER THE WATERFRONT	GREEN	JOHN	HEYMAN	AWAY FROM THE CITY THAT HURTS	1933	
I DON'T MIND WALKIN' IN THE RA	RICH	MAX	HOFFMAN	IT'S RAINING I HOPE IT RAINS A	1930	
I DON'T WANT TO WALK WITHOUT Y	STYNE	JULE	LOESSER	ALL OUR FRIENDS KEEP KNOCKING	1941	
I HAVE EYES	RAINGER	RALPH	ROBIN	IF I WERE IN THE MOONLIGHT ON	1938	
I LIKE MOUNTAIN MUSIC	WELDON	FRANK	CAVANAUGH	WHEN THE SUN GOES DOWN	1933	
I LIKE TO DO THINGS FOR YOU	AGER	MILTON	YELLEN	THO' YOU ARE A GREAT BIG GROWN	1930	
I LOVE YOU TRUELY	JACOBS-BOND	CARRIE	JACOBS-BON	I LOVE YOU TRUELY, TRUELY DEAR	1938	
I PROMISE YOU	OAKLAND	BEN	LERNER	MY LOVE FOR YOU IS AS GREAT	1938	
I SAW A MAIDEN FAIR	BUSH	GRACE	Trad.	I SAW A MAIDEN FAIR SITTING	1942	*CHRISTMAS
I SAW THE PLACE	ENNERS	HENRIETTA	TIPPETT	I SAW THE PLACE WHERE JESUS PR	1945	
I SHOULD CARE	CAHN	SAMMY	STORDAHL	I KNOW I SHOULD PITY ME BUT I	1944	+WESTON
I WILL GIVE THANKS UNTO THEE	TIPTON	CAMBELL	BIBLE	I WILL GIVE THANKS UNTO THEE	1936	
I WILL LIFT UP MINE EYES UNTO	PARKER	CLIFTON	BIBLE	I WILL LIFT UP MINE EYES UNTO	1934	
I WISH YOU LOVE	TRENET	CHARLES	TRENET	I WISH YOU BLUEBIRDS IN THE SP	1946	+IN FRE

Fig. 1. Portion of Page from Title Index.

TITLE THAT'S FOR ME

COMPOSER'S LAST NAME RODGERS

COMPOSER'S FIRST NAME RICHARD

YEAR 1945

LYRICIST HAMMERSTEI

FIRST LINE I SAW YOU STANDING IN THE SUN

COMMENTS *STATE FAI

Fig. 2. Input Work Sheet.

quired more than a day's time. Each disk, representing one of the time periods, was used to produce four alphabetical lists—composer, lyricist, title, first line.

After receiving the printouts, I separated the sheets to put all title lists together in ascending order of year, all composer lists the same, etc. These lists were then placed in two printout-sized notebooks. Title and first-line lists went together in one notebook, composer and lyricist lists in another. Tabs were attached to facilitate finding a list for a certain date. If a patron knows an approximate year for a piece of music, then this arrangement works very well. The worse that can happen is that the patron will have to search through six alphabetical lists all of which are in one notebook. The arrangement seems to be working well for us.

STORAGE

Students helped organize each "decade" stack of music into individual years, then alphabetized within each year. Acid-free storage boxes have been ordered to house the sheet-music collection. The boxes will be labeled by year(s) and stacked on the shelf in our Special Collections area. Alphabetical arrangement by title within each year makes access convenient.

PATRON USE

In the few months that our index has been available, patrons and staff have made varied use of it. We used it to find some Christmas music for a Christmas display we were mounting in the library. A voice student who wanted to sing a medley of forties and fifties songs at her parents' anniversary party was very pleased with our arrangement. Senior citizens have been poring over it lately to find songs about travel for a local contest. It helped a history student locate the words and music for "Lili Marlene." We do not allow the music to be checked out. It may be used in the library and it may be photocopied.

The index was finished too late to be included in any professor's lesson plans and assignments for the fall semester of 1982. We have advertised the index by mounting a sheet-music display in our lobby and by mention in our quarterly newsletter, *Jots and Titles*. Thus we hope that our professors will soon include the sheet-music collection as a resource suggestion to their students.

UPDATE

We plan to update our index every summer when people and machines have less demands on them. This will depend, however, on how much new music we are able to add to our collection to make the update worth while.

NEXT PROJECT

Next summer we hope to have the time and resources to do a similar index for our piano sheet music. This project will be simpler than the vocal-music project since lyricist and first-line categories will not apply. ■■

PERLINE at Risley

Taube Marks

Managing a scientific and technical serials collection with a small but able staff is by no means a new scenario. Claiming and checking-in are time-consuming (if not tedious) jobs when done diligently, but even diligence is not enough to ensure that all journals are received and claimed before it is "too late." If staff have to cope with the extra burden of routing, the management problems become even more difficult. When the library serves a highly select group of physicists and engineers dealing with the complex problems of nuclear power development and nuclear safety on behalf of the entire United Kingdom, it is essential that the library run the most efficient information and library service possible.

Blackwell's PERLINE has been assisting the library at UKAEA NPDE Risley (United Kingdom Atomic Energy Authority Nuclear Power Development Establishment) since July 1982, to carefully husband resources while providing a more accurate method of handling checking-in and claiming. With the aid of a DEC PDP 11/23–based computer and Blackwell-supplied software called PEARL, Risley NPDE staff record arrivals, generate claims, and have a complete record of serials and their past receipt. The library has improved financial information, an easy-to-use file of all serials subscriptions, and a means for producing and transmitting correspondence (be it a claim, a query, order, etc.) electronically to Blackwell's.

After an initial three-day training period in Oxford with Frank Graham, the senior librarian, and his two assistants, Jan Wells and Angela Ratcliffe, the task of preparing the data for entry began. Jan and Angela created an input sheet for each of the 900 + journals (1800 + subscriptions, 28,000 + journal parts per year) noting the precise title, alternate titles, ISSN, notes, frequency, issue notation, cycle of publication, and supplier. The entry sheets were sent to Blackwell Technical Services (BTS) where current and previous prices, order numbers, and further publisher information were added. A keying exercise was undertaken by BTS on behalf of Risley, to create the first serial file. However, before serial input could begin, it was necessary to construct auxiliary parameter files for publisher, frequency cycle, distributor, and suppliers (since the software and hardware relates to all serials regardless of vendor).

File building for future libraries will be easier since universal title, frequency, publisher, and supplier information can be reused so that the next library will need only key local data to complete a record. Machine-readable records will be available by means of the telecommunications software FIBER (a front-end processor that works in conjunction with Blackwell's ICL 2956 main frame), as will be the facility to access Blackwell's Book and Periodicals and Periodicals Subscription File by teletype-compatible terminals.

With the aid of PERLINE, the staff are, issue by issue, abandoning their "well matured" Kalamazoo cards for checking-in with the aid of the minicomputer and a terminal. They will no longer have to thumb through cards to determine, often after the arrival of the next issue when it is too late to claim, that an item is overdue. Rather, they can simply review claims by viewing an overdue file under the following categories:

- overdue issues awaiting claim
- damaged issues awaiting claim
- missing issues awaiting claim
- outstanding claims
- overdue claims
- all categories

These can be subdivided further under the following headings:

- publisher
- supplier
- subscription type
- serial
- all serials

Retrieval of serials for checking in (or browsing) is simple yet flexible so that skilled and unskilled searcher alike can find the title with the minimum amount of keying. Search keys include ISSN and ISBN (with all controls validated on entry), co-

Taube Marks is research assistant, Blackwell Technical Services, Ltd., Oxford, England.

den, locally generated code, title, and up to two alternative titles. Title searches are normally based on four-letter truncations of up to the first five significant words in the title. The use of abbreviated common words can extend the search beyond the first five words.

The check-in module (as opposed to the complete serial record) has been designed to present the user with the minimum information consistent with the correct identification of the serial. Attention is drawn to the center of the screen where details of the next expected issue are displayed. If the issue received is the issue predicted, two keystrokes complete the transaction. If the item received is not the item indicated, the keying of the issue specification records the item in hand and the system, at the same time, calculates the missing issues. The user has the opportunity to claim the issues or defer the claim until later. While checking in, the printer generates routing slips, complete with reader/patron names, addresses, and other details. The patron file can be adjusted at each check-in to bypass or resequence the routing list. User-friendly alerts notify staff that the number of receipts has exceeded the number of subscriptions when extra copies are checked in.

For each issue there is a historical record of past journal receipts with the title, issue specification, publication date, expected date, date received, quantity, and status. In time, the library builds an accurate record of the delivery and publication pattern of each journal and can see an entire year's receipt (and more) in an orderly display with action and status displayed.

All housekeeping activities are online without reference to Blackwell's main frame in Oxford. Risley need not worry whether the main frame is available. They can use the system twenty-four hours a day, seven days a week. This is important since there are occasions when they work in the evening and on Sunday. However, access to Blackwell's is simple, and they can use the combination of PEARL and FIBER software to send claims and messages and receive acknowledgments, letters, updates (as in the Periodicals Information Bulletin), and other information. On the same dial-up line, they can query the Blackwell Periodical File, check the status of their orders, and place orders for new and/or additional subscriptions, sample copies, and monographs.

The installation of PERLINE at Risley is not only the first PERLINE site, but also the first completely automated serials control system in the UK. PERLINE offers the possibility of an automated network for serials control among larger organizations with remote sites, and gives the user the greatest degree of flexibility and control over his/her own serial database. It also offers online assistance to smaller libraries previously locked out of automation on a cost basis. ■■

In-Depth: University of California MELVYL, 2

This special section concludes the series of articles begun in the December 1982 issue of ITAL.

Computing Resources for an Online Catalog

Clifford A. Lynch, Mary E. Engle, and Mark H. Needleman

Library automation has come of age. As part of that maturity, it has become clear that successful library automation projects often require dedicated resources on a large scale. Computing support for major library automation projects is often taken for granted—it is either preordained (because the project is using a computer center established for other purposes) or so highly constrained by the scope and budget of the project that there are few choices.

In the near future, as large library automation projects become more common, computing support will be a key issue. The finest user interface can be subverted by unreliable computing support or poor response time. Cumbersome and forbidding sign-on sequences can predispose a user to reject a system out of hand. At the same time, technological developments are now rapidly changing once-visionary proposals into economically viable realities. This has certainly been the case at the University of California, where the online catalog systems specified in the 1977 *Plan for Development*[1] would have been much too costly several years earlier, in 1970, but by 1980 were cost-effective compared to existing manual approaches. Online catalogs will become one of the major artifacts created by 1980s information-processing technology.

Clifford A. Lynch is manager of the Computing Resources Group, Mary E. Engle is an ADP systems analyst, and Mark H. Needleman is a senior programmer, Division of Library Automation, University of California.

INITIAL DECISIONS

In the late 1970s, it was rare to see a computer center set up from scratch to support a new application. In designing the computing support for the UC online catalog, we had an unusual advantage: there was no history to enchain us. There was no hardware in place. No real decisions had been made about the structure of the online system, although there was a sizable investment in the development of PL/1-based batch programs for database consolidation, normalization, and loading.

The commitment to dedicated computing resources for the catalog meant that the catalog came first, last, and foremost. When designing online applications that use hardware installed primarily for other activities, capacity planning is often impossible since capacity is prone to reassignment by administrative fiat.[2] Fortunately, we did not face this problem, and therefore could plan system capacity and growth rationally. We could first develop an application that met user needs on a relatively small scale, then tune as we enlarged. Control of response time was in our own hands.

We really faced only two constraints. The first was to preserve our investment in batch PL/1 programs, representing years of work. The second was our timetable and budget. The authors of the *Plan for Development* were not designing paper systems: the plan had timetables for implementation, not just vague statements about research and development. Open-ended research and development was a luxury we could not afford. While we could optimize the system for our application, it had to be constructed from off-the-shelf technology.

The university has nine campuses spread over the state of California with almost 100 libraries. DLA first had to decide where to put computing resources in relation to the libraries and how to deliver service from the computer center (or centers) to the end users.

There are two classic models for computing configurations: distributed and centralized. Both presuppose telecommunications networks, but the networks have very different characteristics. In the distributed processing approach, computers would be put on some or all of the nine campuses, each running a copy of the catalog and database. In the centralized or monolithic approach, one large processor would run all systems and applications software. This computer would be centrally placed, and terminals on each campus would tie into it in a star network.

The Promises and Pitfalls of Distributed Processing

At first glance, distributed processing has appeal. Since the costs of computing power have dropped faster than the costs of telecommunications, it might be reasonable to put local data and computing power at each site, or perhaps at four or five locations, with some of the smaller campuses sharing processors. Assuming that most searches are for books on the local campus, segmenting the database according to each campus' holdings would minimize telecommunications costs and give the best response time. The library patron would still be able to access other campus' holdings, when necessary, via telecommunications links. This approach would provide a very robust system: if one machine failed, most queries still could be satisfied at each of the other campuses. So the hypothesis goes.

Unfortunately, the nature of the application does not lend itself to this type of structure. The largest portion of our database is composed not of bibliographic records themselves, but of index files, such as subject and name indexes—the access points to the data. These indexes are indispensable to every user, regardless of location. The database cannot reasonably be split, since each user must be able to access all of the indexes. We would have had to replicate virtually all of the index structure with each local copy of the database. This would mean putting billions of bytes of online storage at each processor, which did not appear to be cost-effective given our projections for disk-storage prices.

In addition, the notion that the distributed system would be more reliable proved somewhat illusory. The online catalog was to become a cornerstone of the UC libraries' operation, and the idea that a single processor failure could disable all catalog activity on a campus the size of Berkeley or UCLA was intolerable. Although we could have put up two copies of each campus' data—one at the local site and one as a backup at another campus—the cost would have been extreme.

It should also be noted that in 1979 we knew of no database management software commercially available that could implement a distributed and partially duplicated database with query routing—at least in any established form. Even today, while much discussion is given to such systems, outside of research establishments there is very little operational software for them. Some major experimental systems such as Computer Corporation of America's SDD-1 have appeared, but they are still research projects.[3]

Vexing management and operational problems also confronted us in a distributed model. Distances between campuses posed logistical problems beyond simple inconvenience for DLA staff, and the alternative—duplicating at least part of the systems and operations staff at each site—seemed no more tractable. Distributed processing might have been fashionable and "futuristic" at the time, but we decided it was poorly suited to our needs, both operationally and economically.

Centralized Processing

The time-honored approach of acquiring a single monolithic processor would resolve the economic and operational problems inherent in the distributed model, but reliability would still be a weak point.

And there were other problems. All of our analysis led us to believe that the catalog would be a totally I/O-bound application; that is, the speed and efficiency of the application would be limited by the speed and efficiency of the input and output operations. It is an elementary but all too often ignored fact of capacity planning that *all processors wait at the same speed*. Thus, a single fast central processing unit (CPU) would spend much of its time waiting for disk reads, making relatively poor use of the costly CPU resource. Using a machine priced for its high CPU speed in an application that made little use of its computational power was a waste.

In the initial stages of developing a system as novel as the online catalog, the favorable reaction of the user community is of prime concern. Once lost, user confidence is very difficult to restore. It would be difficult to test both the systems and applications software throughout the development phases with a single processor since we would frequently have to take our only machine out of service. Catalog availability would be greatly reduced and failure recovery delayed if hardware components had to be replaced. We feared that inevitable downtime and erratic availability during the development phases would lead the user community to believe that the catalog was not a reliable tool.

The Local Processor Complex

We needed an alternative approach that would combine the cost-effectiveness and manageability of the centralized model with the promised reliability of the distributed paradigm. We achieved the necessary compromise by using a number of moderately sized, relatively inexpensive processors (CPUs) linked in a single processor complex at a central location.

In this model, several processors at one location are connected, all sharing access to the data stored on multiple-access disks. If there are hardware problems with one machine, the catalog can be switched to another. To the end user this should appear as nearly 100 percent uptime. With this system we have gained the reliability that we sought.

For maintenance, a specific processor can be drained of work so that no new catalog sessions are scheduled. Multiple machines allow rapid and easy testing not available with other alternatives. The result has been very little disruption during development and, hopefully, greater end-user confidence in the catalog's reliability.

With all processors in a single location, it becomes both necessary and cost-effective to install support facilities to insulate the site from environmental problems. A Halon fire-protection system protects our machine room and tape-storage area. Nightly, each disk is copied to tape, and weekly the tapes are sent to off-site storage. Since even brief power outages or sudden fluctuations in voltage can disrupt service, we have introduced an uninterruptible power supply (UPS). The UPS both conditions the current and backs it up in the event of a brief power failure. Ultimately, additional steps such as fossil fuel generators and fully redundant air-conditioning may also be justified.

CHOICE OF MACHINE ARCHITECTURE

The size of the processors was dictated by our belief that the catalog would be highly I/O intensive. Because main-memory costs were dropping rapidly, we designed the system to make heavy use of main memory to reduce I/O activity wherever possible. Thus, the catalog would also be highly memory-intensive (it currently operates in about four to five megabytes), necessitating either large minicomputers or medium-size mainframes. In addition, we felt we had to have multiple independent I/O paths (i.e., channels or I/O processors) in order to support the traffic load on the database. This led us to favor medium-scale mainframes rather than large minicomputers, since mainframes usually offer such independent channel capabilities.

Having made some basic decisions about processor size, memory size, and I/O requirements, our next consideration was computer architecture. In the field of library automation today, many different machine types are in use, with varying degrees of success. We chose the IBM 360/370 architecture for a number of reasons. Much of the programming we had already done (which would become the support and maintenance programs for the online catalog) was written in PL/1, a language that is well supported on IBM machines. Also, the

IBM architecture is a common industry standard and is relatively stable and well documented. It is easy to hire people familiar with the IBM 360/370 environment.

Of equal importance is the vast amount of software written for the 360/370, some of it by IBM and some by other groups such as universities, independent software sources, or members of user groups. Much of this software is available at little or no cost. The software in use at DLA that is not provided by IBM includes MILTEN, a telecommunications system; WYLBUR, a powerful text-editor and job-submission system; and SPITBOL, a SNOBOL4 compiler, an advanced high-performance string processing language. Because of this diversity, we had considerable choice on the one hand and the ability to build on previous work and investment on the other.

Over the past decade, both IBM and its user community have made an enormous financial commitment to the 360/370 architecture. This, we felt, ensured that it would be actively with us for many years to come. We were further encouraged by the fact that IBM was not the sole source of machines with 360/370 architecture. Today's plug-compatible machine manufacturer (PCM) market is extensive, allowing us greater freedom in procurement and some degree of independence from the vagaries of any single vendor's marketing strategy. We preferred not to rely on a single supplier, favoring the advantage of choosing among many PCMs for competitive pricing and technology.

Having defined our computing requirements rather specifically, we set out to select our hardware from among the options available in 1979. We needed 4341-class machines. This choice was based on physical size constraints and price/performance considerations. We wanted as much incremental upgradability as possible, so we could adapt the machines closely to our needs by adding channels and memory as events dictated.

In 1979, when we requested bids, there were really only two vendors with 4341-class machines in production: IBM and Magnuson. IBM's 4341 processor was limited to four megabytes of memory and six channels. We did not feel comfortable with these constraints (which have been removed since then) or with the delivery dates being quoted (in excess of a year during the initial order period for IBM's 4341).

Magnuson, on the other hand, offered competitively priced processors in their M/80 series expandable to sixteen megabytes of main memory and sixteen I/O channels. Magnuson offered other advantages as well. Their machine made heavy use of microcode. This not only provided a hedge against obsolescence but also opened the intriguing possibility of using custom microcode to tune the machines to our specific application.

Magnuson also offered all of the traditional problems of a small, high-technology company: financial, management, and technological uncertainty. (Magnuson was, in fact, still privately held when we bought our first two machines.) The company did not have a massive support staff, which perhaps concerned us less than some other customers; given our operating system decisions (discussed below), we required relatively little support other than basic hardware maintenance. We felt that we could live with these risks. To date, while Magnuson has in fact weathered considerable financial and management turmoil, their machines have continued to function reliably and to meet DLA's needs.

The DLA computer complex now includes:

- Three Magnuson M80/43 CPUs, each with eight megabytes of main memory and six I/O channels
- Eight Control Data Corporation 33302 disk drives, totaling 1.6 gigabytes of storage
- Fourteen Control Data Corporation 33502 disk drives, totaling 8.9 gigabytes of storage

Each disk is accessible by two completely separate channel/controller paths from each of the three CPUs. The three CPUs are fully interconnected with channel-to-channel adapters. In addition, the DLA complex includes various tape drives and printers and a pair of COMTEN 3670 communicating controllers.

SELECTING AN OPERATING SYSTEM

Having chosen the IBM architecture, our next major set of decisions concerned the software environment. A key component of this environment is the machine's operating system. This is the program that "runs" the computer. It provides job and resource scheduling and management, input-output management, device management, error recovery, security, accounting, timing, and a host of other services.

IBM has several different operating systems for the 370-generation machines and plug-compatible CPUs. Among the most prominent are OS/360, OS/VS2 (MVS), VS1, and VM/370.[4] All of these, except for OS/360, belong to the family of virtual memory systems that are very much in fashion today and are widely used both in industry and in education.

Our choice of OS/360 for the operating system was both controversial and unusual in light of the fact that it was an older operating system no longer supported by IBM. Originally developed in the mid-1960s for the System 360 architecture, and last released in 1972, it sprang from a generation of basic systems software that preceded the concept of virtual memory. Newer operating systems have long since dominated the large-scale IBM installations. An examination of the virtues of OS/360 in the context of DLA's computing environment will help put this decision in perspective.

Stability of the system and control of our own software destiny were paramount considerations. To keep the user community's confidence, the system had to be dependable. With more than ten years of use at a great many computing installations, OS/360 had become both stable and reliable. IBM's newer operating systems, MVS and VM, are still changing. Any change to a program as complex as an operating system introduces potential destabilization. Local modifications must be reworked to fit each new release, and in some cases interfaces to applications programs may also be affected. Coping with this constant change puts great demands on the systems programming staff and disrupts applications development and public service. We preferred to take responsibility for operating system stability, rather than relying on a vendor. A number of our staff already had considerable experience with OS/360 at other installations, and we believed that we could implement it quickly, maintain it competently, and tailor it to our needs.

OS/360 is in the public domain; that is, it is essentially free from IBM. It has a host of supporting software, including compilers, modifications, utilities, teleprocessing systems, and performance monitors. Much of this software was written by user installations and is readily available from computer user groups such as SHARE. The financial benefits to this are significant. The savings that we realized in operating-system and supporting-systems software gave us more money for critical hardware: processors, disks, and telecommunications.

In light of our choice of a local processor complex configuration of four to five CPUs, OS/360 spared us from having to buy expensive licenses (for a system like MVS, for example) for each CPU. We can plan for the future unencumbered by extemporaneous conversions or the capacity and compatibility problems that haunt the systems programmer, the installation planner, and the budget manager as each new release appears. Since we literally know where we stand with OS/360, our development and production schedules need not suffer excessive disruption, and we can carry out our mission supported by a relatively small systems-programming staff.

Since we lack vendor support for OS/360, we always face the danger of encountering problems we cannot resolve. We have decided to meet this challenge by building a staff capable of handling quickly and efficiently the problems that inevitably arise. We have not, to date, encountered software problems that we could not resolve.

One of the advantages of the new operating systems is that they are continually updated to support new computer technology such as graphics terminals and laser printers. Although with OS/360 little of this software support will be forthcoming from the vendor, we feel that we will be able to add it when there is good reason to do so. This is a valuable discipline from a management point of view, since it forces us to evaluate carefully the benefits of new devices rather than being swept along with the excitement of being among the first installations to use a new piece of hardware.

So far we have discussed the managerial and financial advantages in running OS/360. There were substantial technical advantages as well. The current trend in operating-system design is toward systems such as MVS, which excel in sharing a very large processor among a very large number of concurrent tasks. In order to do this, they rely on techniques such as virtual memory, which have a rather substantial fixed overhead. We wanted to dedicate processors to a few very large applications. In this situation, the overheads introduced in the newer operating systems are not offset by the ability to share hardware among a large number of tasks. OS/360 offers a viable alternative to this bind.

The ability to modify the system at will gives us major performance benefits. Unquestionably, the newer operating systems have a number of performance improvements (partially to compensate for new overheads); we have retrofitted many of these into OS/360 without the accompanying overhead. For our applications mix, the system we are now running performs perhaps 40 to 50 percent better than "standard" OS/360.

In addition to OS/360, we chose to implement an additional level of control through the Asymmetric Multiprocessing System (ASP),[5] a program of roughly the same vintage as OS/360 (a successor system, JES3, is available for the MVS environment). This allows us to control all of our processors as a single, loosely coupled processor complex with easy and largely automated recovery from hardware failures. Because of this, we need only a single operator per shift for our entire computing complex.

TELECOMMUNICATIONS SOFTWARE

Along with the operating system, the telecommunications software is crucial to the online catalog. It transmits all of the data the users type at their terminals to the applications program (the catalog) and ensures that the catalog's responses are returned to the proper terminal. In theory, the catalog program could handle these services itself. However, this would require very close interaction with the operating system and the hardware, not easily done in a high-level language like PL/1 (in which the catalog is written). The catalog program would have to have been much more complex, requiring more development time and staff.

We decided on MILTEN, an existing telecommunications software system originally written at Stanford University as part of the MILTEN/WYLBUR/ORVYL system.[6] MILTEN was already in use at DLA with the WYLBUR text-editing system, on which we were developing our software. MILTEN enables terminals to switch between different applications programs on the same computer very easily, so that we can support test versions of the catalog simultaneously with the production version. MILTEN also can support future expansion as new UC library automation systems are implemented. Besides transmitting data to and from terminals, MILTEN provides sign-on/sign-off procedures, security checking, statistical data on individual terminal use, and control functions for the telecommunications network.

Another major advantage to MILTEN is its availability in IBM 370 source code, enabling us to modify as needed to support the requirements of a particular library environment. Since MILTEN now supports only ASCII and 2741-type terminals, it is small, parsimonious with computing resources, and manageable. It is important to note that the only terminals we intended to support were generic ASCII terminals and generic packet-switched terminals, although we expected to provide a high level of support and control for these. This ruled out many commercial systems oriented to 3270-type terminals.

DLA has modified MILTEN to support the online catalog by adding new and more user-friendly messages and broadcast facilities for the public catalog terminals, by improving diagnostic facilities to support the catalog, and by defining dedicated catalog lines that do not require manual sign-on/sign-off sequences. Most commercial systems are weak in these features, although the need for them cannot be overstressed. In a public-access system, a complex log-on sequence is enough to drive many users away. We have also begun to improve support for users who dial the catalog from their personal microcomputers, which are becoming increasingly commonplace.

Packet-Switching: A Path for Growth

The *Plan for Development* mandates a network of more than 1,000 terminals. In order to implement such a network, our current telecommunications network will have to be expanded and redesigned. We will discuss the software considerations here. An accompanying article on telecommunications will discuss the network.

In our present network, each terminal is connected to the CPU across a dedicated path of its own. This path consists of a port on the communications controller front-end and a subchannel address on the CPU channel (actually, in order to provide redundancy, we duplicate this subchannel connection on each CPU in our complex).

This system has worked well for the prototype with about 100 terminals on the campuses, but it cannot support the full future network. IBM 370 architecture has a maximum of 256 I/O devices (subchannels) per channel. Each set of 256 terminals needs a channel on each processor, 256 ports on the communications controller, and 256 subchannel interfaces to each processor. There are technical problems with supporting multiple channels through a single copy of MILTEN, but the dominant considerations are economic. Channels are costly, and since the number of channels that can be connected to a given machine is limited, adding more channels for MILTEN would reduce the number we could add in other areas.

This would also require adding vast numbers of ports on the communications controller, since an additional port would be required for each new terminal connected. Also, because traditional telephone communications can transmit only a limited number of data paths per phone line, we would have to add a significant number of phone lines to the campuses along with the companion hardware that controls the data flow at both ends. This represents a substantial expense without even providing the redundancy that a large-scale network needs for acceptable service. For a network of this size, the cost of providing a backup for each critical segment in the data-transmission path is prohibitive.

We plan instead to convert our telecommunications network to packet-switched technology. A packet is essentially data to be transmitted over the network (in our case either the information that the catalog sends to the user or the responses the user types at the terminal), surrounded by transmission control information such as origin and destination information, routing information, error and redundancy checks, and receipt acknowledgements. With the origin/destination information embedded in the packets, we will need only one port and CPU connection to control multiple terminals, reaping benefits in simplicity, reliability, and low cost.

There are several packet-switching protocols currently available, the most prominent being the X.25 protocol[7] and IBM's System Network Architecture (SNA).[8] A protocol is simply a set of rules to perform various communications functions in a standardized fashion. We selected a packet-switched network protocol called TCP/IP (Transmission Control Protocol/Internet Protocol), developed for the Department of Defense ARPANET network designed by Bolt Beranek and Newman.[9-11]

The ARPANET technology has been in use for about ten years and has proven reliable. It has many features absent in the other protocols, among them dynamic routing—where each packet responds to current network traffic, providing load leveling. Packets are automatically rerouted by alternate paths to bypass nodes that have failed. The TCP/IP protocol allows internetwork communication, so that a packet can be sent from one network to a final destination at a node in an external network, passing through one or more intermediate networks along the way. Additionally, it supports communication with other networks using different protocols and among different types of computers within the same network.

While the ARPANET protocols have the benefit of being vendor-independent and are consequently very flexible, they require considerable software development to implement in our host computers. Here, again, our ability to freely modify OS/360 and MILTEN to mesh with this new software is very important.

SOFTWARE PACKAGES FOR THE ONLINE CATALOG

In an ideal world we could have purchased much more of the base software for the online catalog rather than writing it ourselves. But due to the unusual features of the online catalog, we could not find adequate commercial software.

One of the key differences between the UC catalog (and public information retrieval systems in general) and more typical business-oriented terminal systems is that the UC catalog is a conversational rather than a transaction-oriented system. That is, it involves an open-ended dialogue with the user, with a great deal of the context carried by the system throughout the session.

For example, a user may set parameters in a session that limit a search to a specific campus' holdings or that select just one of three possible display modes. MELVYL remembers your last command so that entering HELP at any given point will return an explanation of where you are in a search and what you can do next. Three errors in a row cause the system to provide HELP automatically.

In a transaction-oriented system, on the other hand, you typically fill in a number of blanks on a screen (or respond to prompts), the data is processed, and a result displayed, then you start from scratch with a new transaction. In order to provide a user-friendly public-access system, we felt that a heavily conversational bias was mandatory.

The easiest way to develop terminal applications in an IBM world is to write programs that deal primarily with a single user, and then use a transaction monitor package to thread multiple users through the single-user program. This approach greatly reduces the complexity that the applications programmer must handle, simplifies debugging, and isolates other users from any problem in the software that a given user's session may uncover. There are a number of transaction-processing systems on the market designed for high-volume handling of large terminal networks; perhaps the best known is IBM's CICS. Unfortunately, the name "transaction processing" is accurate: these systems are designed to handle transactional rather than conversational interactions. For our situation, they also would have needed extensive changes in order to interface to MILTEN.

We were therefore faced with the archetypal system designer's problem: we could purchase something that would not really meet our needs, or we could develop a custom system to do exactly what we wanted. Since the software at issue was the user interface—the heart of the system, and the key to its success or failure—we decided on a custom system.

The Multi-User Interface

This led to some rather unusual applications development. The applications program interfaced directly to MILTEN, and as such became aware that it was dealing with many users, rather than just "seeing" one and letting the transaction monitor be concerned with multi-threading the program. Therefore, we had to develop a rather complex and delicate software module called a "dispatcher" or "controller" to keep track of the many users on the catalog at any given moment and to decide which terminal to service next. While this required substantial programming, the end result was an extremely successful user interface: a system that because it was custom designed could support a large number of terminals with minimal overhead. It remains to be seen how difficult it will be to maintain and modify this dispatcher module as demands on it change. Certainly it has been one of the more complex and error-prone parts of the system to date; problems tend to be sensitive to timing and hard to reproduce, but far from unmanageable.

The Database Management System

We also needed a database management system (DBMS). Again, there were a number of database systems on the market that could do most of what we needed. None were perfect for our application, but if the system was to be ready on schedule, we did not have time to develop our own DBMS along with everything else. Since the user does not interact with the DBMS system directly, we thought we could mask unpleasant features by programming around them in the user interface. We selected Software AG's ADABAS system, a network data model DBMS that has been in the field for a number of years and appeared relatively mature and stable.[12] It had been used successfully in bibliographic applications, most notably at the Washington Library Network.

ADABAS has in fact worked, but its implementation and use have not been easy. We have had to program around the system's shortcomings both in the user interface and in our database design. Not surprisingly, it takes much longer to fix ADABAS problems than those in the systems we maintain ourselves. Because we do not have access to the source code for ADABAS, performance and tuning have been difficult, and at this point ADABAS and the disk I/O subsystem are the major bottlenecks in our catalog's performance—it is often hard to find out which of the two is really causing the problem.

ADABAS is the major resource consumer in the MILTEN/CATALOG/ADABAS suite of programs. The following shows resource consumption for a typical (fairly light) afternoon's catalog use:

	CPU Time	*I/O Operations*	*Region*
MILTEN	889 secs.	496,572	250 K
ADABAS	6,370 secs.	490,028	1500 K
CATALOG	1,683 secs.	2,000	3650 K

THE PROBLEMS OF OVERHEAD AND RESPONSE TIME

At this point DLA has a successful system in operation, and we are now working to expand it from a 100-terminal prototype to a full production system. This involves some changes to the database design, conversion to packet-switching telecommunications, and purchase of more terminals and computing hardware.

We have two major objectives in system configuration, tuning, and development: reducing overhead and improving response time. These are distinct but related issues. Our success with them will determine both the ongoing success of the system and the cost of expanding it. A number of technological developments on the horizon will affect these issues considerably.

Reducing system overhead means effectively reducing the CPU time and I/O operations necessary to carry out a search or, more generally, the time to process the response to one prompt and issue the next. (Our overhead is fairly low except in actual search processing, and the possibilities for reducing overhead in other areas seem limited.) Reducing overhead can also be thought of as processing more transactions per minute of computer time, which in turn means processing more users per CPU.

Improving response time, on the other hand, means reducing the elapsed time between initiating a search and having the results in hand. It is important to understand that reducing CPU overhead will not necessarily reduce response time unless the machine is running close to capacity. Recent monitoring at DLA suggests that CPU overhead is not our immediate problem. Statistics from a typical period show the following pattern of CPU utilization (again, a relatively quiet day):

Operating System	1.0% of CPU
ADABAS	17.0% of CPU
Catalog	4.7% of CPU
MILTEN	2.5% of CPU
Available (free for other purposes)	74.8% of CPU

In periods of heavy use, our observation is that all figures extrapolate linearly, except ADABAS CPU time, which grows somewhat faster.

These figures suggest that response time (which is currently running in the five- to ten-second range for typical searches) is our major problem right now. Overhead will become important mainly in figuring the cost of expanding the system. At present, we could either double or halve our CPU overhead with only a negligible effect on response time. Put another way, buying a faster CPU won't provide faster response for our users: all CPUs wait at the same speed.

The Disk Bottleneck

The bottleneck for response time is the

disk subsystem. In order to improve response time, we will have to make disk accesses either faster or fewer.

Overcoming the disk bottleneck has been a major theme in computer systems for library automation since the 1960s. Malinconico[13] gives a fascinating historical perspective on this problem, which, as we shall see, may still be the major obstacle to implementing successful large-scale library-automation systems.

As long as we use ADABAS, reducing the number of disk accesses does not appear to be promising. We are already giving the DBMS as much main memory as it seems to be able to use constructively for buffering and are circumventing it for some operations. With changes in the file design, we could reduce the number of disk reads, but we would need a lot more disk space (which we presently cannot afford or house). Short of re-examining and replacing the algorithms the DBMS uses to access the database, which would involve replacing the DBMS, there seems to be little we can do to reduce disk accesses. There is an alternative that might also help response time—carrying out more disk accesses in parallel for a given search—but this is equally inaccessible to us since we cannot control ADABAS' access and search resolution algorithms.

An alternative is to make disk accesses faster. There are four major components to disk-access time: queuing delay, seek time, latency time, and transfer time.

We intend to add more channels and controllers to reduce queuing. The new generation of disks coming onto the market now (IBM 3380 and equivalents) offer faster transfer time (3 megabytes/second as opposed to 1.5 megabytes/second) and improved seek and latency time. The faster transfer rates will further reduce queuing delay. We project a continued, although not radical, improvement in this technology over the coming years. However, other approaches that may be available in the next few years could further alleviate these bottlenecks.

Drums

Drums or their modern solid-state analogs are one option for faster access to data. Drums eliminate seek time by placing a read head over each track on the device. Solid-state memories improve matters further by not using rotating media at all, eliminating latency time as well as seek time, and speeding access to the point where channel interface, electrical protocol, and transfer rate become the limiting factors. Our problem with these approaches is that well over a gigabyte of our database has very high access, and at current prices it would cost too much to house that much data on drums or solid-state memory. While we expect these costs to drop, it seems unlikely that several gigabytes of such storage will become affordable in the near future.

Database Machines

A far more radical alternative is some form of database machine. There are several of these on the market today, although none of them exploits the concept to its fullest. The concept of a database machine[14] involves the following considerations.

First, when a DBMS is running on a general-purpose CPU along with everything else, the CPU spends much of its time deciding what to do next because of the high frequency of I/O interrupts generated by the DBMS. Each time a block of an index is read in, the DBMS must be dispatched by the operating system, decide which block to read next, initiate another read, and pass control back to the operating system so that it can continue doing whatever it was doing before the read was completed.

Second, while a CPU is a general-purpose device, the DBMS makes use of only a small subset of the facilities available on the CPU. Thus, an improvement in a small group of facilities (or a custom hardware implementation of them) would greatly reduce the CPU time used by the DBMS.

Finally, by using special file organizations, it may be possible to place intelligence in the disk subsystem itself so that much of the search logic can be carried out in the disk area itself, without involving the CPU.

Unfortunately, these techniques are directed primarily to overhead reduction.

Only the final one, which implies building more intelligent and consequently "faster" disks (for specific types of access), offers DLA any real reduction in response time. While a database machine might very well reduce overhead, it might actually increase response time unless it made heavy use of intelligent disks (which none of the existing systems, such as Software AG's database machine, actually use). A database machine would create other problems, notably a single point of failure for the whole disk subsystem. At present, although we intend to keep a close watch on new developments, we do not feel that database machines offer the answer to improving response time.

Improving Performance through Microcode

Returning to the overhead problem, one very simple solution is to purchase faster CPUs. Undoubtedly, as the price/performance ratio continues to shift in our favor, we will increase the speed of our machines. As the proponents of the database machine concept have pointed out, however, this will not necessarily result in large reductions in CPU utilization.

For example, a faster floating point processor would have almost no effect on our response time or our overhead, since our application does not emphasize calculation. Consequently, while faster processors will help our overall system performance (since, for example, our batch loading and database maintenance programs will benefit substantially), we may need to seek methods to reduce overhead that are more specific to the problem than CPU upgrades but cheaper and less radical than database machines.

An attractive and very reasonable alternative to the database machine is to use microcoding to build into the CPU some of the functions that would be performed by a database machine. As mentioned earlier, the Magnuson is a microcoded machine. This means that the CPU uses a very low-level programming language that talks directly to the hardware to simulate the functions of the IBM 360/370 architecture. This is a very flexible approach in that as the 370 architecture expands in definition, new microcode can be added to the machine to perform those enhanced functions.

This has intriguing implications for the online catalog at DLA. Having customized both the operating system and application software, we now see the possibility of rewriting certain critical sections of them in microcode to greatly increase throughput and system efficiency. This is another instance in which our choice of an unsupported operating system like OS/360 has turned out to be of benefit. We can add microcode without fearing the need to redo it for new operating system releases.

With regard to the applications program, certain heavily exercised portions of the catalog may lend themselves to being microcoded in order to improve their efficiency. Microcoding can also be used to extend the IBM machine instruction set with new instructions that would take over current software functions. These new instructions could then be incorporated into existing and new software, replacing current routines. This would greatly increase the speed and efficiency of the existing software and would eliminate the need to rewrite the same or similar routines in many different programs. Since the functions are similar but not exactly the same in the several programs, the microcode design would have to be general enough for overall applicability, yet small and fast enough to justify its existence. There are design limitations in what can be done and how large a microcode routine can be. But microcoding extended features does present an exciting alternative for various types of queue searching and maintenance, binary tree searching and maintenance, and table scanning and lookup.

Microcoding also shows promise for system monitoring. Since it operates at a very low machine level, it is possible to write routines that monitor how the hardware itself is performing. It is also possible to get below the software and monitor its behavior as well. This would add measurement and tuning tools at a level not present in most of the available software packages or, if present, only at the cost of severely degrading system throughput and response time.

Microcoding has potential for monitoring I/O subsystems. The DLA system uses disks that are shared among multiple CPUs and accessed by multiple paths, both to protect against failures and to increase overall system throughput by allowing multiple simultaneous access. We expect that microcoded diagnostic routines at the hardware level can give us basic statistics about channel loading on the various I/O channels, and the frequency with which specific I/O operations return busy responses from a particular channel and must be retransmitted along an alternate path. These diagnostics will help us configure and tune our I/O subsystem to provide the fastest possible access to all disks in the system.

As microcoding begins to play a larger role in our application, our choice of the Magnuson M80 machine has been sustained, since it is heavily microcoded and can be adapted to provide many of the features discussed above. As a small company, Magnuson can be expected to be more responsive to customers. Also, since many of the features we need may be interesting to others, Magnuson may develop them to improve their product and increase their share of the plug-compatible market. In fact, Magnuson has already started to move in this direction.

Reducing Disk Costs

Data Compression

As stated above, building an online catalog requires massive amounts of disk storage to hold the data and indexes. Since the IBM EBCDIC character code uses an eight-bit pattern to describe each character, the cost for storing data on a disk is fixed. An alternative way to store data is to use a variable-length representation code with an algorithm that weighs the frequency of occurrence of characters and stores more frequently occurring characters in fewer bits. Unfortunately, these bit-oriented compression and decompression algorithms run very slowly on typical byte-oriented processors. Performing compression and decompression in software, therefore, would be much too slow, considering the amount of such activity necessary for savings of any magnitude with a large bibliographic database.

Using microcoding for this data conversion is a feasible alternative. It would be possible to write microcoded instructions to convert data between a compressed format and the standard EBCDIC character representations. This has two major advantages: first, data compression and decompression can occur at much faster speeds, making it feasible to implement; second, it becomes relatively transparent to the application program, which at most would only have to call a small subroutine to compress and decompress data. Thus, we can use data compression in many applications without recoding or repackaging the rather complicated routines for each application that wanted to use them.

Microcode is not the complete answer to this problem since in order to achieve the necessary data decompression speed, special hardware may be needed to support the microcoded instructions. In addition, to take full advantage of the benefits of data compression, the DBMS would have to be modified to exploit the data compression routines.[15]

Optical Disks

Since the cost of computing power has rapidly declined, the dominant costs of a computing facility such as DLA's are now telecommunications and data storage. While data compression can reduce data-storage requirements as much as 30 to 40 percent, ultimately a cheaper storage medium must be found if library automation is to be economically viable for electronic document storage and retrieval.

While the price of magnetic storage will continue to drop, the cost reduction needed for new developments requires a new medium such as optical disks. These offer the possibility of huge cost savings in systems like DLA's and introduce the potential for exciting new types of systems as well.

Optical disks provide much cheaper storage for nonvolatile data—that is, for data that changes very little over time. In our case, the bibliographic records themselves change little over time, if at all, compared to the index records. Although new bibliographic records are continually added to the database, once there, they are not constantly updated as are the index entries, for example. Optical disks are attractive because of high data capacity (with a density thirty times that of magnetic media), low price, archival quality (since they are unaffected by magnetic or radiation disturbance), and high data-transfer rates. Additionally, they impose few environmental demands since the storage surface is encased in plexiglass and thus is unaffected by dust, smoke, and fingerprints.[16]

Data is stored on optical disks with a laser beam that melts pits into a thin tellurium layer on the disk surface. The pits create a different reflectivity than the unburned surface when a lower-powered laser beam is passed over them. Once a pit is burned it cannot be erased or undone, so, at present, the optical disk is primarily a write-once medium. There are several experimental projects under way to allow multiple writing of optical disks, but the perfection of these techniques appears to be at least two years away. There are also various theoretical techniques that would allow a certain number of repeated writes on write-once optical disks by using specialized encoding techniques,[17] but to the best of our knowledge these have yet to be exploited commercially. Currently, both write-once and repeated-write optical disks exist in prototype, and we expect to see the first generation commercially available in late 1983 or early 1984.

The optical disk poses several interesting problems as a storage medium. The laser requires a settling (or focusing) time each time it is moved to do a read. Once positioned, however, data transfer is extremely fast. This means that the optical disk behaves very much like a magnetic disk with a very long seek time and a very fast data-transfer rate. Current index structures, which are very efficient on magnetic disks, may not work well on optical disks, and new data structures may have to be developed to exploit the devices to their fullest if they are to be used to store indexes as well as data. In addition, it is not at all clear how to integrate optical disks and magnetic disks into a unified medium for housing a database. However, the potential cost savings from converting some of the database to optical media are too significant to be ignored.

DIRECTIONS FOR COMPUTING TECHNOLOGY

In building the UC online catalog we made a number of discoveries about what was and was not available on the marketplace. We have discussed how current technological developments might help DLA address its immediate problems: weaknesses in commercially available database systems, the disk bottleneck, and the costs of massive amounts of disk storage. All of these are active development areas in the marketplace.

In the course of designing and implementing our system, however, we have identified many facilities simply not available in the IBM world, as far as we could determine. In many cases they can be found in other architectures, particularly in minicomputer environments. Such facilities could have a major impact on how processor complexes are assembled and what they would cost. In many cases, our inability to obtain such facilities has led to reduced reliability, needless complexity, or increased costs in our system.

Fault Tolerance

In general, too little attention is paid to fault tolerance in the IBM world. Any major error seems to bring down the system. Fault tolerance is the ability to either reconfigure around the error or restart the system automatically. We have found that redundant CPUs are an excellent approach to this problem, although there are serious problems in the amount of manual switching necessary to move from a failed processor to a functioning one and altogether too much discontinuity. There is no reason why processors in a complex cannot monitor and restart each other when a failure is detected. IBM's recent announcement of its Remote Operator Console Facility for the 4341 has at last removed the major barrier to such support.

Similarly, there is no reason why software/hardware failure monitors cannot be built into machines so that they could restart themselves after a failure. Such features would greatly reduce the time required to restore a complex to operating status and eliminate much of the operational complexity in a local processor complex. The same considerations apply to communications front-end processors, which seem to be designed with a great disregard for automated restart, switching, and graceful degradation capabilities.

Disks present a special and more serious problem. With the reappearance of fixed media devices such as the 3350 and 3380, more and more data can be contained in each disk unit. The cost of providing redundant disks is generally prohibitive. Any failure in the electronics or media generally renders massive amounts of the database inaccessible. DLA still maintains mountable media (3330 type disks) for certain critical operating-system data, but the cost and floor-space limitations prohibit us from using this approach to the database itself. Virtually all of our catalog outages of more than a few minutes are due to disk failures. We are in need of a better means of recovery from disk failures.

Local Processor Complexes and Shared Memory

The basic software for local processor complexes (ASP/JES3) is quite good. However, the channel-to-channel adapter normally used to connect CPUs into a local complex becomes a major problem when one wishes to connect four or five processors for full redundancy (five processors require ten channel-to-channel adapters!). Some form of integrated channel-to-channel adapter/switching device would alleviate the problem. IBM's recent 3088 system coupler announcement appears to be a major step in this direction.

There is a desperate need for a shared-memory capability between loosely coupled processors. The new possibilities that this would open up are almost endless: sophisticated recovery and fault-tolerant processor complexes, low-overhead distributed computing and data sharing, multi-CPU database update with minimal problems, and the ability to share very large tables.

Intelligent Disks

Serious consideration should be given to moving some of the basic DBMS functions out into the disk subsystem. The trend to date is primarily the development of special-purpose CPUs to act as database machines. This may be a reflection of the marketplace: since only IBM currently manufactures the full line of equipment from operating system to CPU and disk subsystem, development by any company other than IBM would require collaborative efforts.

CONCLUSION

This article has tried to give some insight into what led DLA to configure its computing resources as they are now and to evaluate the choices that were fruitful and those that were not. In the main we feel we chose well for our purposes, given the technology available. Much of our experience seems to run counter to generally accepted thinking within the industry; for example, our experience with OS/360, MILTEN, and a multi-user applications program has been very successful. Our experience with commercial software has been somewhat less so.

In some cases, such as the disk bottleneck, we face the same obstacles that much of the rest of the industry does. In other areas, such as in our need for shared memory, our needs seem to be different from those of the industry as a whole. Perhaps we are a bellwether: certainly the reliability and performance requirements of public-access information-retrieval systems are at least as rigorous as those of many commercial systems. We certainly feel that if we could obtain some of the technology mentioned above, we could build a better system (or, more accurately, evolve a better one).

Our decisions, given their historical context, produced a viable system. Some approaches we rejected may well work for others. Our system is still evolving, and many of the decisions we have discussed will be re-examined in coming years, if for no other reason than that the effective life span of most computing equipment is only five to seven years. New developments in computer technology make things possible in library automation that could not have been envisioned a short while ago. UC's online catalog is a testament to the rate of change computing is bringing about in library automation.

It is interesting to reflect on the interaction between user needs as we have enunciated them, technological developments, and the reification of technology in the marketplace by vendors in the industry. We had to satisfy many of our needs through our own efforts. Other needs were satisfied accidentally; that is, we benefited from changes demanded by other industry segments in the user community.

It is clear that library automation and computing have a great deal to offer each other. To date the transfer has been mostly in one direction: library automation has exploited new computing developments. With the rise of large-scale public-access information-retrieval systems, it appears that library automation, because of its immense scale and rigorous service-level requirements, must continually push the state of the art in computing. It seems surprising that so few people outside the library community realize that library automation applications offer not only a fertile field for new directions but also one of the finest proving grounds for the computing developments of the coming decade.

REFERENCES

1. University of California, Office of the Executive Director of Universitywide Library Planning, *The University of California Libraries: A Plan for Development, 1978–1988* (Berkeley: The University, 1977), 210p.
2. For a discussion of this problem, see Clifford Lynch and Edwin Brownrigg, "Implementing Library Automation Plans in a University Computing Environment," *Planning for Computing in Higher Education* (Boulder, Colo.: Westview Press, 1980), p.215–25.
3. J. B. Rothone and others, "Introduction to a System for Distributed Databases (SDD-1)," *Association for Computing Machinery, Transactions on Database Systems* 5, no.1:1–17 (March 1980).
4. For an IBM view of the evolution of these systems and the references cited there for details on each operating system, see *IBM Journal of Research and Development* 25, no.5 (Sept. 1981).
5. IBM Corporation, *IBM System/360 and System/370 ASP Version 3 Asymmetric Multiprocessing System: General Information Manual* (White Plains, N.Y.: The Corporation, 1972), 45p. (Order no. GH20-1173-1)
6. R. Fajman and J. Borgelt, "WYLBUR: An Interactive Text Editor and Remote Job Entry System," *Communications of the ACM* 16, no.5:314–22 (May 1973).
7. CCITT, "Provisional Recommendation X.25," *Data Communications Standards* (New York: McGraw-Hill, 1978).
8. IBM Corporation, *Systems Network Architecture Format and Protocol Reference Manual: Architectural Logic* (White Plains, N.Y.: The Corporation, 1978). (Order no. SC30-3112)
9. *ARPANET Protocol Handbook* (Menlo Park, Calif.: Network Information Center, SRI International, 1982), 549p.
10. *INTERNET Transition Workbook* (Menlo Park, Calif.: Network Information Center, SRI International, 1982), 512p.
11. Jonathan B. Postel, Carl A. Sunshine, and Danny Cohen, "The ARPA Internet Protocol," *Computer Networks* 5:261–71 (July 1981).
12. Software AG, *ADABAS General Information Manual* (Reston, Va.: Software AG, 1979), 52p.
13. Michael S. Malinconico, "Mass Storage Technology and File Organization," *Journal of Library Automation* 13, no.2:77–87 (June 1980).
14. For a general discussion of database-machine concepts, see Fred Maryansk, "Backend Database Systems," *Computing Surveys* 12, no.1:3–27 (March 1980) and *Computer* 12, no.3 (March 1980). Theme issue on database machines.
15. For a more detailed discussion of data compression for bibliographic databases, see Clifford Lynch and Edwin Brownrigg, "Applications of Data Compression Techniques to a Large Bibliographic Database," *Proceedings of the 7th International Conference on Very Large Databases* (Cannes, France: IEEE Computer Society Press, 1981), p.435–47.
16. David Maier, "Using Write-Once Memory for Database Storage," *Proceedings of the Association for Computing Machinery, Symposium on Principles of Database Systems* (New York: ACM, 1982), p.239–46.
17. Ronald L. Rivest and Adi Shamir, "How to Reuse a 'Write-Once' Memory," *Association for Computing Machinery, 1982 Fourteenth Annual ACM Symposium on Theory of Computing, San Francisco, California, May 5–7, 1982* (New York: ACM, 1982), p.105–13. ■■

Telecommunications for an Online Catalog

David L. Shaughnessy and Clifford A. Lynch

The University of California (UC) has made a major long-term commitment to the concept of an online union catalog for public access. In designing a telecommunications network to support the catalog, we had to consider the entire context of the UC telecommunications, computing, and library service environments—not only as they exist today but as they will evolve in the coming decade. The widespread use of computer networking will significantly alter the way the university operates in the 1980s. The library system is part of the core of the university, and the catalog is the key to the library. A major goal of conversion to an online catalog was to increase access to library holdings; similarly, the goal of the supporting network had to be widespread accessibility to the new online catalog. Our challenge was to plan and deploy a network capable of growing and meshing with other university facilities in the years to come.

Thus, we approached the design of DLA's communications network with not only the three standard considerations of budget, technical requirements—such as terminal types to be supported, and operational reliability, but also a rather unusual concern for adaptability to the changing overall university computer and communications environment.

This article examines both the specifications and the environment for DLA's telecommunications network and provides a history of and prospectus for its implementation.

THE UNIVERSITY ENVIRONMENT

It is vital to understand the university environment as it affects the telecommunications system, since this environment has a major impact on both the underlying transmission media and the protocols we have chosen for our network.

Telecommunications and computer networking at UC are not standardized. This is, in many ways, natural for a very large, multicampus research institution, but the implications of such a catholic approach can be serious for a large public access network covering the nine campuses.

University-wide telecommunications planning is carried out by the Office of the Assistant Vice-President—Information Systems and Computing, within Systemwide Administration. Historically, there have been a number of studies of intercampus telecommunications needs, but little implementation of facilities for data traffic. Even today, while the university is seriously re-examining telecommunications in light of increasing costs and new technology, the primary focus at the system-wide level remains voice. This emphasis is based on current costs and usage. Although the consensus is that system-wide data communications will be very important in the future, DLA's activities are the first concrete manifestation of this future. As the major system-wide data communications user, DLA is active in the University Telecommunications Planning Task Force.

Intercampus data communications is still largely unexplored; while there are many special-purpose links from one campus to another, DLA will represent the first system-wide network. Intracampus communications facilities—local networks and single-campus computer networks—are blossoming throughout the UC system. In some cases, even individual departments are building their own networks. Typically, these local activities are administered or coordinated by computer centers and telecommunications staff on each campus. Systemwide Administration is involved only in an advisory role.

The adoption of a uniform approach to telecommunications, as has occurred in many large corporations, is not workable in UC's environment, at least for the near future. The state of the art is changing rapidly, and a consensus on the "right" way to

---

David L. Shaughnessy is a development engineer and Clifford A. Lynch is manager of the Computing Resources Group, both with the Division of Library Automation, University of California.

do things does not yet exist about telecommunications standards. DLA has neither the authority nor the inclination to establish standards by fiat. But it is clear that computer-communication networks will play an increasingly critical role in the university's research and instructional activities. For the online catalog to reach its full potential, it must be linked to this environment, bringing the library directly to the broadest user community.

Perhaps here it is appropriate to review the University of California library system as a whole. There are nine campuses distributed about the state, four of them in the north central area (along with the UC Systemwide Administration offices where DLA is located), and the remaining five in the south (see figure 1). Among the campuses, there is great diversity in the distribution of library facilities. The Berkeley campus, for example, has more than thirty branches located in different buildings. The San Francisco campus has its libraries in two physically connected buildings. The Santa Cruz campus holdings are concentrated in two buildings at some distance from one another, and the library at Scripps Institution of Oceanography is located miles from the central library at the San Diego campus. Thus each campus presents a different set of "last mile" problems—delivering data from the campus terminus of the long-haul network to the actual terminals.

Variations in campus library facilities further complicate planning. Libraries are among the first buildings to be erected on a campus; many of them are very old. Only a few of the newer buildings were designed to accommodate large numbers of terminals. The old brick and stone monuments offer severely limited numbers of twisted pairs (telephone wiring) for connecting dispersed terminals to multiplexers or concentrators. If a rooftop microwave or satellite dish is installed, costly new cable conduits may be necessary. Even electrical receptacles for terminals may be a problem in some cases. Yet, since these buildings are usually the main libraries on their campuses, they typically need large numbers of terminals early in the deployment cycle.

Local facilities are managed locally.

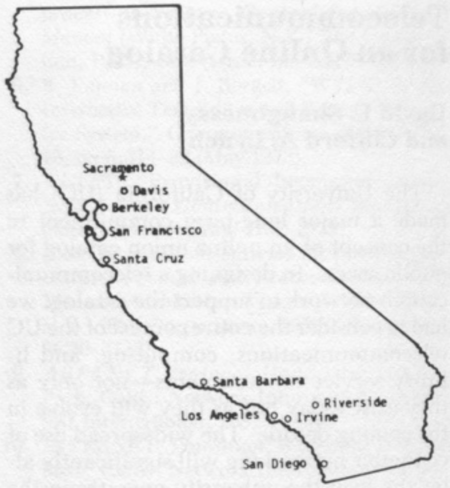

Fig. 1. Distribution of University of California Campuses.

Each campus has independent telecommunications and facilities management offices. DLA has no authority to request changes in local facilities or telecommunications directly; rather such requisitions must be made by each library to its local campus service groups on DLA's behalf, considerably complicating both technical and administrative aspects of campus installations.

BASIC SPECIFICATIONS AND REQUIREMENTS

The basic terminal for MELVYL, the online catalog, is any half-duplex, line-by-line ASCII terminal. The reasons for choosing simple ASCII terminals are twofold. First, the online catalog was developed to accommodate untrained (and possibly computer-intimidated) users. We wanted terminals that were as simple as possible. We assumed that most patrons had at least passing familiarity with a typewriter, and that the closer the terminal was to a typewriter, the more comfortable the user would be with it. Despite the rising popularity of video games, we felt that complex, full-screen terminals like the IBM 3270 might be threatening to the user; in fact, we have observed programmers of many years' experience reduced to helplessness by the complexities of systems based on such terminals.

Economics also supported the use of ASCII terminals. They are inexpensive, particularly in large lots, and prices have dropped steadily over the last few years. (UC buys vast quantities of such terminals; we have been able to take advantage of quantity pricing agreements negotiated by the UC purchasing department.) They are readily serviceable and easy to install. From a facilities point of view, they are tractable—for example, they do not require coaxial cable or cluster controllers, which would further aggravate the wiring problems in older libraries. They make fairly modest use of bandwidth, given our application. In this regard, it is important to remember that we send perhaps fifty lines out from the computer for each line typed in by the user. There is no filling out of forms, for example, where the ability to read only entered or altered fields would significantly reduce data transfer.

Our choice of ASCII terminals did, however, add to the complexity of our telecommunications system. ASCII terminals are manufactured by many vendors. Unlike the IBM 3270, where a single vendor's product acts as a de facto standard, the ASCII terminal is far from standardized. ASCII terminal specifications address only the electronic interface (RS-232) and the transmission code for printable characters (ASCII code). In fact, ANSI (American National Standards Institute) standards X3.4 and X3.41 only address character codes, and specifically decline to assign any standard meaning to special control and graphics characters or to escape sequences. What little commonality exists in these sequences is due to the prevalence of a few popular terminals (such as the DEC VT100 series) in the marketplace, and there are even ASCII terminals on the market now that emulate the characteristics of other vendors' popular ASCII terminals. This is largely because most system software in use today cannot cope with the extreme diversity of terminal characteristics.

Because of this diversity, we had to provide a very fine level of parameterization in the telecommunications software. The system must be aware of, for instance, the line size, the number of lines per screen, and the screen-clear sequence for each type of terminal. Hardware tabs, where available, should be exploited both to reduce the number of characters sent and to improve display speed. For hard-copy terminals, idles may be needed at times to compensate for mechanical limitations in the print mechanism: these should be set only when required to allow the terminal to print as fast as possible. Such considerations are essential in developing a high-quality user interface.

Our preliminary tests indicated that 120 characters per second (cps) would be an adequate speed for data transmission. Normally each screen of data is preceded by a screen clear and then written down from the top of the screen to avoid having the text scroll upwards while the screen is being filled. This allows the user to begin reading the screen while it is still being written. We find that most users, particularly casual ones, are comfortable with 120 cps under these circumstances. They cannot read much faster than this.*

We expected considerable use of the catalog from the increasingly ubiquitous personal computer, and ASCII terminal support allows us to handle this group of users nicely, since almost every microcomputer on the market has an ASCII terminal emulator available for it. In fact, ASCII transmission protocols represent the lowest common denominator compatible with most UC computers and terminals. Potentially, then, MELVYL would be accessible to this entire terminal population (possibly several

*There is actually considerable debate about optimal speed, with proponents of both slower—30 cps—and faster display rates. In the prototype network we had to standardize on a single speed to control complexity. In the production network, the terminal's data rate is agreed upon by the Terminal Access Controller (TAC) and the terminal, and the packet-switching protocols offer sufficiently good flow control to allow us to overcommit trunk bandwidth. In the future, users may be able to set the speed at which they want the terminal to run. Alternatively, we can probably run most terminals at 960 cps, if desired (the limitation being wire lengths between the terminals and the TACs), providing a bursty but very fast display.

thousand terminals) as soon as the appropriate switching, port selection, and local network interface facilities were in place.

Besides terminal types, there were other basic requirements. In order to provide a user-friendly system, we had to avoid complex, cumbersome log-on sequences and dialogues. Messages indicating system unavailability had to be simple and helpful. The system had to be insensitive to activities at remote locations (such as terminals being disconnected or powered on and off). Note that this final requirement strongly discouraged chains of terminals that were polled in some fashion, since an interruption anywhere in the chain could disable many terminals.

Other than these needs, and the general requirement for high reliability and low cost, we came to the network design with few a priori requirements. In the interests of maximum flexibility and compatibility, however, we were also inclined to avoid standards unique to any single main-frame vendor. Staying with generally available, multi-sourced equipment also helped insulate us from planned obsolescence, or from disruptions motivated by any one vendor's marketing strategy.

THE PROTOTYPE NETWORK

Just as the online catalog went through a pre-prototype period for library staff training and evaluation, and a prototype period to observe operating characteristics and obtain feedback from patrons and library staff, so too the deployment of the communications network was to go through stages in its development. At the outset, DLA was faced with a major dilemma. Data communications networks do not scale up well. In a small network with eight or ten terminals clustered at each campus, the major objective is to reduce the total cost of the hardware to support those terminals. Multiplexers, for instance, have a base cost for, say, eight terminals and an incremental cost per terminal; the base, rather than the incremental cost, dominates. Since line bandwidth was being purchased in 9600 baud (960 cps) quanta, optimizing line usage had little economic effect. In a production network with hundreds of terminals at each campus, the economics turn upside down. Doubling or tripling the base cost for the cluster makes little difference, since this cost is swamped by the incremental terminal connection factor and the cost of bandwidth.

A network capable of economically supporting the ultimate configuration of more than 1,000 widely scattered terminals either had to have this capability designed in from the the beginning (leading to a tremendously expensive prototype, since equipment geared to providing economies of scale with a huge number of terminals would be amortized over a very small terminal population), or the prototype and final networks would have to be of different designs. In the latter case, a conversion from the prototype to the production network would be required.

Because of our profound lack of knowledge about the catalog's telecommunications-related operating characteristics, we decided on conventional methods for the initial network. This was economical for a small network and seemed relatively straightforward to implement. Most important, it would provide us with operational data on which to base the design of the production network. We selected equipment with as much flexibility as possible so that we could reuse it where practical, but we did not plan to have it fit into a predetermined slot in the production network. At that time we had no assurance that our needs could be met within budget constraints by any of the existing data communications technologies.

We had already learned from earlier experiences that obtaining conventional data circuits from the various telephone companies that serve the UC campuses is not as simple as might be expected. When the phone company at DLA on the Berkeley campus had to deal with a different phone company serving one of the campuses, a new level of cooperative planning was needed. In order to complete a connection, the engineered data path must reach a common interchange point on the border of the two service areas. In the first attempt to provide service to the Santa Barbara campus, the two telephone companies failed to connect with one another, and the line had to be re-engineered. We also learned that DLA was the first UC department to attempt networking on this scale through the system-wide telecommunications office. Previously, each campus department needing data communications had worked through its local campus telecommunications office. No procedures or forms had been developed to coordinate orders, or to track or recharge the cost, for example, of local area circuits on remote campuses.

To support the prototype, we settled on a simple "star" network. This is not really a network in the true sense of the word. One location (DLA) serves as the center hub of a group of data paths radiating out to each campus like the spokes of a wheel (see figure 2). It is the easiest design to implement and control, but the most vulnerable to failures. Only costly total duplication provides an alternate path in the event of a line or facility failure. To compensate for this inherent design limitation, we elected to install equipment capable of recovering and correcting for all but complete failure of the data path.

The typical terminal installation for the prototype consists of a cluster of about eight terminals connected to DLA over a 9600-baud leased line through a multiplexer and high-speed modem. The terminals run at 120 characters per second. Given an error-

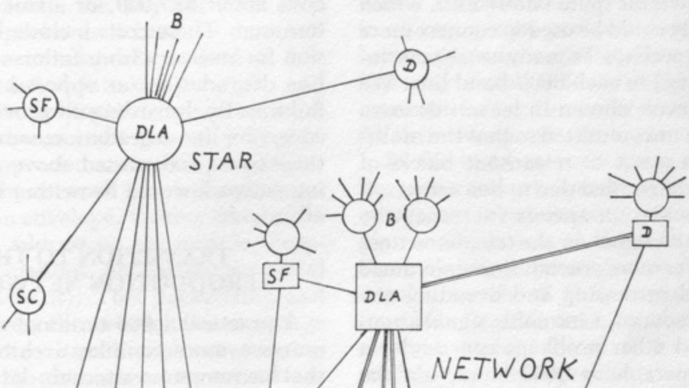

Double lines are for heavier traffic.

Fig. 2. Models of MELVYL Prototype and Production Configurations.

free transmission line, this configuration would permit a virtually continuous flow of data to or from all eight terminals at once, using the time-honored time division multiplexing technique. This is a means of sharing a line among several lower-speed devices by partitioning the line capacity into equal slots based on the number of terminals, irrelevant of activity or lack of activity on a given terminal.

In fact, the flow of data is far from continuous. This allows us to exploit a more sophisticated line-sharing scheme called statistical multiplexing, whereby terminals are dynamically assigned line capacity when they have data to send or receive. The statistical multiplexing technique leaves us with considerable spare bandwidth, which theoretically could be used to connect more terminals (perhaps as many as sixteen instead of eight) to each 9600-baud line. We have, however, chosen to leave this extra bandwidth uncommitted so that the multiplexers can use it to retransmit blocks of data that are garbled due to line errors.

As transmission speeds increase, the quality of the signal on the telephone lines becomes ever more crucial, since tolerances in the modem sensing and decoding circuitry are reduced. Line noise, signal attenuation, and other problems can render a circuit inoperable at 9600 baud, but the line may be marginally acceptable at a slower speed such as 4800 baud. Because of this, we selected high-speed modems that have an alternate, slower transmission speed. Such modems allow us to make further use of the extra capacity afforded by statistical multiplexing because data can be transmitted at a slower rate when too many errors are detected in high-speed operation. This fall-back capability keeps us operational while we wait for the phone company to determine the nature and source of the degradation, to schedule service, and to make actual repairs.

The star network design works well for the 100 terminals used for the prototype catalog, but is insufficient for more than the most modest growth. If we were to add a large number of terminals to the star network, we would be faced with two undesirable options: we could use the full capacity of the multiplexers by increasing the number of terminals per high-speed telephone line (thereby either losing the margin for error recovery or maintaining it by slowing the data rate of each terminal); or we could install more telephone lines and modem-multiplexer sets. Even if the first option were selected, the terminal population would soon exceed the finite capacity of the equipment (about 180–200 terminals), and more lines would be needed. Economics quickly becomes the limiting factor.

The simplicity of the star design makes it easy to predict growth costs, and those costs are high. A telephone line between DLA and UC San Diego costs between $600 and $700 per month, with rates escalating rapidly. A pair of modems and multiplexers costs about $23,000, or about $1,500 per terminal. These costs include little provision for recovery from failures other than line degradation (as opposed to outright failure). By discarding all provision for recovery for line degradation, we could lower these costs, as discussed above. The resulting network would be neither reliable nor affordable.

TRANSITION TO THE PRODUCTION NETWORK

A practical 1,000-terminal network requires a more complex architecture: one that has numerous alternate data paths carried over a variety of media and that incorporates sophisticated and automated error recovery and reconfiguration capabilities. In building the production network, we are emphasizing reliability, operability, extensibility, and control of costs through economy of scale.

Automation of the network management function, wherever possible, is extremely important. In our 100-terminal prototype network, virtually all reconfiguration (modem speed changes, switching out inoperative lines where alternate paths existed, etc.) was done manually. The need for human intervention greatly slowed our ability to respond to problems. Someone had to recognize a problem existed, isolate it, and contact a DLA telecommunications staff member to initiate corrective action. In the 1,000-terminal production network, the staff size to carry out such manual procedures would be impractical; moreover, statistics suggest that in a network that size, *something* is almost always wrong, which would create a perpetual stream of interruptions for the telecommunications staff.

This new emphasis on operational aspects requires changes at every level of the network: the data paths, the hardware, and the protocols. The media must be broadband, reliable, and cost-effective; as we shall discuss, both satellite and terrestrial microwave are attractive for the long-distance links. For distribution at a given campus, we expect to use such techniques as radio and very high frequency (21-22 gigahertz) microwave, in addition to traditional local telephone loops.

Given our requirements for reliability and extensibility, the only approach is to implement packet switching. In a packet-switching network, data is broken into packets, inserted into the network, and routed to its destination by the network control hardware and software, guided by addresses appended at the point where the data enter the network (much like placing a letter into an envelope, writing the address and return address on the outside, dropping it in a mailbox, and letting the postal system deliver it). This packetizing and routing process allows us to take advantage of a complex multipath mesh network; the character-by-character asynchronous transmission we employed in the prototype, in contrast, is limited to point-to-point connections.

A packet-switched network requires hardware, software, and protocols to provide pathfinding, routing, and congestion control, to correct transmission errors, and to detect and avoid inoperative paths. In a sense, the protocols serve as the specifications for the internal operation of the network and the external interfaces to it, providing a definition of the functions to be accomplished by the hardware and software components of the network.

PROTOCOL ISSUES

A protocol is a set of specifications for standardized data formats, transmission sequences, and responses, enabling communication among a group of sites. Protocols were developed originally to describe and formalize packet-switching rules, but in recent years the term *protocol* has taken on a broader sense; it describes more general communications activities, such as computer channel interfaces and videotex transmission. (Cohen and Yemini give a nice presentation of the general concept of a protocol.)[1] There are a number of packet-switching protocols in use today, including X.25, originally developed in Europe by CCITT and widely used there,[2] Systems Network Architecture (SNA), developed by IBM and heavily used for networking IBM (and some non-IBM) equipment,[3] and TCP/IP, developed by Bolt Beranek and Newman (BBN) and the Defense Advanced Research Projects Agency (DARPA) for the U.S. Department of Defense ARPANET.[4-6] Xerox ETHERNET is an example of a local-area packet protocol designed to allow equipment that is located within a limited area to communicate at high speed.[7] A number of other computer vendors also offer their own protocols.

In the last two years there has been a flurry of standards activity in the data communications community, and much euphoric discussion of the ISO model for Open Systems Interconnection.[8,9] Unfortunately, much of this euphoria is ill-founded. The ISO model is just that: a conceptual model. (Even as a conceptual model, it has some problems when applied to anything other than the X.25 protocols it was designed to model. Consider, for example, fitting the higher-level PUP protocol[10] for Xerox ETHERNET into the ISO model. See Nessett for another example of these problems.)[11] From the point of view of actually interconnecting local and intercampus networks to fully exploit our catalog, the Open Systems Interconnection model merely formalizes a problem.

Our solution—or more accurately, our approach, which we believe will provide a framework within which specific solutions can be developed—is to exploit the concept of the internet that has evolved over the last few years.[12,13] The internet concept assumes that there will be many incompatible networks and that what is needed is a means of concatenating them through gateways. In this sense it mirrors the administrative and technical nature of telecommunications at UC.

DLA will use the U.S. Department of Defense TCP/IP protocols for its backbone network. We believe that these protocols will facilitate interconnections with the many other networks we expect to see in the UC system in coming years. They provide a flexibility of interconnection that we could not hope to find with the SNA or X.25 protocols.

There are major philosophical and conceptual differences among SNA, X.25, and TCP/IP. SNA is a centralized protocol oriented to full-screen terminal or computer-to-computer interactions. It is primarily for IBM equipment. Routing in SNA is established at session connection. This reduces a network's ability to adapt quickly to changes in traffic patterns and to make optimal use of bandwidth, and limits the ability to recover from failures in network links. The software interfaces to SNA are extremely complex and are still under active development by IBM.

For implementing a network, X.25 does not offer a complete set of protocols; it must be supplemented by higher-level protocols, which are not yet widely accepted in the U.S. X.25 is also limited in that it is strictly a virtual circuit protocol designed with large public data networks in mind.

TCP/IP is really two protocols: Transmission Control Protocol and Internet Protocol. TCP is a reliable end-to-end protocol that goes beyond X.25 to address some considerations dealt with in the higher-level CCITT protocols. However, it is the Internet Protocol that really distinguishes TCP/IP from X.25 or SNA. While SNA and X.25 are concerned primarily with a single network,* IP deals specifically with the problems of network interconnection. Internet addresses include both a network specification and a destination address within the target network. In addition, TCP/IP is a very robust protocol, designed to operate reliably over potentially unreliable transmission paths.

*CCITT does offer X.75 in conjunction with X.25, but this protocol only addresses the problem of linking relatively homogeneous X.25 nets through concatenated virtual circuits. This approach is considerably more limited than IP.[14]

Without a fully realized internet concept, there are problems in interfacing multiple networks even when they operate under the same protocol—particularly when the networks are managed by different organizations. Addressing and routing are major problems. At present, protocols other than TCP/IP typically do not deal with the interconnection of heterogeneous networks at all, except through a limited technique called encapsulation. The idea of encapsulation is to allow information to be passed between two homogeneous networks across a network that uses an incompatible protocol by sending the entire packet as data across the intermediate network. However, encapsulation does not provide true internet addressing and routing, and it is subject to packet size and overhead problems.[15]

Encapsulation is not an answer to DLA's requirements. Long-term, we expect the catalog to interface with many different networks: various campus local networks provided by a number of vendors (ETHERNETs and others), general campus networks (based perhaps on X.25), and special-purpose networks (such as videotex). We must be able to address the terminals in all of these networks and carry out the requisite protocol conversions for each of them.

HARDWARE: PACKET SWITCHES, TERMINAL INTERFACES, AND GATEWAYS

The network is made up of nodes linked to each other through various kinds of dedicated data paths. Each node is a small computer called an Interface Message Processor (IMP) that functions as a packet switch, routing and managing message traffic. Each IMP also monitors both the transmission quality of the data paths connected to it and the status of its neighboring nodes.[16,17] Attached to each node can be one or more "host" computers that allow terminals to connect with the network. Terminals connected to the local host can access resources there, or open connections from the local host to remote hosts in the network via the local host's IMP connection. The key point here is that a program on the local host is connecting to a service program on the remote host across the network, on behalf of the user. The connection is between computer programs on the two machines, with the program on the user's local host *mediating* on his or her behalf. The "rules" for the interchange between the two computer programs are set out by the TCP/IP protocols plus a higher-level protocol called TELNET,[18] which deals with remote terminal service.

There is no reason for such computer-to-computer links to be limited to supporting service to a terminal. Other higher-level protocols are provided for activities such as file transfer between two hosts, for example. With this approach, we have gained the capability to link other machines for the distributed processing of bibliographic data. This will be important in future applications. In fact, the network may even be able to support some nonlibrary applications with little or no impact on its primary function.

There is also a special-purpose host, called the Terminal Access Controller (TAC), which allows terminals without a local host computer to access remote hosts. Basically, the TAC provides TCP/IP and TELNET on behalf of each terminal just as a full-capability host computer would, but offers no local processing capabilities. Initially, these TACs will provide network access for the dedicated online catalog terminals in the campus libraries. Auto-answer modems can be attached to a few ports to provide local dial-up capability. Other ports could be connected to a campus port selector, allowing campus terminals connected to this port selector to access the catalog.

A gateway is a computer that provides a link between two different networks.[19,20] Ultimately, the catalog will gateway to a number of other computers. Any campus machine that supports TCP/IP and TELNET will be able to access the catalog on behalf of any terminal that can connect to that local computer through whatever means are in use locally—dial-up, direct connection, or some form of local network. (One implication of this is that any terminal on one local TCP/IP–capable host could also connect to any other such campus host on our network.) Additionally, local networks could be linked directly to the IMP through small, dedicated gateway computers. In some cases, local network vendors have already announced their intention to provide such machines; in other cases, they can be obtained from third parties or developed specially for the purpose.

The implementation and operation of gateways and internetwork connections are active research areas. There are an enormous number of possible topologies that can occur, particularly when portions of the networks involved fail. Sunshine[21] and Davies and Bates[22] provide an excellent overview of some of the possible problems, which go far beyond the scope of this paper. One important point, however, should be made: as well as distributing access to services, internetworking can actually increase the robustness of a network. Multiply interconnected networks, with properly designed internet routing, actually function as additional data paths. Figure 3 gives the flavor of this effect. In this example, we have three linked networks. Suppose that within network *A* nodes *X* and *Y* have only a single link *L*. (There are no other paths that can go from *X* to *Y*.) If link *L* fails, network *A* is partitioned into two subnetworks, one containing node *X* and the other node *Y*. In this case, the two subnetworks can continue to communicate, albeit inefficiently, by passing through the internet gateways to networks *B* and *C*, which in turn are internetworked.

Bolt Beranek and Newman, the original developer and manufacturer of the ARPANET technology, is helping DLA to apply this technology to its network.[23] A few modifications are needed, but most are relatively easy to implement. The purpose of these changes will be to bypass the TAC host selection and sign-on sequences for catalog terminals in order to simplify the user interface and to provide load leveling and automatic reconnection to another host when the currently active host fails. In addition, software will be redistributed. In the ARPANET, one commonly finds a minicomputer acting as a bridge between the network and large local hosts, which reduces the local software development necessary to connect the host to the network. Although we may use this as a short-term

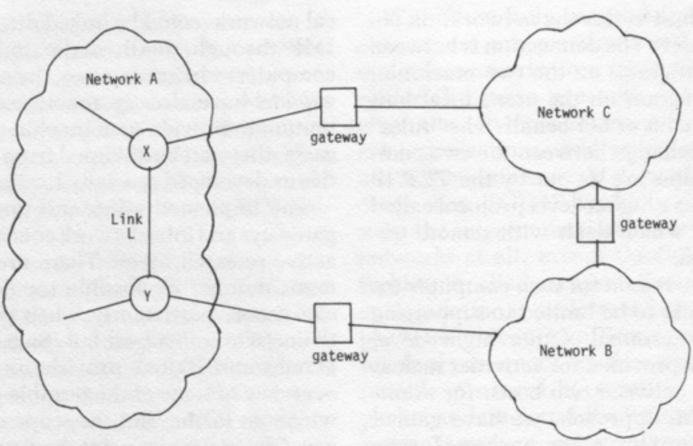

Fig. 3. Example of Internet Routing Configuration.

expedient, DLA's long-term plan calls for full implementation of TCP/IP directly on the host computers. This will maximize reliability, throughput, and control, and keep costs to a minimum.

To facilitate management of the network, a Network Control Center (NCC) will be installed at DLA. This will consist of a BBN C/70 minicomputer that will receive and track statistics and status reports from IMPs throughout the network, notifying the operations staff of events such as line outages or node failures. Thus we will not only circumvent problems automatically where possible, but will have automated problem identification and tracking capabilities. The NCC machine also hosts a variety of diagnostic programs that can verify repair work or perform detailed problem diagnosis. For example, when the telephone company reports that a line has been repaired, we will be able to use these utilities to verify correct operation before signing off on the repair call.

THE UNDERLYING DATA PATHS

The physical data paths, or transmission facilities, must be considered along with changes in load, architecture, and protocol. Neither the reliability nor the capacity of standard 9.6 KB telephone lines suffice for the production network.

We can no longer assume that conventional telephone lines are the solution to the needs of data networking facilities. Deregulation of telecommunications and telephone companies is changing the nature and costs of both local and long-distance facilities, especially in data applications. Bell Telephone is now marketing its Dataphone Digital Service (DDS) in preference to conventional analog data service, but there is at present no similar digital service in non-Bell areas (Los Angeles and Santa Barbara). Although the conventional facilities offered by telephone companies are usually not too expensive to install, they are among the most expensive in ongoing use costs. Data network companies offer services that are relatively inexpensive both to install and operate on a small scale, but most of them are limited in protocol and capacity. At our expected traffic volume and terminal count, the operating cost soon would become prohibitive. For a low operating cost, we must follow the example of many private corporations with geographically dispersed operations—the University (or DLA) must develop its own communications facilities.

The first and most obvious replacement for wire is terrestrial microwave. This transmission medium, though often set apart from other media, has been used rather transparently for many years. Television stations use it to transmit programming from their studios to their remote transmission towers. Telephone companies use it to carry phone conversations where no telephone wires could possibly be strung. We would not be unique in using microwave to solve our communications problem, but we would be near the leading edge in the current application of this technology.

As a trial application, we decided to link the Berkeley and Santa Barbara campuses by microwave. In cooperation with another Systemwide Administration office, DLA obtained access facilities and channels from the Public Safety Microwave System of the California State Department of General Services. During the trial, operating costs were about half of the usual charges from the telephone companies (UC Santa Barbara is in a non-Bell service area). At this time, DLA is using a microwave circuit with capacity (9600 baud) equivalent to the telephone companies' "voice grade" circuits. The quality of its data transmission is equal to that of the phone lines, at a lower cost per mile. The higher-speed (56,000 baud) microwave circuits used by the other UC office may not even have been deliverable by the phone companies.

Although the average operability of our microwave circuit is roughly equivalent to a phone line, the problems that are occurring are entirely different. The transmission quality of the telephone lines is continuously drifting within operable margins, which means that the equipment must perform retransmissions to transport error-free data, but our microwave circuit operates with a near-perfect data quality, interrupted by total dropouts. For any measured period of time, the amount of good data passed will be greater on the microwave circuit (with both the microwave and telephone circuits operating at the same speed), but the microwave circuit's occasional need to switch to an alternate path eliminates much of this advantage. In addition, DLA has a lower priority than other users because the network is maintained primarily for public-safety agencies, and thus repairs on DLA circuits may have to wait.

A second alternative would be to develop a microwave network just for the university, either jointly with or solely by the state General Services Department. Because it would be an entirely new system, UC would have the option of choosing either an analog system, such as the one we now use, or a less-complex digital system, which would allow easier data integration. A private microwave system would be costly to install, however, because capital construction costs are high for placing microwave repeater stations in remote locations. Microwave requires line-of-sight transmission, and mountaintops are ideal for these stations because the distance between stations can be maximized. However, reliability in a microwave network is costly and hard to obtain, since it calls for redundant chains of repeaters between the two terminal points of a connection.

The third choice is a more esoteric application of microwave technology: a satellite-based network.[24] Satellite communication is nothing more, after all, than a microwave network with a single repeater station located about 22,300 miles above the earth in geosynchronous orbit, rather than on top of a mountain. Since the path is directed upwards, rather than parallel to the earth's surface, the radio beam is less affected by atmospheric conditions such as inversion layers. At microwave frequencies, signal attenuation by the ionosphere is minimal, although at some of the higher frequencies starting to come into use, extreme weather conditions can cause some signal scatter problems.

Although the cost per earth station is higher than most microwave installations, there is no need for the large number of repeaters required by terrestrial microwave systems, so more funds can be directed toward redundancy to avoid outages due to component failure. Earth stations (satellite dishes) can be placed on the roof of the building containing the computer, or in a nearby parking lot, rather than in a poorly accessible remote location. This means lower construction and maintenance costs. Since UC would not own the satellite supporting the DLA application, or even the transponder, there would be an ongoing cost for our portion of the satellite's transmission bandwidth. (Because of the enormous capacity of even a single satellite transponder, it is not cost-effective to acquire dedicated facilities even at the transponder level.) The cost, however, would not be as high as that charged by a telephone company for the same bandwidth.

Satellite communications offer significant dividends in reliability. The space segment itself is highly reliable, and in the event of transponder failure on board the satellite, a spare transponder could be switched rapidly into operation on command from the ground-control facility. We know of no unexpected total failure in an operating communications satellite.

The broadcast nature of satellite communications has intriguing implications. Each ground station can listen to all traffic passing from the satellite to any other ground station in the system. By providing an inexpensive link between each campus and its closest neighbor over a medium other than satellite, we would have full alternate-path capability: if one earth station failed, data could be routed over the terrestrial link to the neighboring ground station. To fully exploit this potential, we probably will have to provide tighter integration between the satellite ground stations and the packet-switches; given the current state of the art, explicit commands would be required to reroute the data paths after failure of a ground station. Current vendors of satellite systems (transponder capacity and earth stations) such as VITALINK present the satellite links to the user as point-to-point links. Although the end point of a link can be moved from one earth station to another through the satellite network control center, we would want to accomplish this rerouting automatically from our packet network control center, when the IMPs connected to the satellite link detected a protracted link failure. This would require the development of specialized software to link the two network control facilities. An alternative approach would be to provide permanent redundant satellite links: although no extra earth station equipment would be required, the extra links would still tie up satellite transponder capacity, since the bandwidth would have to be reserved.

There is promise of even more effective use of satellites in coming years, as experimental technology in the BBN/DARPA SATNET project is refined and becomes commercially available. This involves a tight coupling of satellite data communications and packet-switches, making essential use of the broadcast nature of the media.[25,26] Satellite transponder bandwidth is allocated dynamically through a reservation scheme. Ultimately, approaches based on this technology, perhaps in conjunction with a network of packet satellites (launched into non-geosynchronous low earth orbit by the space shuttle), may lead to further drops in satellite bandwidth costs.[27]

Broadcast satellite transmission greatly simplifies disaster recovery for the host computer site: rather than having to reroute an entire network in order to switch to a backup site, we could just install a—perhaps truck-mounted, mobile—satellite dish.

As described, we have several choices for a long-distance backbone network. We also have a number of options for the distribution networks on each campus that supports the terminals. Some campuses have existing, unused coaxial cable connections between buildings that could carry large amounts of data. Other buildings are within line of sight from each other, which makes them suitable for very high frequency microwave links, or perhaps infrared or laser light links, any of which would be less expensive than digging trenches and laying new cable.

We are also investigating short-distance, low-power radio for communicating between a centrally located campus controller and all campus terminals, or, on a smaller scale, between a controller and terminals within a building.[28] If this can reasonably be developed, building renovation for installing terminal data cables could be eliminated. This work is usually costly, always disruptive, and ultimately inflexible. A radio-linked terminal could be moved from place to place with no more difficulty than finding a place to plug in the power cord. Terminals at present are tied to the location of data cables or telephone lines, creating problems when they need to be relocated to accommodate shifting usage patterns or remodeling.

SUMMARY AND PROSPECTUS

A telecommunications network represents an enormous, long-term commitment. Although the effective life span of a computer is typically about five to seven years, the operational life span of a major telecommunications network is usually measured in decades. Life spans of this magnitude are necessary to successfully amortize the huge initial costs. Consequently, it is essential to develop a network that will avoid early obsolescence. This issue is partly addressed through the use of highly extensible, well-defined protocols such as TCP/IP, which enable a network to evolve as its internet environment grows. TCP/IP is largely independent—or more accurately, adaptable—to characteristics such as propagation delay, which appear in switches between terrestrial and satellite transmission media. (This is not always the case; many current local networking protocols are intimately bound to the characteristics of the supporting transmission media. Fortunately, DLA's present concerns are with long-haul telecommunications—local distribution is being treated in a more ad hoc fashion, pending developments on the campuses and further technology development by DLA.)

Orderly, nondisruptive growth is a major consideration in DLA's network. Because of the scaling considerations discussed above, it will be necessary to carry out a massive conversion over a period of eighteen months in our transition from prototype to production network. The growth pattern of the network is interesting in that it appears to grow from the center out, with equipment formerly comprising backbone links being recycled into local use on the periphery of the network. We will be able to run some facilities in parallel, but a certain amount of discontinuity seems inevitable. Moving terminals from prototype network multiplexers to production TACs is a reasonably fast process. Without installing peak-load data paths to some locations specifically to provide for the conversion, however, we may face more serious service interruptions in moving long lines to DLA over from the multiplexers to the IMPs, since this will be done by DLA staff visiting the campuses for installation work in most cases.

The telecommunications network for the online catalog has two logical components: the long-haul backbone and a set of nine distinct local networking and distribution problems. At this point, we believe that the long-haul part of the network, while expensive and complex, can be implemented without overwhelming technical difficulty. Because of budget constraints, delivery problems, or conversion planning, it may be necessary to implement the large network incrementally, and more slowly than necessary, but we are confident that the project will go forward without major surprises. The major technological questions are those of optimization—for example, how do we effectively exploit the broadcast nature of satellite communications for maximum redundancy at lowest cost?

The local networks of the system are another matter altogether. At present, we are engineering special-purpose, custom solutions for each terminal cluster. Because of the enormous amount of detail that must be tracked for each installation, this is a slow, staff-intensive, and costly process. We are trying to develop a series of building blocks to handle local distribution problems, but much of the technology is only now becoming available in the marketplace.

In the long run, avoiding wiring of any kind wherever possible seems to be the key to success in local distribution. This eliminates a whole series of facilities-management problems, and allows local terminals to be reconfigured at will. The capabilities we seem to need have much in common with tactical military data communications systems. In the coming years, we expect local distribution to be the main arena for DLA research and development. We will have to develop a local distribution system that supports gateways to campus networks through the Internet Protocol and that also supports both fixed-terminal clusters and movable packet-radio-based terminals dedicated to the online catalog.

Many questions remain unanswered. If we waited for answers to all of them we would never implement a network. Instead we are trying to develop an extensible, adaptable network that is hospitable to interconnection with others. We have not demanded uniformity in order to simplify the implementation. Given the University of

California telecommunications environment, these choices seem realistic. The result is a network that we believe is capable of evolving to meet the needs of the future—unforeseen as well as planned.

REFERENCES

1. Danny Cohen and Yechiam Yemini, "Protocols for Dating Coordination," *Proceedings of the Fourth Berkeley Conference on Distributed Data Management and Computer Networks, August 28–30, 1979* (Berkeley, Calif.: Lawrence Berkeley Laboratory, Univ. of California), p.179–90.
2. CCITT, "Provisional Recommendation X.25, " *Data Communications Standards* (New York: McGraw-Hill, 1978).
3. IBM Corporation, *Systems Network Architecture Format and Protocol Reference Manual: Architechural Logic* (White Plains, N.Y.: The Corporation, 1978). (Order No. SC30-3112.)
4. *ARPANET Protocol Handbook* (Menlo Park, Calif.: Network Information Center, SRI International, 1982).
5. *INTERNET Transition Workbook* (Menlo Park, Calif.: Network Information Center, SRI International, 1982).
6. Jonathan B. Postel, Carl A. Sunshine, and Danny Cohen, "The ARPA Internet Protocol," *Computer Networks* 5, no.4:261–71 (July 1981).
7. *The ETHERNET: A Local Area Network Data Link Layer and Physical Layer Specifications* (Digital Equipment Corporation, Intel Corporation, Xerox Corporation, 1980).
8. R. desJardins, "Overview and Status of the ISO Reference Model of Open System Interconnection," *Computer Communication Review* 11, no.2:10–14 (April 1981).
9. "Open System Interconnection—Basic Reference Model (ISO Draft Proposal 7498; ISO TC97/SC16)," *Computer Communication Review* 11, no.2:15–65 (April 1981).
10. Xerox Corporation, *INTERNET Transport Protocols* (Stamford, Conn.: The Corporation, 1981).
11. D. M. Nessett, "Hyperchannel Architecture: A Case Study of Some Inadequacies in the ISO-OSI Reference Model," a final report for the U.S. Department of Commerce, National Bureau of Standards, Institute for Computer Sciences and Technology (Livermore, Calif.: Lawrence Livermore Laboratory, Univ. of California, 1981). (NTIS Document #DE 81028497.)
12. Alan Sheltzer, Robert Hinden, and Mike Brescia, "Connecting Different Types of Networks with Gateways," *Data Communications* 12, no.8:111–22 (Aug. 1982).
13. Postel, "ARPA Internet," p.261–71.
14. Jonathan B. Postel, "Internetwork Protocol Approaches," *IEEE Transactions on Communications* 28, no.4:604–11 (April 1980).
15. John Shoch, Danny Cohen, and Edward Taft, "Mutual Encapsulation of Internetwork Protocols," *Computer Networks* 5, no.4:287–300 (July 1981).
16. *Report No. 1822: Specifications for Interconnection of a Host and an IMP* (Cambridge, Mass.: Bolt Beranek and Newman, Inc., 1978).
17. J. M. McQuillan, *Report No. 2831: Adaptive Routing Algorithms for Distributed Computer Networks* (Cambridge, Mass.: Bolt Beranek and Newman, Inc., 1974). (Ph.D. dissertation in applied mathematics, Harvard University; available from National Technical Information Service, AD 781467.)
18. *INTERNET Workbook.*
19. Virginia Strazisar, *How to Build a Gateway*, IEN 109 (Cambridge, Mass.: Bolt Beranek and Newman, Inc., 1979).
20. Sheltzer, "Connecting," p.111–22.
21. C. A. Sunshine, "Addressing Problems in Multi-Network Systems," *Proceedings INFOCOM 82* (Silver Spring, Md.: IEEE Press, 1982), p.12–18.
22. B. H. Davies and A. S. Bates, "Internetworking in the Military Environment," ibid., p.19–31.
23. M. Wingfield and R. Edmiston, *Report No. 4810: Telecommunications Study for the University of California Division of Library Automation* (Cambridge, Mass.: Bolt Beranek and Newman, Inc., 1981).
24. James Martin, *Satellite Communications Systems* (Englewood Cliffs, N.J.: Prentice-Hall, 1978).
25. Robert E. Kahn, "The Introduction of Packet Satellite Communications," *Proceedings of the 1979 National Telecommunications Conference, Washington, D.C.* 8:45.1–45.8.
26. Irwin Mark Jacobs, Richard Binder, and Estil V. Hoversten, "General Purpose Packet Satellite Networks," *Proceedings of the IEEE* 66, no.11:1448–67 (Nov. 1978).
27. Ke Y. Park, "Global Packet Radios Using Low Altitude Satellites," *IEEE 1980 National Telecommunications Conference, Houston* (New York: IEEE, 1980), p.70.5.1–70.5.5.
28. Robert E. Kahn and others, "Advances in Packet Radio Technology," *Proceedings of the IEEE* 66, no.11:1468–96 (Nov. 1978). ■■

Terminals for a User-Friendly Library System

Bennett J. Price

A video display terminal (VDT) is the library user's point of contact with MELVYL, the University of California's online library catalog. MELVYL's user-friendly interface is designed to mask the sophisticated software needed to manipulate complex bibliographic files. The VDT must not hinder the user-friendly interface by interposing forbidding technology and hardware. Consequently, the Division of Library Automation (DLA) has spent and will continue to spend much time and attention in obtaining user-friendly terminals.

This article explains the criteria that DLA uses to select terminals, the modifications we make to them, and, finally, our experience to date with terminals.

SELECTION CRITERIA

To begin with, DLA defined three sets of selection criteria: (1) technical and electronic, (2) functional and ergonomic, and (3) maintainability and vendor support.

Technical and Electronic Criteria

The technical requirements for MELVYL's terminals are simple. DLA will eventually supply 600 VDTs to a system designed to support 1,600 terminals. Many of the additional 1,000 terminals are already on the nine campuses in departmental offices, computer centers, laboratories, typing rooms, and other such locations. MELVYL is designed to be compatible with these VDTs, which are largely glass Teletypes like the Lear Siegler ADM-3. A "glass Teletype" is the "TV" version of a terminal that can send and receive the ASCII character set but not much else. Notably absent in such a "dumb" terminal are features like insert, delete line, reverse video, backspace, and cursor control keys.

In technical terms, DLA's terminals must support half-duplex, asynchronous, character-by-character transmission at 300 or 1200 baud mark parity through an RS-232C interface. The screen must be capable of displaying twenty-four lines of eighty upper- and lowercase characters (the ninety-six-character ASCII set). The technical requirements for MELVYL's terminals are easily met; in fact it is more difficult to find terminals that do not meet them than it is to find terminals that do. Well-known terminals that will not work with MELVYL are the IBM 3270, the OCLC-Beehive 105, and the RLIN-Zentec ZMS-90.

Functional and Ergonomic Criteria

DLA looks for terminals that are easy to use. Many have been investigated and tried. Our minimum requirements are as follows:

1. The twelve-inch display must be sharp, both in center screen and at the corners, bright, and without flicker.

2. The picture tube must have some sort of antiglare and contrast-enhancement treatment. Antiglare screens diffuse reflected light; they make it difficult or impossible to see reflections of fluorescent lights, bright windows, or yourself on the screen. Contrast enhancement gives the unlit screen a dark gray or black appearance and makes the display more legible in bright light.

3. Characters must be well formed; numerals *1* and *0* must be distinguishable from letters *l* and capital *O*. The descenders on the letters *p*, *g*, *y*, and *q* must drop below the line. This requires more than a five-by-seven dot matrix.

4. The keyboard must be integral to the terminal. Detachable keyboards (like those on the OCLC and RLIN terminals) are usually considered desirable, but MELVYL terminals go into areas that often are unsupervised, in comparison to the controlled atmosphere of an office. They are placed on tables that might be loaded down with books, jackets, raincoats, umbrellas, or lunch bags. An integral keyboard has a much better chance of survival under these

---

Bennett J. Price is a senior systems analyst, Division of Library Automation, University of California.

conditions. (In practice, an integral keyboard usually means that the screen cannot be tilted or rotated to suit the user's gaze.) The detachable keyboard is desirable for someone who must work at a terminal for hours, but to us the persuasive factors were that catalog users normally would not spend hours at the terminal and that the keyboard could not get knocked off the table. Of course, a separable keyboard can be bolted down, but this increases the installation cost, may ruin furniture, and makes repairs difficult. Some vendors do offer metal plates to attach separable keyboards to the rest of the terminal.

5. Of all the keys on VDTs that are not alphanumeric, DLA uses only four—shift, backspace, break, and return. In order that the VDT itself be easy to use without special instructions, no terminal is acceptable if it requires two fingers to activate any of the latter three functions. Thus, VDTs must have a backspace key—"CONTROL H" is not sufficient. Similarly, those terminals that require a shift or control key to be pressed simultaneously with the break key to initiate a break are excluded. If the keycaps for these function buttons are of a color different from those on the alphanumeric keys, it is considered a bonus.

6. DLA has not yet determined how best to provide printing, whether of brief citations or full bibliographies. A printer port, however, is required just in case it is later decided to attach printers directly to the terminals, rather than locating all printers remotely. No demands are made on the type of port (pass-through versus buffered, for example); the VDT simply must have some print facility.

7. On most VDTs the switches used to set baud rate, parity, and the like are in the back and relatively inaccessible. Any terminal that has these switches in a highly accessible position is disqualified unless access to them can be prevented easily.

The following features are considered desirable but not essential.

Nonthreatening Appearance

Some VDTs are considerably less imposing in appearance than others. Their color, shape, size, or lines belie their technology. We do not want to scare patrons away from MELVYL before they even sit down. Some terminals have fewer keys than others; this too is considered desirable. However, a separate numeric keypad seems more a benefit than a deficiency inasmuch as we plan to offer call-number or class-number searching.

Case and Keyboard Design

DLA looks for VDTs whose cases do not have flat surfaces (on which to place and then knock over coffee cups). Similarly, the technology of certain keyboard designs makes them much less vulnerable to spilled liquids than others.

Scroll/Soft Scroll

Scrolling occurs when a twenty-fifth line is written to the bottom of a filled screen capable of holding only twenty-four lines; the top line on the screen then pops up or scrolls off the screen. "Soft scrolling," a nice but nonessential feature, allows the display to smoothly "float" up rather than pop up each time a new line is written on the bottom of a screen.

Quietness

Fans make terminals somewhat noisy; fans with sleeve bearings (instead of ball bearings) often become quite noisy after a few years of operation. All of the terminals we have investigated are convection cooled, not fan cooled. Some VDTs are like typewriters when the end of a line is approached—they emit the high-technology version of a bell—a beep. Terminals that do this are down-rated as are those that beep each time the break key is depressed. Terminals that simulate the "clacking" of typewriter keys with electronic key "clicks" are similarly down-rated, unless the "clicker" can be turned off. Disconnecting the beeper or speaker inside the terminal is not a good solution for us because on occasion MELVYL does use beeps to accompany a few important messages.

Screen Phosphor

The screen phosphor is a material applied to the inside face of the picture tube (the cathode ray tube or CRT). When electrons strike the phosphor, it glows. Different phosphors glow with different colors. Color TVs have red, green, and blue phosphors, black-and-white TVs have a white phosphor. VDTs are commonly offered with white or green phosphors, and some manufacturers offer yellow or amber phosphors as well. We give no preference to any phosphor color. Another phosphor characteristic that is considered important, however, is persistence: that is, how long the phosphor continues to glow after the electron beam has been turned off. When used with scrolling displays, long-persistence phosphors are objectionable because new lines are being written atop lines that have not yet disappeared. Although MELVYL usually clears a screen completely and then writes from the top down, some scrolling is used, particularly in MELVYL'S LOOKUP mode (the user interface for beginning users).

A feature sought but not found in the low-cost VDTs we have investigated so far is an automatic display dimmer. Such a dimmer will cut the screen intensity in half or turn the screen off if the terminal is not used for five minutes or so. A display that remains unchanged on the screen for long periods of time may eventually burn a pattern in the phosphor. On some VDTs this can be prevented by programming those screens that appear for hours on end to display in half-intensity.

Visual Attributes

Increasingly, inexpensive dumb terminals are being offered with a variety of visual attributes. For example, characters, lines, or screens can appear in reverse video (black characters on a light background), with half intensity, or with underlining or blinking. Because signing on and secret passwords are generally not needed for MELVYL, we have little need of the blanking (i.e., no display) attribute. Many terminals have all of these attributes. Some VDTs allow the host computer to control the blink rate or even to change the size of the characters displayed. In the latter case an important message might be twice or quadruple the size of regular messages. All of these visual attributes are considered desirable provided we do not have to pay extra for them. We have not implemented any of them to date.

Keyboard Features

1. A "typeamatic" keyboard is one that automatically repeats a keystroke when a key is depressed for more than a second or so; this feature is commonly found on electric typewriters for the space bar, backspace, and underscore keys. On VDTs the "typeamatic" function typically is found on every alphanumeric key. DLA down-rates terminals on which the "typeamatic" feature cannot be turned off. We think the consequence of a heavy finger should be only one character.

2. Keyboard bounce—the multiple entry of a character when only one character is desired—is due to poor keyboard design (or the failure of a good keyboard). Few of the terminals we have investigated show this defect.

3. "N-key rollover" specifies what happens when a second (or third or *N*th) key is depressed while one or more other keys are already depressed. Terminals lacking rollover capability lock out subsequent depressions. For example, in typing the word *cat* on a terminal with two-key rollover, if the *a* is pressed before *c* is released, *ca* will appear on the screen. Without the N-key rollover function, only the *c* would appear on the screen. Poor typists with heavy fingers do better with keyboards that lock out rollovers; most typists with any skill at all seem to find the rollover feature desirable, provided the keyboard's action is firm enough to somewhat hinder accidental depressions. We prefer multikey rollover VDTs on the assumption that most users can type a bit and that many can type well. (For similar reasons, we accept the "qwertyuiop" keyboard rather than the "abcdef . . ." keyboard.)

4. Entry rate is the speed with which characters can be correctly entered. It is in part determined by the shape and location of the keys and keyboard and the mechanical motions built into the key switches (length of stroke, pressure, end-of-stroke cushion, confirmation of closure, etc.). In some terminals, however, the electronic

design of the keyboard will cause false entries when the entry speed is high. High entry speed is common in skilled typists, but even unskilled typists will often generate certain letter combinations very quickly. We have found that most terminals can keep up, but some are quite poor in this respect.

5. Keyboard lock (or keyboard disable) is a feature that permits the host computer to enable and disable the keyboard (the keys do not actually become locked so that they cannot be depressed). Patrons sometimes begin typing commands before the computer has signaled its readiness to receive them with its prompting character. A VDT whose keyboard can be disabled helps prevent these premature responses. (In a half-duplex system like MELVYL, every keystroke appears on the screen, unless the keyboard is disabled, regardless of whether or not the computer receives it.)

Protected Fields

These are helpful for full-screen fill-in-the-blanks data entry (as in RLIN). Protected fields prevent the operator from changing the information within the protected field (such as the mnemonics of the 008 field) and define the spots that the tab and back-tab keys skip over. Protected fields are often assigned a visual attribute different than the rest of the screen. MELVYL might be able to use protected fields in the future to prevent backspacing beyond the prompt arrow, but to date no use has been made of this capability. Two of the three brands of terminals purchased for the MELVYL prototype provide protected fields.

Cursor Addressing

This feature has two aspects. First, cursor "writing" allows the host computer to place the cursor anywhere on the screen; after the cursor is so placed, the host may then begin writing on the screen, or control can be turned back to the operator. The other aspect of cursor addressing is cursor "reading" or "locating"; this allows the host to "ask" the VDT where its cursor is located (and often, what characters precede the cursor). Some terminals can do cursor write but not cursor read. MELVYL's current screen formats do not require cursor addressing, but we consider this a desirable feature (as long as we do not pay extra for it).

Cursor Appearance

A solid nonblinking block is considered the most desirable as it is easier to see than an underscore and less annoying than a blinking block. A blinking cursor is often misinterpreted to mean that the system is responding or acting upon a request; the solid cursor is not so interpreted. If a VDT has a blinking cursor, the blinking should stop when the keyboard is used; otherwise the cursor is difficult to locate, particularly when using "typeamatic" spacing or backspacing. Desirable, but not found, so far, is a backspace key that moves a "destructive" cursor, i.e., one that erases what is on the screen while moving back, rather than simply moving back. Patrons are sometimes confused with the operation of the backspace because it does not erase what is on the screen, though it does erase what has been sent to the computer. Thus *united* when corrected to *union* appears on the screen as *uniond*, but the computer reads *union*.

Programmable Keys

Unfortunately, none of the inexpensive terminals we have investigated so far have keys that can be programmed to hold strings such as START, END, or FIND. (Some inexpensive VDTs do come with function keys that send out a string like SOH @ CR with one keystroke; in a full-duplex system these special codes might be defined as START or END and the appropriate alphanumeric string could be written to the screen by the host.)

Still other features, important for a fill-in-the-blanks approach but unimportant for MELVYL, are forward and reverse tabs and tabs that can be set by the user to physically appear anywhere on the screen. The desirability of a reverse tab key, for example, must be weighed against the desire to have as few keys as possible on the keyboard.

Only one manufacturer offers a 120-volt

AC convenience receptacle and an illuminated on/off switch; both very useful features.

Maintainability/Vendor Support

An easy-to-use terminal helps no one if it does not work. DLA has devoted considerable attention to finding terminals that are both reliable and easy to repair. An absolute requirement is that the VDT have local (i.e., both northern and southern California) service facilities. DLA does not require these facilities to be elaborate. If there is no local service station, the vendor must agree to stock a spare terminal or logic board during the warranty period that can be swapped for a defective unit (warranties run from three months to one year). In addition, we look for terminals with major components that are easily accessible and removable. The opinions and reactions of many owners and servicers of various brands of terminals have been collected, as have vendor figures for mean time between failures and mean time to repair.

Self-test features are valuable aids to troubleshooting, particularly in determining whether a terminal's failure is due to the terminal or due to the communications link. Two of the three brands of terminals that DLA purchased for the prototype have self-test features; easy access to this feature is not possible on one of these two, however, because we have deactivated the escape and control keys necessary to initiate the test.

MODIFICATIONS TO THE TERMINALS

In order to make the VDTs easy to use, DLA makes minor modifications to them. All keys that our users never need are deactivated, and their keycaps either are replaced with blank keycaps or covered with sheets of plastic. Backspace keys that are identified by a left-pointing arrow have the keycap replaced with a cap reading *backspace*. The numeric portion of the keycap for the embedded numeric keypad of one manufacturer was quite distracting; these keycaps have been replaced with caps similar in appearance to the rest of the keyboard. (An embedded numeric keypad is one shared with the keys activated by the first three fingers of the right hand. When a "number" or "numeric" button is depressed the regular keyboard is altered so that the right hand controls not a typewriter but a ten-key to sixteen-key adding-machine keyboard.)

In addition, where easily possible and appropriate, keyclicks, blinking cursors, and "typeamatic" keyboards are disabled. Set-and-forget switches that are too easily accessible are locked shut and "patron-proofed."

A small plastic sign is glued above the keyboard of all terminals. It reads, in small letters, "U.C. Prototype On-Line Catalog," and then, in much larger letters, "To Use—Type START and Press RETURN." *start* is the only "log on" required for terminals permanently wired to MELVYL; *start* ends the previous session, resets the default parameters, and brings up the salutation screen. The cost of these modifications is minor, ranging from nothing to $50 per terminal. Rather than modifying the terminals ourselves, we have them modified by the vendors or manufacturers; this saves us from having to store, uncrate, recrate, and reship dozens of large boxes.

INSTALLATION OF TERMINALS

DLA and the campus libraries share the costs and responsibility for installing and repairing terminals. For the prototype, DLA offered each campus ten terminals for public locations. Eight of them were wired directly to MELVYL and operate at 1200 baud (120 characters per second, or cps). Two 300 baud (30 cps) dial-up modems were supplied for the other two terminals. DLA paid for all equipment and telecommunication charges, except for those charges related to dial-up activity. DLA also pays for all repairs on terminals and telecommunications equipment, except for on-site terminal repairs. The campuses and DLA agreed to split fifty-fifty the cost of an eleventh terminal, a spare to be used should one of the first ten be in repair.

Local wiring, whether for electrical sockets or for data cables from terminals to telecommunications equipment, is each campus' responsibility, as are the costs of furniture and any security devices.

For the prototype, DLA recommended three terminals to the campuses, the Micro-Term ACT-5A, the Perkin-Elmer Bantam 550B, and the Televideo 912C. In late 1980, these cost $650–$700 each at the 100+ price available to the university. We pointed out the pluses and minuses of each brand and urged each campus to select the brand or brands that could get local-campus service support most easily. Three campuses selected Micro-Term terminals, two selected Perkin-Elmer, and one selected Televideo. Two campuses selected half Micro-Term and half Perkin-Elmer, and one campus is using its own terminals. In sum, in the first quarter of 1981, eighty-five modified terminals were delivered to the eight campuses initially participating in the prototype test of MELVYL.

Most of the terminals went into public areas in a campus' main library, usually near its main catalog. Other popular locations in the main libraries were the reference room, the documents department, and in or near ILL/ILB. A few were placed in open stacks, and many more in public areas of branch libraries. Campuses with few branches tended to put all eight "hard-wired" terminals in the main library, and campuses with many branch libraries tended to distribute "hard-wired" VDTs more widely. Rarely did a branch library have more than one terminal. One campus purchased seven additional terminals for its undergraduate library. In a few locations, campuses have provided printers for patron use.

Most terminals were placed on tables or surfaces twenty-seven to thirty inches high; that is, on surfaces suitable for use while seated. Some VDTs were placed on surfaces at waist height, suited to stand-up use and its concomitant faster turnover, but it appears that patrons will go to considerable lengths to drag high stools to the terminal rather than stand. It may be that software will be necessary to keep certain terminals for "express" or "one-search-only" use.

In the first four months of 1981, DLA visited seven of the campuses to connect all of the equipment for the prototype; at the eighth campus one enterprising staff member did it herself with some advice from us over the phone. We found that where campus computer centers had prepared the data cables, all was well. Electricians, however, did not do so well, and some of their work had to be redone because DLA's instructions to the campuses were insufficiently detailed for someone not familiar with data processing equipment.

TERMINAL PERFORMANCE AND REPAIRS

The prototype, running since August 1981 in public areas, has helped us learn much about service and reliability. DLA prepared a checkout guide for each terminal and furnished a simple diagnostic tool and spare fuses for each terminal location.

Reliability of the terminals DLA recommended has not been as good as we had hoped. In the first year of operation, about one-third of the units failed; fortunately these were largely still under warranty. After one year, the Perkin-Elmer 550B is proving somewhat more reliable than the other two brands (unfortunately, Perkin-Elmer has discontinued this model). From a pool of eighty-five terminals over a year old, many of which are rarely turned off, we are experiencing about one failure every two to three weeks.

It has been equally discouraging to find that factory or third-party depot service is satisfactory only if a large pool of spares is readily available. Repairs typically take one to two hours (at $60 to $75 per hour), but turnaround time is seldom shorter than a month, even when the terminal is hand delivered to the service depot by the library. To keep ten terminals working under these conditions thus requires two or three spares, not one. At some campuses, the local computer center cannot efficiently repair any of the three brands DLA purchased for the prototype, or will repair only the terminals it owns.

Fortunately, DLA has been able to arrange for depot repair with the Digital Systems Group (DSG) of the Crocker Nuclear Laboratory on the Davis campus. DSG is providing the university with turnaround time of one or two days at $30 an hour. They have been repairing the Micro-Terms since their warranty ran out in January of 1982, and in August 1982 began to do our Televideo repairs. We expect that the Perkin-Elmers will be similarly protected. We are fortunate, therefore, to be free of manufacturer or third-party service, an advantage unavailable to many institutions.

Because DLA does not provide funds to the campuses for on-site service, each campus must bear the burden of arranging for VDT repair. On those campuses with a library repair and maintenance unit, or with staff interested in troubleshooting, repairs seem to go fairly smoothly. Spare terminals are quickly put into place while the defective terminal or its logic board is sent out for repair. On campuses that do not have maintenance units, repairs appear to cause more problems. This is an area that will need improvement when the number of terminals on a campus is increased from ten to dozens.

CONCLUSION

Since we purchased our first eighty-five terminals in December 1981, our investigation and evaluation has continued. We have concluded that obtaining VDTs from multiple sources is far too difficult and unproductive for the marginal benefits it may gain one campus or another. Provided that depot repairs can be done quickly, and that sufficient spares can be kept at the campus, there is little advantage to simultaneous purchasing of VDTs from multiple sources. In mid-1982 we purchased 125 Televideo 910+ 's at $515 each. After taxes, shipping, and vendor modifications, the delivered cost will be about $575 per unit—about $125 less than the least expensive of the three models we purchased eighteen months earlier. These Televideos will be modified in the same ways as the terminals already purchased: their unneeded keys will be deactivated and replaced with blank keycaps and a small sign ("To Use—Type START & Press RETURN") will be affixed. Our evaluation of terminals is ongoing, and it is quite possible that our next large purchase will be from another manufacturer. ■■

Monitoring and Evaluating MELVYL

Ray R. Larson and Vicki Graham

In order for an information system, such as an online catalog, to provide effective service to its intended users, two things are required: first, the system must be flexible enough to change over time in accordance with the needs of its users, and second, some way must be found to determine those needs, providing feedback to the system design and development process.

In designing public access online catalogs, this process of feedback and refinement is even more critical than in other types of information systems, both because of the relative "youth" of online catalogs (with all of the unknowns that implies) and the wide range in experience of potential users.

Most online information retrieval systems (such as DIALOG, ORBIT, OCLC, and RLIN) have been developed with the implicit assumption that the users will be a well-defined group of persons who may be trained in the system's use. That assumption cannot be made about online public access catalogs.

From the beginning, designers of the University of California On-Line Union Catalog, now called MELVYL, recognized the need for evaluation and feedback, and decided that the evaluation process would need to incorporate both observations of actual user activity at the terminal (by transaction monitoring), and questionnaires to determine the motivations and needs of users. To make such data collection cost-effective and efficient, the University of California's Division of Library Automation (DLA) incorporated a software transaction monitor into the "Patron Interface" system and developed a subsystem to administer questionnaires online.

Transaction monitors have been used in

Ray R. Larson is a research programmer/analyst and Vicki Graham is a documentalist/bibliographer, Division of Library Automation, University of California.

a number of studies of online information retrieval systems to record user activity.[1] Monitoring is particularly well suited to studies of public access online catalogs, where the user's privacy is protected (there are no identifying account numbers).

The ultimate goal in designing the monitoring system for MELVYL was to capture enough information to be able to fully reconstruct any user session, including all commands, errors, and system responses, without intruding on the user or degrading system performance.

THE TRANSACTION MONITORING SUBSYSTEM

Each subsystem (such as display or search) of the MELVYL user interface program has routines for collecting data on the use and performance of the catalog. Each subsystem creates a record of its specific activity and passes the record to a common subroutine that adds information (such as time and date, terminal location, and session identification number) and accumulates the completed records in a buffer area in the computer's main memory. When this buffer area is full, it is copied to a disk file. When the disk file is filled, all "writing" from the buffer is switched to a second file and the first file is processed through a series of programs to provide statistical reports on system use.[2]

Transaction File Contents

The MELVYL transaction log file contains information on virtually every activity performed by the catalog, including: (1) copies of each command received from users; (2) reports on the results of each search; (3) reports on each display; (4) reports on user errors; (5) reports on help requests; and (6) reports on system activity.

Raw transaction data are collated and interpreted by a batch program producing nineteen different output files. Each record in each of these files contains identification information, such as session number, date, time, and command sequence number within the session. The identification information allows all records from the various files to be linked as necessary. The files and their contents are described below.

1. The session summary file contains one record for each user session. Session summaries provide totals for the number of searches, displays, help requests, errors, etc., along with identification information such as date, time, terminal, campus, search mode, etc.
2. The search summary file contains one record for each search, including such information as:
   a. Which indexes were used.
   b. System response time.
   c. Total items retrieved.
   d. Boolean logic and keyword modifications used (e.g., truncation).
   e. Search mode (COMMAND or LOOKUP).
   f. Whether the search was "normal" or "long." ("Normal" searches are those that can be resolved by the catalog search subsystem fairly easily—i.e., those with fewer than 1,000 "hits." "Long" searches are those with more than 1,000 "hits" that require special processing by the search subsystem.)
3. The display summary file contains one record for each "Display" screen transmitted to users. Each record includes the following information:
   a. Display format (i.e., brief, long, review, MARC).
   b. Number of records available for display.
   c. Number of records on this screen.
   d. Response time from display request to actual display.
   e. User "think time" from display to next command issued.
4. The browse summary file contains one record for each "Browse" (when a "Browse" is followed by "Select," the results are stored in the search summary file). Each record includes:
   a. The index browsed.
   b. Boolean logic and keyword modifications employed.
   c. Response time from the browse command to display of headings.
5. The help summary file contains records for each "Help" request. The records have a code indicating the situation in the user session at the time help was requested. If a glossary term was used in the "Help" command, it is in-

cluded in the record. (Situation codes include conditions such as "search just completed," "display just completed," "user just logged on," "user made an error," etc.)

6. The error summary file contains records for each user error. The records contain a code indicating the nature of the error. (For example, the user entered a "Find" command with no index specified; the user tried to "Display" when nothing was found by the previous search, etc.)
7. The command transition file includes one record for each transition between commands issued by users. It includes the following information:
   a. Code for the previous command issued (e.g., a LOOKUP mode search).
   b. Code for the current command (e.g., a LOOKUP mode display).
   c. Elapsed time between the two commands.
8. The search text file contains transcriptions of every syntactically valid search issued by users.
9. The browse text file contains transcriptions of every syntactically valid browse command issued by users.
10. The error text file contains transcriptions of every command that resulted in an error.
11. Five questionnaire record files store the user's responses to online questionnaires, and some additional information concerning how long it took to complete the questionnaire.
12. Two system activity files contain information about overall activity on the MELVYL system. One file records the times and dates that the catalog is available for use. The other contains summary information about catalog activity measured in (approximately) five-minute blocks. This includes:
    a. Number of log-ons.
    b. Number of log-offs.
    c. Number of users currently logged on at the end of the block.
    d. Number of commands issued by users.
    e. Average response time for searches during the block.
    f. Amount of time the system was idle during the block.

Three other files collect information on the general information, prompt, and help screens transmitted to users, transitions between catalog subsystems, and raw transaction log records based on various criteria. These three files contain little useful information for general analyses of MELVYL use and are only produced when required for special testing.

THE ONLINE QUESTIONNAIRE SUBSYSTEM

The questions used in the online questionnaire were formulated by a consortium of investigators from OCLC, Inc., the Research Libraries Group (RLG), Joseph Matthews and Associates, the University of California, and the Library of Congress for the nationwide Online Public Access Catalog Evaluation Project sponsored by the Council on Library Resources. The same questions—on machine-readable paper questionnaires—were asked of users of fourteen online catalogs other than MELVYL.

Since DLA acted as the data processing center for all questionnaires from the CLR project, we were able to compare the relative merits of paper and online questionnaires. The online questionnaire subsystem was designed with a simple menu selection and fill-in-the-blanks approach. It is incorporated into both the LOOKUP and COMMAND modes of the patron interface system. Each question is stored as an individual screen. When the questionnaire subsystem is activated, users are asked to select from a list of responses and type in the number of their choice (see figure 1).

The log-on software of the catalog permits a request for questionnaire participation to be issued automatically to every *n*th user logging onto the system. During the pilot cycle of data collection (10/26/81 to 11/22/81), the sampling parameter was set to ask every fifteenth user to participate. During this period 1,155 user questionnaires were collected, 47 percent of which were completely answered. The remainder contained partial data (where the user gave up before finishing the questionnaire). Following the pilot cycle, the CLR question-

```
1.  I came to this computer search with

        1. A complete author's name.
        2. Part of an author's name.
        3. A complete title.
        4. Part of a title.
        5. A topic word or words.
        6. A subject heading or headings.
        7. A complete call number.
        8. Part of a call number.

Type all numbers that apply, separated by spaces,
and then press RETURN.
->
```

Fig. 1. A Typical Screen from the Online Questionnaire.

naire was revised and shortened for the final cycle of data collection (4/5/82 to 5/31/82). For the final cycle, the sampling parameter was set to twenty-five, and 1,259 questionnaires were collected, 72.2 percent of which were complete. The distribution of questionnaires by UC campus and terminal location for both cycles reflected the normal use patterns of the catalog, indicating that the random sampling parameter was effective, providing a representative sample of users.

Data collection with an online questionnaire differs in a number of ways from data collection with a machine-readable paper questionnaire. There are advantages and disadvantages to each method, though we found that the advantages weigh on the side of the online questionnaire.

Once the paper questionnaires were printed, no further changes could be made without reprinting. The online questionnaire is much more flexible: it can be changed during a test cycle by adding new screens or by deleting or modifying existing screens. This can be done quickly and easily without affecting other parts of the questionnaire.

Paper questionnaires involve a series of administrative tasks. Data collectors have to be trained, data collection time blocks must be scheduled, and questionnaires must be distributed to sites. Serial number control must be maintained, and completed questionnaires have to be transported from sites to a central location for optical scanning. During scanning, problems can arise for a variety of reasons, including improper marking of responses, stray marks or fingerprints picked up by the machine, bent corners, incomplete data in the administrative section, and mechanical problems from the scanner itself.

The online questionnaire requires no administrative paperwork. Data collectors are unnecessary as the request for participation in the study appears automatically to every *n*th user of the catalog. Biases due to selection of respondents are eliminated. (Unlike human data collectors, the online questionnaire has no judgment over whom it will select, nor does it make subjective assessments of who "should" be asked to participate.) Raw data does not have to be transported from sites, and is already stored in a usable form. Administrative data such as serial numbers, time, and location are automatically supplied. Enormous quantities of data can be collected quickly and anonymously, and can be linked directly to the computer catalog session of the respondent, making it possible to analyze responses in the context of the actual searching behavior.

An online questionnaire has one major disadvantage: terminals that should be available for catalog searching are temporarily tied up by questionnaire respondents. The pilot questionnaire took about fifteen minutes, and may have caused queuing problems at catalog terminals. The final version of the questionnaire was considerably shorter, but still tied up terminals. A second disadvantage is in the questionnaire's sequential presentation: respondents cannot skip questions or sections and return to them, nor can they change a response once they have gone on to the next question.

TRANSACTION FILE AND QUESTIONNAIRE LINKING

The program that processes the transaction monitor data identifies all transaction records for users who have answered the questionnaire. Each of these transaction records may be selected subsequently from the transaction files and passed to another program which creates a new record for each questionnaire composed of: (1) the entire questionnaire record; and (2) selected data from session summary records, search summary records, browse summary records, display summary records, help summary records, and error summary records.

The linked questionnaire/transaction file provides data for comparing users' attitudes about MELVYL with what they actually did during their sessions. For example, in one of the questions users are asked whether (and to what degree) they agree or disagree with the statement, "The computer responds too slowly." User responses to this question can be compared with the actual response times of their sessions. This comparison can be further modified by other questionnaire information, such as amount of experience with MELVYL, experience with computers in general, etc., to provide estimates of the point at which response time becomes "too slow" for the average user.

WHO USES MELVYL?

Users of the system vary widely in age, educational level, and academic discipline. The data show users ranging from under fourteen years of age to over sixty-five, with educational levels from fewer than five years of school to more than sixteen. Table 1 shows the percentage of respondents grouped by age, level of education, and academic status.

The majority of users at UC are undergraduates between the ages of fifteen and twenty-four. This is consistent with the overall makeup of the student population at the nine campuses of the University of California, and does not necessarily imply that use of the computer catalog is affected by age or academic level.

Users come from a wide variety of academic disciplines. Table 2 shows users by academic area and the main focus of their academic work. The figures in this table generally reflect campus statistics on total student enrollment, academic discipline, and academic level.[3]

Questionnaire respondents were drawn from all nine UC campuses. Table 3 shows the number of respondents by campus, and compares these figures to the data collected by the transaction monitor showing catalog sessions by campus. Some of these figures seem disproportionately high or low considering the actual size of a particular campus. Los Angeles, the largest campus, has a fairly low percentage of total catalog use, and San Diego has the highest though it is a

Table 1. Age, Education, and Academic Status of MELVYL Users

	Percent
Age Group	
14 and under	1.7
15–19 years	23.4
20–24 years	46.8
25–34 years	18.8
35–44 years	5.4
45–54 years	1.7
55–64 years	0.8
65 and over	1.4
Current or Highest Grade Completed	
Grade school	2.6
High school	4.9
Some college	59.5
College graduate	33.1
Present University Affiliation	
Freshman/sophomore	23.2
Junior/senior	39.0
Graduate-masters level	6.4
Graduate-doctoral level	7.6
Graduate-professional school	1.9
Faculty	3.2
Staff	2.6
Other status	16.1

Table 2. Academic Disciplines of MELVYL Users

Academic Area	Percent
Art and humanities	24.2
Physical and biological sciences	22.7
Social sciences	19.3
Business and management	4.7
Education	2.0
Engineering	14.1
Medical and health sciences	5.8
Law	2.9
Major undeclared	4.0
Interdisciplinary	0.2

Table 3. Distribution of MELVYL Use by Campus

Campus	Percentage
Berkeley	17.8
Davis	11.3
Irvine	4.8
Los Angeles	14.6
Riverside	5.2
San Diego	22.7
San Francisco	1.4
Santa Barbara	12.2
Santa Cruz	6.5

smaller campus. Apparent inconsistencies of this sort in this table can be attributed to the number and placement of computer catalog terminals in the libraries and branches, and the percentage of campus holdings included in the catalog.

Use of Library, Computer Catalogs, and Card Catalogs

In the demographic section of the questionnaire, respondents were asked about their use of the library and its catalogs, their use of other computer terminals, and their familiarity with computer technology. The library is used daily by 42.3 percent of questionnaire respondents, weekly by 42.9 percent, and monthly by 8.8 percent. Others use it four times a year or less.

Table 4 compares frequency of use of the computer catalog with frequency of use of other library catalogs for both cycles of data collection.

More respondents are new users of the computer catalog than are new users of the card catalog. Since the computer catalog is still relatively new, this is not surprising. During the pilot cycle of data collection, the percentage of new users of the computer catalog was quite high. Of those new users, 87.0 percent said that they were very likely to use the computer catalog in the future, and 11.0 percent were somewhat likely to use it. Data collected in the final cycle show that the number of new users had already dropped considerably. Use of the computer catalog is expected to continue to increase as more patrons become familiar with it and as more catalog records are included in the database.

Most users found out about MELVYL by seeing a terminal in the library (81.0 percent) and learned to use the computer catalog by themselves (29.9 percent) or by using instructions printed on the terminal screen (56.4 percent). (See table 5.) Very few learned from library staff or other library-organized programs. Though 34.2 percent of respondents never used any kind of computer terminal before, most seemed to be comfortable enough with computer technology to approach MELVYL without help, and were able to learn to use the system by following instructions on the terminal screen. The conclusion that library staff assistance or training programs are unnecessary should not be drawn from this. Although patrons do not seem to want outside help, and claim to be satisfied with their search results, there are indications that they are not using the system to its full potential and often seem unaware of features that might benefit them.

Search Characteristics and Information

Respondents to the questionnaire use the catalog for a number of reasons ranging from gathering information for course work to pursuing personal hobbies or interests. Despite the many ways the information will be used, the majority of users come with a subject heading *or* topic word and search using subject headings *and* topic words. This is consistent with the data gathered by the transaction monitor (see table 6). (The terms *subject heading* and *topic word* were presented in the CLR questionnaire without definition, relying on the respondents' own perceptions of difference between the terms. Many respondents appear to find the two categories indistinguishable and indicated both in their responses.)

HOW IS MELVYL USED?

MELVYL's transaction monitoring subsystem has been operating almost continuously since June 1981, producing an average of 52,000 raw transaction records every eight hours. This rather awesome mass of

Table 4. Frequency of Use of Library Catalogs

	% Computer		% Other	
Frequency	Pilot Study	Final Study	Pilot Study	Final Study
Every visit	14.8	13.9	17.7	16.6
Almost every visit	33.5	34.1	32.2	29.4
Occasionally	29.1	35.6	31.2	38.0
Rarely	3.9	7.2	14.1	11.1
Not before today	18.7	9.2	4.8	4.9

Table 5. Sources of Assistance in the Use of Terminals

	Percent of Cases
I got help in doing this search from:	
Printed matter or signs	17.2
Terminal screen instructions	43.9
Library staff member	4.4
Person nearby	4.0
I did not get help	44.2
I use another computer system:	
Daily	14.0
Weekly	20.1
Monthly	11.2
Four times a year	12.5
Once a year	7.9
Never	34.2
I learned how to use this computer catalog:	
From a friend	8.3
By printed instructions	31.0
By instructions on terminal	56.4
From library staff	5.4
From library course	2.0
By myself	29.9

Table 6. Search Characteristics

	Percent of Cases
I came to the search with:	
Complete author's name	35.3
Part author's name	12.4
Complete title	27.0
Part title	9.9
Topic word or words	34.0
Subject heading	49.0
Complete call number	2.0
Part call number	1.3
I searched for what I wanted by:	
Complete author's name	32.9
Part author's name	15.1
Complete title	25.3
Part title	11.3
Topic word or words	36.4
Subject headings	47.9

information is accumulated on magnetic tape (about twenty-five reels). The following discussion of MELVYL use is based on a single month's data collected during May 1982. It should be noted that the interpretations are based on the complete transaction records for the month, not a sample.

The Average User Session

During May 1982 the transaction subsystem collected records from 65,468 sessions (a mean of 2,111.87 sessions per day). This section describes the "average" user session during that month.

The majority of users (59.2 percent) chose LOOKUP mode for their sessions, and the remainder chose COMMAND mode.

The mean length of time that a user spends at the terminal (from log-on to log-off) is 8 minutes, 41 seconds (521.361 seconds).

The mean number of commands (of all types, i.e., find, display, and help commands, and also responses to prompts, carriage returns to continue displays, etc.) issued in a session was 22.297.

The user conducts a mean of 5.66 searches per session. (The maximum observed was 224 searches in one session. This is probably an example of the "novelty factor" often observed with first-time users of the catalog. New users seem to look at the catalog as if it's a video game, and "play" with it—for instance, by searching for every author they can think of.)

Only 14.28 percent of the sessions used the help facilities. Those who did use "Help" issued a mean of 1.59 unqualified help requests and 2.63 help requests with a glossary term (e.g., "HELP INDEXES").

Only 2.60 percent of the users made the same error three times in a row and were automatically sent a help screen. (The MELVYL help subsystem has been designed to aid users who get "lost" by automatically sending appropriate help screens when the same error is made three times consecutively.)

A mean of 10.45 display screens were sent to users per session.

A rather surprising 75.22 percent of the users made no errors at all during their sessions. (This includes command syntax errors, logical errors—such as trying to "display" when no records have been retrieved, and issuing unrecognizable commands.) The other 24.77 percent made an average of 2.85 errors during their sessions.

Only 3.66 percent of the users made use of the browse command, but those who did used it an average of 2.66 times.

Displays

During May 1982, 428,154 records of

displays of retrieved records were collected; of these, 40.91 percent were from COMMAND and 59.09 percent from LOOKUP mode sessions.

For the LOOKUP sessions, 44.5 percent of the record displays were in review format (the automatic default), 40.7 percent were in long format, and 14.8 percent were in brief format. Users took an average of 32.93 seconds to absorb the information on the screen and issue the next command.

For users who chose COMMAND mode, 66.3 percent of their displays were in brief format (the default), 14.0 percent were in long format, and 6.0 percent were in review format. Of the remainder, 9.9 percent of the display screens sent to COMMAND mode users were displays of lists of headings retrieved by browse commands, 3.0 percent were displays of individual fields or groups of fields from the MARC record, and only .8 percent were full MARC displays. The tendency for COMMAND mode users to stay with the default (brief) display, while LOOKUP mode users go to long or brief format from the default (review) display, can be attributed largely to the lack of call numbers in review format displays.

For COMMAND mode displays, the user took an average of 22.49 seconds to absorb the information in the display and issue a new command.

Searching Patterns

Perhaps the most interesting and significant observations from the transaction monitor data are the ways people search for things in an online catalog.

There have been many studies of how people use card catalogs in libraries, but studies of online catalogs are just beginning. The early results presented here indicate that online catalogs are being used in a qualitatively different fashion from card catalogs.

In a major study at Yale University,[4] 73 percent of card catalog users were searching for "known items" and only 16 percent were conducting subject searches. Lipetz suggested that many of the so-called known-item searches were actually disguised subject searches, making the "real" breakdown 56 percent known item and 33 percent subject searches. A study by Palmer at the University of Michigan found a similar 70 percent known-item searches.[5]

Swanson, in reviewing sixteen catalog use studies (representing approximately 8,000 searches), suggests that the average searching pattern is 60 percent known item and 38 percent subject searching.[6]

These studies have obvious limitations: the researchers cannot easily directly observe what the user is doing, and must rely on interviews or questionnaires where the users describe their intentions in searching the catalog. By monitoring, however, we are able to observe directly what indexes the users are searching.

The searching patterns of users of the online catalog are rather dramatically different from those observed in the card catalog studies. Table 7 shows the combinations of indexes used by COMMAND mode users in searching MELVYL during May 1982. The most commonly used index, by far, is the subject index (SU). When the possibility of "disguised" subject searching is considered

Table 7. Indexes Used in COMMAND Mode Searches, May 1982

Indexes Used	Number	Percent
SU	36707	51.6
PA	15325	21.5
TI	13378	18.8
TI & PA	2077	2.9
CA	1306	1.8
SE	934	1.3
SU & TI	494	0.7
SU & PA	403	0.6
UT	290	0.4
TI & CA	128	0.2
SU & CA	56	0.1
SE & SU	22	0.0
SU & PA & TI	19	0.0
PA & CA	16	0.0
SE & PA	8	0.0
SE & TI	5	0.0
SE & CA	4	0.0
SU & PA & CA & TI	2	0.0
TI & PA & CA	1	0.0
UT & CA	1	0.0

Key:
- CA = Corporate Author Index.
- PA = Personal Author Index.
- SI = Series Index.
- SU = Subject Index.
- TI = Title Index.
- UT = Uniform Title Index.

(e.g., searching for topical terms in the title [TI] index) the actual percentage of subject searches may be closer to 60 percent.

Table 8 shows the combinations of indexes used in LOOKUP mode searches during May 1982. In LOOKUP mode the user is given a choice of either an author-title or subject search. MELVYL automatically constructs the actual search statement so that an author-title for author *X* and title *Y* becomes: "FIND PA X OR CA X AND TI Y" (that is, Personal Author *X* OR Corporate Author *X* AND Title *Y*). MELVYL constructs subject search statements so that searches for subject Z become: "FIND SU Z OR TI Z." In LOOKUP mode, where users are given this clear choice of known-item (author-title) or subject searching, the differences between card catalog and online catalog use patterns are most apparent.

Table 8. Indexes Used in LOOKUP Mode Searches, May 1982

Indexes Used	Number	Percent
SU & TI	71019	63.4
PA & CA	17240	15.4
TI & PA & CA	14780	13.2
TI	9013	8.0

The differences between searching patterns in the online catalog and the card catalog suggest (what many librarians have suspected) that subject searching in card catalogs is limited not because users don't want to do subject searches, but because the way subjects are presented and arranged in card catalogs is not conducive to use. As Mooer's law of information retrieval states, "An information retrieval system will tend *not* to be used when it is more painful and troublesome for a customer to have information than for him not to have it."[7] When users have a better way to get at subjects, they will tend to do subject searches, and the online catalog certainly can provide a less "painful and troublesome" method.

One conclusion that we can draw from these early observations is that the online catalog is not going to be used just like the card catalog, and online catalog designers cannot rely entirely on studies of card catalog use. We do not have to re-create the card catalog online; we can create a new, more sophisticated, and at the same time less "painful" way of exploring the riches of our libraries.

WHAT DO USERS THINK OF MELVYL SERVICE AND FEATURES?

Overall User Evaluation

Response to MELVYL is very favorable. Respondents are satisfied with the results of their searches and intend to continue using the computer catalog.

The figures in table 9 are very interesting, considering that only 13.5 percent of users actually found all or most of what they were looking for, 11.8 percent found more, and 31.5 percent found nothing. What users seem to be expressing is a favorable attitude towards the *capabilities* and *potential* of the system once the database is complete. A look at the series of questions comparing the computer catalog with the card catalog (table 10) verifies this. (Table 10 contains data from the pilot cycle. This series of questions was eliminated in the final questionnaire.) The computer catalog came out on top for all categories except one: to search among all books in the library. Most respondents found it faster, and preferred it for finding books published in recent years, for finding a few books on a topic, and for searching for a specific book. Most people thought that there was no difference in the difficulty of learning to use the two catalogs. Of those who found a difference, more thought it was easier to learn to use the computer catalog than to learn to use the card catalog. In the final cycle of data collection, these questions were con-

Table 9. User Attitudes toward MELVYL

	Percent
In relation to what I was looking for this search was:	
Very satisfactory	32.7
Somewhat satisfactory	33.5
Somewhat unsatisfactory	14.9
Very unsatisfactory	18.9
My general attitude toward the computer catalog is:	
Very favorable	70.8
Somewhat favorable	22.1
Somewhat unfavorable	3.2
Very unfavorable	3.9

Table 10. User Comparisons of the Online and Manual Catalogs

	Computer Catalog Superior	No Difference	Card Catalog Superior
Comparison between Computer and Card Catalogs (Pilot Cycle Data)			
Searching speed	78.0	11.5	10.5
To search through all books in the library	44.6	8.0	47.4
To search for a specific item	47.4	18.5	34.1
To search for recently published books	58.6	24.2	17.2
To find a few books on a topic	69.8	13.2	17.0
To scan through several titles	75.5	11.2	13.3
To learn to use without assistance	37.3	40.9	21.8
To prepare a comprehensive bibliography	49.5	22.3	28.2

	Percent
Computer Catalog Compared to Other Library Catalogs (Cycle Two Data)	
Better	68.3
Equal	17.3
Worse	14.4

densed to one general question comparing the computer and card catalogs. The computer catalog was considered better by 68.3 percent of respondents.

User evaluation does not differ significantly according to the mode (COMMAND or LOOKUP) being used, though a slightly higher percentage of COMMAND mode users found more than they were looking for, and a slightly lower percentage found nothing. Satisfaction with the catalog depends on a number of variables, some demographic, and some related to catalog features. No single variable has been identified that can account for differences in levels of satisfaction.

Use of Features; Problems

A major section of the questionnaire asked for opinions on a variety of issues pertaining to the catalog, including ease or difficulty in use of the terminal and its features, availability and usefulness of instructional aids, and legibility of displays. Respondents were given a list of statements about these issues and were asked to indicate how strongly they agreed or disagreed with them. Responses to this section show that users have no difficulty with most features of the system. They find both bibliographic and instructional displays understandable, and are comfortable with the library environment. Some users report having trouble remembering the limitations of the prototype database (e.g., only partial coverage of holdings). Others find it difficult to increase the results of their search. (Both of these problems will be partially solved by increasing the database.)

User Priorities for Improvement

Although users are favorably disposed towards the catalog, they are also eager to offer suggestions for improvement. Table 11 shows the results of three questions asking respondents to indicate their priorities for service improvements, additional catalog features and capabilities, and types of material they would like to see added to the database.

On top of the list for added features is the request for an online thesaurus or some other method of viewing a list of related search words. Many respondents also asked for the ability to search through a book's table of contents and index. While such a capability is still impractical to implement, the desire for it reinforces the observation that users tend to view the catalog as an information retrieval tool, rather than a device for locating specific known items. Both of these requests reflect a desire for increased sophistication in subject searching.

Many users request that more of the library's other books and materials be included in the catalog. Heading this list is a request for journals and magazines, followed closely by a request for the library's older books (many of which are being

Table 11. User Priorities for MELVYL Improvements

	Percent of Cases
Additional System Features Desired	
View a list of search words related to my search word	57.3
Know if a book is checked out	57.0
Search table of contents or index of a book	49.3
Print search results	45.5
Search by journal title abbreviations	26.3
Know library location of book	24.4
Limit search results by language	20.6
Search by call number	20.5
Search for illustrations and bibliographies	19.3
Limit search by date of publication	18.5
Change the order in which items are displayed	13.2
Computer Catalog Service Improvements Desired	
More terminals	63.6
Terminals in locations other than near card catalog	36.9
Terminals outside library	35.7
Command chart posted at terminal	30.7
Manual at terminal	16.9
Training sessions	11.5
Manual for purchase	10.3
A-V training program	9.3
Additional Kinds of Materials Desired	
Journal titles	64.0
More older books	39.6
Government publications	37.1
Newspapers	34.9
Dissertations	34.5
Technical reports	24.8
Motion picture films	24.4
Records or tapes	23.0
Music scores	13.5
Maps	13.2
Manuscripts	10.8

added now by retrospective conversion). Users also would like the catalog to include government publications, technical reports, dissertations, and newspapers.

Users give high priority to the ability to obtain printed copies of their search results. In addition, they would like to know whether or not a book is checked out before they leave the terminal. They also would like to see the number of terminals increased, both in locations outside the library as well as locations in the library away from the card catalog. Users would prefer a chart of commands posted at the terminal rather than more staff assistance or other library-organized training programs.

Other Observations

Given the opportunity to type in comments, respondents were glad to offer further suggestions and opinions. DLA now has an extensive file of comments and ideas gathered directly from catalog users. They range from requests for more complex system features to ideas for improving service, to pleas for expansion of the database:

"I can hardly wait until more records are added!"

"More of everything, just more, more, more!!!!!"

Responses are overwhelmingly positive. Most users who ask for improvements include praise for the existing system in their comments:

"Put more of these in the libraries. They make life much!!!!!!! easier."

"The catalog is great!!!!!! Please expand the system."

"This is fun!"

"This system is much easier to use than the card catalog."

"I hope that this prototype turns into the real thing!!!!"

"Thank you very much computer, it was grand."

The information collected from and about users of the MELVYL prototype is proving to be a valuable resource in the design of the catalog. In place of nebulous speculations about hypothetical users, we have a wealth of hard data from actual users of the prototype system. We know where the users have problems, and what features they like. Perhaps the most important observations are that the users like MELVYL, that they are not intimidated by "the computer," and that they plan to continue using it.

ACKNOWLEDGMENT

This research was supported, in part, by a grant from the Council on Library Resources.

REFERENCES

1. W. D. Penniman and W. D. Dominick, "Monitoring and Evaluation of On-Line Information System Usage," *Information Processing and Management* 16:17–35 (1980).
2. Ray R. Larson, *Evaluating Public Access On-Line Catalogs: Phase I, Development and Testing of Data Collection and Analysis Tools. Final Report to the Council on Library Resources* (Berkeley: Univ. of California, Division of Library Automation, July 1981), 68p.
3. University of California, Office of Information Systems and Computing, *Statistical Summary of Students and Staff of the University of California, 1979–1980* (Berkeley: Univ. of California, 1980), 121p.
4. Ben-Ami Lipetz, "Catalog Use in a Large Research Library," *Library Quarterly* 42:121–39 (Jan. 1972).
5. R. P. Palmer, *Computerizing the Card Catalog in the University Library: A Survey of User Requirements* (Littleton, Colo.: Libraries Unlimited, Inc., 1972), 141p.
6. D. R. Swanson, "Requirements Study for Future Catalogs," *Library Quarterly* 42:305–15 (July 1972).
7. C. N. Mooers, "Mooers Law or, Why Some Retrieval Systems Are Used and Others Are Not," *American Documentation* 11, no.3:ii (July 1960). ■■

Online Catalogs: Through a Glass Darkly

Edwin B. Brownrigg and Clifford A. Lynch

Other articles about DLA in the current and previous issues have reviewed the past and described the present. In this article we turn our attention toward the future and attempt to draw conclusions from our experiences. We will make generalizations and offer speculations. Some of these may seem outrageous; some will undoubtedly prove incorrect. Prognostication is a black art at best. Perhaps the greatest danger is caution—the twin errors that Arthur Clarke characterizes as "failure of imagination" and "failure of nerve" in his collection of cautionary essays, *Profiles of the Future*.

We believe that it is important to try, however speculatively, to foresee the future. Fifteen or twenty years ago very little thought was being given to the possibilities of online catalogs. This lack of foresight has left us with tremendous problems. In some sense we may have lost control of the proliferation of information: while the volume of publication has increased tremendously, the cataloging and classification practices of the last decades have left us impoverished. Although the technology to provide access to this material is finally at hand, we find ourselves in a poor position to take advantage of it.

Online catalogs serving single libraries, groups of libraries such as the University of California, or library consortia within a relatively restricted geographical radius are now an operational reality. It seems safe to assume that such systems will proliferate over the next decade. Patron acceptance to date has been overwhelming. For the purposes of this paper, we will treat the expanding availability of such systems as a given. Our focus will be on the discrepancies between the services such systems offer and what the library patron really seems to

Edwin B. Brownrigg is the director of the Division of Library Automation (DLA), University of California, and Clifford A. Lynch is the manager of the Computing Resources Group at DLA.

want. We will also probe some of the gaps that we believe will develop in the fabric of library access, as such systems spread in an uncoordinated and independent fashion.

DLA's online catalog has been an experiment. While future system designers can undoubtedly learn valuable lessons from its successes and failures, our experience offers insights that are far broader in scope. Tied up in the UC online catalog experience is a commentary on the library practices of the last half of the twentieth century. Perhaps our experience can be used to point to directions for the evolution of library automation (and hence librarianship). At the very least it can highlight policy alternatives and problems. Ultimately, the user response to the UC system raises philosophical questions about the relationship that has developed to link the user, the librarian, and the body of information the library historically has housed.

THE ACCESS RENAISSANCE

For a century, and in some cases longer, the large research libraries have been building formidable collections. A librarian is awed by the endless ranges of books housed in our marble monuments. Generally, however, these libraries have closed their stacks to patrons, and the original catalog cards have not held up under decades of use. Often, too, the prominence of the authors of the older works has dimmed. If the cataloger's subject analysis of such works was not deep, material pertinent to research may never be uncovered.

Invaluable information made inaccessible by the darker ages of library science can be brought to light again when librarians redirect their energy from just "collecting" books to augmenting the means for scholars to find them. Research is a serendipitous act. If a researcher knew exactly what he was looking for, he would not be doing research. Because the researcher often is denied direct access to the books themselves, he must take an indirect path to them through bibliographic apparatus. Due to the sheer magnitude of the bibliographic nexus, the printed media of card, book, and fiche catalogs have severely fragmented this indirect path. With the passage of time, the bibliographic path has become a maze.

With the advent of information technology, librarians have been given an opportunity to solve this access problem. The solution involves some basic ingredients. Among the more arduous is converting old cataloging data into machine-readable form (retrospective conversion). Among the more expensive is building the computing engines that bring order and harmony to the bibliographic fury. Among the more creative will be novel ways to index and retrieve bibliographic information. Among the more dramatic will be delivering documents electronically through the same medium—the online catalog.

The primary orientation of the library community is shifting (and must shift) from developing collections to providing access. Online catalogs quickly provide major improvements in this area, placing the catalog at the user's fingertips at any time of day or night through dial-up access and remote terminals. The more sophisticated online catalogs provide searching capabilities that are vastly superior to card catalogs, a much faster and more attractive user interface, and, at least to a limited extent, the ability to search for materials in a number of libraries simultaneously. This is unquestionably a major achievement, but one must be cautious before taking too much satisfaction in it. If the studies of patron access and other comments on online catalogs have any clear theme, it is that many of our assumptions about what type of access the user desires are without foundation. In a very real sense we have provided a modern, computerized version of a manual system that seems to have had some rather serious flaws. Automated systems allow the user to circumvent some of these problems, but as much by unforeseen application of user ingenuity as by any great insight on the part of the system designers.

CATALOGING FOR ACCESS

There has been a long-standing assumption that the typical library patron is performing a known-item search. It is difficult to tell, at this late date, whether this was the impetus for or the result of historical cataloging practices. The myth of known-item searching has finally been put to rest by the recent Council on Library Resources (CLR) study of online catalog use.[1] In an online environment at least, unknown-item searching is very common.

Data from the CLR study suggests that a major part of patron searching activity is subject searching for unknown items. It is evident from inspecting the transaction logs from the DLA catalog, or from a few hours of experimentation with any of the systems that support subject searching, that such searching presents major difficulties, particularly to the typical user who has not been initiated into the mysteries of the LC subject heading scheme. Part of the problem is lack of subject authority control that could provide better standardization of terms, and more alias access points to popular terms. A much larger part of the problem is that the typical user seems to view the subject thesaurus as a partially hierarchical scheme in which terms can be selected to either generalize or refine a given set of subject terms. In order to mirror this conception in an online catalog, it will be necessary to provide access to syndetic relationship structures, as found in the eighth edition of the *LC Subject Headings*, and a means for the user to navigate among them. Taking a broader view, it seems that increased attention to enriching subject classification and cross-referencing is perhaps the single highest payoff area for improved user access.

For example, in an open-stacks library the patron can browse through the biology section. Within this section can be found books on all subclassifications of biology. In an online catalog, however, asking for books on the subject of biology will not provide books specifically covering subfields of biology such as microbiology. The "genericness" or "specificity" of a given subject term is a difficult concept for the user to grasp, and an extremely frustrating one. Users approaching a subject either from a very naive point of view or a very sophisticated point of view may have great difficulty in locating desired material.

The other side of this problem is keeping retrievals to a manageable number. We still have little experience with really large online catalog databases. Even with MELVYL's modest 750,000 records, many searches retrieve 50 to 100 records. With a database of, say, 4,000,000 titles, the typical search may retrieve several hundred titles. This is unmanageable. The user's major problem may prove to be refining the search results to a usable size. In this situation, very specific subject headings—far more specific than many of those currently in use—will be needed, as will the ability to broaden or narrow the search. With access to the syndetic structures implicit in an adequate subject authority file, the user potentially could direct the online catalog to broaden automatically a search result or to provide alternatives (narrower terms and subdivisions) for refining it.

Much of the existing nonprint material collected by libraries is poorly served by current online catalogs. Part of this weakness stems from the very reasonable tendency to concentrate on the retrieval of print media first, since this constitutes the bulk of library use. The nonprint media in question include graphics (photograph collections), maps, music, sound recordings, films, etc. In all cases there is a general theme: very little thought has been given to how to organize and catalog these materials so as to be helpful to the typical user. Most or all of these areas have specialized classification schemes, but in most cases they are not oriented toward public access retrieval. How, in an online catalog, should a user indicate he wants a map of known ore deposits on the Australian continent, or a record of the Brandenburg concertos performed by a baroque chamber music ensemble rather than by a modern orchestra? Until such problems are solved, there is little to be gained by loading most nonprint MARC records into an online catalog. In fact a crude approach may well be detrimental if the indexing is done incorrectly. Suppose, for example, that the default is to include records of all types in the database, and a user asks for material on Africa. Along with the standard huge result, hundreds or even thousands of maps will be retrieved.

In this connection, it should also be pointed out that the user interface techniques used in today's online catalogs are still extremely crude from at least two points of view. The user directs the catalog with one of two approaches—a menu-based selection, which is initially appealing to the naive user but rapidly becomes a major liability, or a command language. While the better command languages are English-like and have a somewhat natural feel to them, they still require the user to learn a language. The English-like nature of the syntax can be insidious, since it creates a notion on the part of the user that the catalog actually understands English. The user is then disappointed at how "stupid" the system is when it fails to cope with apparently minor variations of the correct command-language syntax.

Considerable progress has been made in recent years toward the goal of computer understanding of natural language. For a limited universe of discourse (that is, a system that can only do a small number of things), it is possible to produce reasonably effective natural-language parsing and analysis programs. Online catalogs, at least today, offer an extremely limited universe of discourse. Ideally, it should be possible to supply a natural-language mode for the casual user. Some experiments have been conducted in this area under DLA's auspices; results to date have been encouraging although far from conclusive.

The second area in which the user interface can be greatly improved is by awareness of content as well as syntax in diagnostics. For example, MELVYL's help facility deals only with what the user has typed in. It does not, for example, offer narrower terms to help refine a search, nor does it examine individual components of Boolean queries, looking for those that produced no matches in order to suggest possible misspellings. More broadly, many systems offer help facilities that explain what the user has done and what can be done next, or how to use specific facilities. In some cases, the user doesn't want help, however, but rather hints on possible alternatives or extensions to what he or she has already done.

REDISCOVERING ANALYTICS

Serials present an interesting problem with regard to subject access. In many fields, especially science and engineering, journal articles rather than books are the primary forum for new developments. While it would be relatively easy with the existing MARC serials database to add serials to an online catalog as a sort of many-volume monograph, for searching purposes this is probably useless (or actually, worse than useless, just unwanted clutter) to the average online catalog user. (It is useful for known-item searching, where the user wants holding or location information, of course, but this is the smallest part of the online catalog revolution.) In order to provide the library patron with better access, we must return to analytic cataloging of journal *articles*, and integrate monographs and *articles* (rather than monographs and journals) into the same database.

For years, librarians have been teaching patrons and each other that serials are different than books and that they should be accessed differently. But different access to serials does not necessarily follow. Rather it has followed merely for practicality, and it has continued to fragment the patrons' bibliographic approach to research. The demise of analytics as a standard cataloging practice left the patron with the notion that the catalog was the primary source of bibliographic information and that access to everything else, including serials, was less than primary. The online catalog removes all technological barriers to revitalizing analytical cataloging. Much that has been bibliographically hidden could be uncovered. Yet it is ironic to note that, as Marshall McLuhan taught, the contents of new media are initially old media. The very large databases that index journal articles are a perfect example. Rather than reintegrating access to journal articles with monographic catalogs, computing technology initially has widened the gulf.

Analytics were discontinued nearly forty years ago by most major libraries. They were, even then, becoming prohibitively costly. Moreover, today most libraries are hard pressed merely to keep up with monographic cataloging.

Since North American libraries have embraced online cataloging, a subtle but revolutionary process has taken place. Cataloging only appears to be done in the library. In reality very little cataloging is performed anymore; copying cataloging from another source has taken the place of cataloging proper. However, the significance of this shift has not to do so much with *how* the "cataloging" is accomplished, but rather *where it is done*. The locus of cataloging has moved out of the libraries. Today it is dispersed or shared, with a large concentration still with the Library of Congress. The cataloging copy is loaded onto the large database utilities and telecommunicated into the libraries. In all likelihood, the shift in the locus of cataloging will continue until the entire phenomenon of cataloging has changed dramatically.

In order to predict the future of the source of cataloging copy, at least two broad trends have to be considered. One is that with computing technology harnessed for cataloging, any financially motivated group can become a distributor of cataloging copy. The other is that libraries alone cannot cope with the cataloging of evanescent information such as journal articles, government documents, and technical reports. The grant-funded CONSER project is ample evidence of the inability of individual libraries to cope with such cataloging. For all its effort, CONSER cataloging is at the journal *title* level, not the journal *article* level. Today, a very large proportion of published material is not cataloged by libraries at all. Rather it is done by publishers of indexes to journals, who *catalog* their content. Only they do not call it cataloging, they call it indexing. Maybe they also call it advertising.

With the online catalog, publishers' indexes of journal articles (analytic cataloging) can be reintegrated with monographic cataloging. The librarian and the patron finally can have bibliographic access to both material formats within the same catalog—just like it once was many years ago, except now handled electronically.

CATALOGING AS ADVERTISING

But will the publishers stop with indexing journal articles? Already many contribute data to Cataloging in Publication (CIP). In the future there could be Cataloging *at* Publication because such practice would sell new monographs faster. (People won't buy something they don't realize exists.) In order to fully appreciate this, it is important to shift from the concept of "book" cataloging, to that of "monographic" cataloging. With the advent of electronic publishing, monographs will not be books per se, but something else—maybe just monographs.

Change will not halt with the locus of cataloging. The locus of usership will change simultaneously. Patrons could cease to be library patrons and become instead publishing-company patrons. Already we see users accessing online catalogs not from libraries, but from their homes and offices. As more and more journal articles are delivered electronically through the online catalog network, libraries will significantly change in at least two ways.

The most revolutionary change in the nature of libraries will be that they will become distribution points or switches for electronic information, both bibliographic and textual. Telecommunications topology for online catalog networks will be determined in great measure by the geographic distribution of libraries. Library buildings will house telecommunications concentrators attached to telephone lines, optical fibers, and a variety of radio devices.

Libraries will also revert to an earlier function: archiving the information stored on paper and photographic media. The frequently used hard-copy information sources themselves will be converted into digital image form, becoming part of the electronic store of information. The statistical data used to select material for conversion will be gathered, of course, from now-traditional library automation systems, such as circulation systems and online catalogs.

Uppermost in the thinking of library directors and budget directors will be the issue of cost. How much will the transition to modernity cost? How will such costs be justified? Will there be offsets? Transition costs in libraries historically have been high, usually because of the need to run a new system in parallel with the old one it is to replace. As for justification, it will inevitably devolve to arguments like those that caused capital investment in aviation, following a century or so of rail transportation.

There should also be substantial cost offsets. Consider, for example, the logical consequence of large portions of journal articles being published exclusively electronically; serials check-in would cease.

But those libraries that do not make the investment will take on the air of old railway stations: a few custodians will guard the halls, but the public will not be waiting, because "the train doesn't go there any more."

DOCUMENT DELIVERY AND ELECTRONIC PUBLISHING

A number of activities are lumped together under the general rubric of document delivery and electronic publishing. It is necessary to begin any discussion of this area by distinguishing among the possible activities.

Document delivery at a terminal is concerned primarily with delivering text (and perhaps graphics as well) that first exist in some other, typically printed, form. The implication is that the material either will be digitized retrospectively for presentation over the computer network, or the publisher of the material will provide it in an alternate form in addition to the printed media.

Electronic publishing, in its pure sense, is a far more radical service than document delivery. In electronic publishing, the computer network becomes the primary medium for the creation, storage, and dissemination of a document. This is a fantastically powerful concept that changes the way people think about information, rather than merely changing the way they get access to it. Some of the implications of electronic publishing are that it becomes cheap to create and store information, and that information can be retrieved selectively in small pieces. Consequently, an electronic journal can be far more comprehensive than a printed journal. In a printed journal, users can ignore the parts of the journal they don't want (and can stop looking at the journal if they are bypassing too much of it). In electronic publishing they only need to select the articles of interest. The "unit" of publication thus becomes the article or news summary rather than the journal itself.

With the advent of electronic publishing, the accidents of physical grouping will be eliminated. The role of the editor will change, and the periodical as we think of it may lose its identity. Instead of buying a periodical because its editor selects and presents high-quality material that is often of interest to the user, with each issue forming a random or specific grouping of material, articles may simply be marked as "endorsed" by a given known and respected editor. Selection rather than organization may become the major function of the editor of the future.

Electronic publishing is not without its dangers. As we said before, research is a serendipitous act. In today's world, if one reads periodicals, one at least glances at many articles far afield from one's main interest. Occasionally one of these articles catches the eye and the reader makes an unexpected connection. (The editor of the journal has thus played a part in arranging this chance meeting.) In an electronic publishing environment, it will be very easy to select and obtain only material in a specific field of interest, eliminating much of the current random browsing. This may further aggravate the fragmentation of many fields, producing ever-larger numbers of very small and insular research communities that don't talk to each other.

Electronic publishing has different implications for popular journals than for scholarly journals. For popular journals it allows for very timely material, since the journal can in some sense be "continuously" published. New material can be made available as it arrives. For the scholarly journal, electronic publishing offers cost control: it reduces production costs and makes it possible to avoid numerous typesetting and editing passes. Electronic publishing also eliminates the multi-year backlogs of unpublished articles now plaguing many major scholarly journals. In the case of scholarly journals, the publisher may well bow out, except perhaps as an indexer. An article could be published onto the network directly by the author. Refereeing could be handled by certifying the document directly (perhaps with a form of public key cryptosystem, both to maintain anonymity and to ensure that an author's claim to acceptance of his work is legitimate).

Aside from technological problems, electronic publishing creates a variety of difficulties that are ultimately sociological in nature. There is the question of acceptance within a research community. Until enough people within that research community are using electronic publishing as a primary means of distributing information, producing a "critical mass," nobody will want to publish electronically. The situation is analogous to the early days of the telephone, when people declined to install phones because there was nobody for them to talk to. There is no incentive for the author to publish electronically (other than the new capabilities of the electronic media) and there are a number of disincentives. How do you cite an electronic document on a publications list? Will peers take such publications seriously in evaluating research activities for promotions and awards?

There are also economic questions. In a network environment, a master copy of a document must exist somewhere. Who puts it there and where is it kept? In academic settings, for example, will each university maintain faculty research documents and theses on its local computer as part of its basic service? What happens when someone moves from one institution to another? In traditional publishing, there is a very real sense that the work passes from the author's control at the point of publication. This is very important in assigning precedence for research and can have major implications in such areas as the award of patents. What will be the point of demarcation in electronic publishing?

Royalties and copyrights also present nightmarish complexities in an electronic publishing environment. As matters stand today, suppliers and users of electronic publishing services cannot even enunciate a clear set of policy alternatives to lawmakers; consequently, it is unlikely that the legal structures will be changed to accommodate electronic publishing successfully in the near future. The best that can be said is that there is a growing awareness of the problem. The December 7, 1982, *Wall Street Journal*, for example, carried a story that opened with the bald statement, "Electronic publishing will be the death of copyright law."

Major technological hurdles remain to be overcome. In several cases, a technology that is currently experimental must pass into cheap mass production. Resources that are now in short supply will have to become virtually unlimited. There are two basic approaches to the transfer of textual information. The common approach today is textual: treat the document as a string of characters, using seven or eight bits per character in some standard code like ASCII or EBCDIC. This approach, while economical and simple to work with, is quite limited. It will not cope with graphics of any kind. Worse, it will not cope with typographic symbols beyond those on a standard typewriter—in particular, the language of science and mathematics. There is no provision for integral signs, Greek letters, sub- and superscripts, and the like.

At the other extreme is the facsimile approach: treat all material as a picture, and use high-resolution graphics terminals to display the information. This is enormously costly in storage space and telecommunications bandwidth—perhaps two orders of magnitude more costly than character coding. For electronic publishing, it is problematic, since the text is usually produced by scanning printed material or by taking the final machine-readable photocomposition output just before it is run on the photocomposer. It is not a convenient form for an author to generate.

Today, very few devices can display or provide hard copy for facsimile-type data, and such devices are very expensive to purchase or operate. We will have to develop a middle ground: a character-set standard for electronic publishing that encodes the standard alphabet efficiently, provides escape mechanisms for the simple nonstandard graphics, and also allows for embedded facsimile-type information. In conjunction with this new character-set standard, we will need presentation-level protocols that map the character set as well as possible onto a host of output devices with a vast range of costs and capabilities, all the way from a simple terminal that can display only the base alphabet directly (on which pictures might be totally suppressed and special symbols listed as *integral* or *epsilon* where they occur) through laser-driven photocomposers and high-resolution monochromatic or color bit map terminals.

Very inexpensive media will be needed to house the endless terabytes of data required. Most likely optical disks will fill this role well, particularly given the low volatility of the information being stored. Machines will be needed to retrospectively convert printed documents. These will evolve from today's optical character recognition technology and machines such as the Kurzweil talking reader.

For information making heavy use of the facsimile mode, considerable telecommunications bandwidth will be required; the almost unlimited cheap bandwidth promised by very speculative technologies such as the recently proposed "space mirror"[2] may be necessary to make transmission of such data on a broad scale economically reasonable. It is worth remembering that Bell Telephone demonstrated the Picturephone at the 1964 World's Fair in New York, but even today teleconferencing is a costly and rather rare practice. Economics can greatly inhibit the spread of technology. If this bandwidth is not forthcoming, it is likely that current trends toward eliminating graphics from publications wherever possible, because of their high costs, will continue at least for the near future.

REMOTE PUBLIC ACCESS TO ONLINE CATALOGS

One policy issue concerns how to accommodate the increasingly common home/office terminal or personal computer that can be used to search the online catalog. To date there has been almost no experience with uncontrolled public access to computer utilities (if indeed uncontrolled access is desired). The problem here is to provide as much access as possible while still permitting fair and wide distribution of limited resources among as many people as possible.

Note that as the services available become more extensive, the typical session length increases. The CLR study has shown that users spend more time with an online catalog than with a card catalog. When documents become available, a typical session may tie up a terminal access path for hours.

Today the resource-allocation problem is solved by one of two methods: one is to provide a fixed amount of resources and let natural rationing occur through contention (similar to the approach used in allocating resources in many academic computing environments). The other is to ignore the problem and hope it will go away. Unfortunately, it will not go away: the public demand for these services is growing in direct proportion to the proliferation of terminal devices.

One possible distribution mechanism for online catalog services is videotext, which uses the home television set as a very common and inexpensive terminal over a split-speed telecommunications link (data is sent from the home at a very low rate and returned to the TV set at a much higher rate). Unfortunately, the type of tree-structured, menu-retrieval approaches found in today's prototype videotex systems are ill-suited to systems such as online catalogs. Similarly, the low-resolution TV screen is poor even at displaying text—a TV screen of text holds much less information than a typical terminal screen, and a TV screen is probably hopeless for reproduction of high-resolution graphics. Barring a major redesign of television into a high-resolution device, we believe that these features will limit the use of videotex systems to a small number of rather specialized applications. Online catalog searching and electronic document distribution will not be among them. These limitations may in fact ultimately cause the consumer marketplace to reject videotex in favor of a display system that is more costly (at least in terms of initial equipment investment), but more flexible.

Alternatively, cheap microelectronics may allow the traditional TV set to be extended. There are two barriers to high-resolution broadcast television: electromagnetic spectrum allocation and the receiver. The spectrum is becoming one of man's most limited resources, and doubling the space required for a television channel (from 3 to 6 MHz) is probably unacceptable. There are no serious technological barriers to a high-resolution TV receiver; only questions of how cheaply they can be mass-produced. We might see high-resolution TV receivers that will function directly as computer-communications displays in high-resolution mode, that will present high-resolution pictures from videodisks, and that will use a microprocessor-driven interpolation procedure to simulate high-resolution images from lower-precision broadcast displays.

THE BIBLIOGRAPHIC INTERNET I: CIRCULATION SYSTEMS

Providing links between online catalogs and circulation systems is an obvious extension of today's online catalog technology. To date, there have been two major impediments: lack of standards and protocols (and in fact a clear definition of the problem to be solved), and a reluctance on the part of circulation-system vendors to accept the fact that they must operate in a service environment containing numerous computers for different purposes. Until now, vendors of library automation services and systems have taken a very defensive role and have discouraged all types of interconnections rather than fostered them. The operational philosophy seems to be that once a library automation system is installed for any purpose, the library becomes a territory of that vendor, and is thus making a commitment to wait (forever, if need be) for that vendor to make available the full spectrum of library applications. In the last five years, with the general acceptance of the ISO models for open systems interconnection, and the spread of plug-compatible vendors in the IBM marketplace, the computer industry seems, at least in principle, to have embraced the concept of cooperation rather than territorial defense. It is time for the library marketplace to take the same steps.

Automated circulation systems are increasingly common. These are typically minicomputer-based. In general, we do not believe that circulation and online catalog functions belong within the same system; circulation is ultimately a very local function, whereas online catalogs benefit from a union or regional approach. The requirements for storage capacity, processing power, and reliability are very different.

If one assumes separate catalog and circulation systems, this means that applications-level protocols must be perfected to connect online catalogs (or, in fact, arbitrary computer systems) to circulation systems. Given the widespread use of minicomputer-based systems from vendors such as CLSI and GEAC, and the increasing use of online catalogs, the development of such protocols and their subsequent implementation by vendors should be a high priority for the library community. At first examination, such protocols may appear relatively simple—the functions are essentially to inquire on the status of an item, and perhaps recall or hold the item once it is identified. The full protocol is complex, however. There are issues of uniquely identifying the item, security issues such as identifying the inquirer/holder/recaller, and network issues. The network issues are particularly vexing: many circulation systems today can be networked, and several nodes on the circulation network may be connected to the online catalog (or online catalog network; see below). This is the first example we know of the need to address the full scope of internetwork issues at the level of applications protocol (other than perhaps mail systems) and is an area of active research in computer science.

THE BIBLIOGRAPHIC INTERNET II: LINKING ONLINE CATALOGS

Finally, there is the issue of interconnecting online catalogs themselves. This is a very complex area and one in which we believe a great deal of misdirected energy has been spent in the last few years. Actually there are two separate issues: intranational and international networking. International (or, more explicitly, multilingual) catalog networking is a largely unexamined issue from the technical point of view, although Fred Kilgour recently gave a very nice summary of the political and economic questions involved.[3] We will first examine intranational interconnection, and then explore the new problems that arise when this is extended into the international environment.

The general model of catalog interconnection as we see it is that a user will transact primarily with a local online catalog system (local in the logical sense that it contains material for the user's "home" library or library union, not necessarily that it runs on a computer housed in the user's local library). On occasion the user will wish to search remote (e.g., distant) online catalogs. On these occasions the local catalog will open a remote *computer-to-computer* connection to the remote catalog, controlled by a program on the local catalog that will mediate access on behalf of the user. This is as opposed to the model wherein the user is connected *directly* (perhaps through some form of telecommunications switching or pass-through) to the remote catalog and conducts searches directly in the command language of the remote catalog. In our model the entire matter of standardized command languages for online catalogs becomes irrelevant, since it is the responsibility of the local catalog to translate its local command language into some generally accepted applications-level protocol that is understood by the remote catalog.

This conceptual model has great implications for the design of applications-level protocols for interconnecting online catalogs. In fact, what is needed is a more general protocol for interconnecting all types of information retrieval systems, since we have already suggested that online abstracting and indexing systems will be integrated into the online catalog. Also, as we shall see later, other types of information retrieval systems may find their way into the mix. The approach used to date in designing applications-level protocols for online catalog interconnection (as seen, for example, in Norway's library network or the CLR linked systems project) has been to try to codify and represent every possible transaction against the catalog database in the online catalog protocol. We believe that this is a limited and inflexible approach ultimately doomed to failure as large numbers of online catalogs and other information retrieval systems are linked. What is really needed is a sort of metalanguage protocol whereby the information retrieval system can describe its capabilities and functions to a remote system so that the mediating program can map local commands into the capabilities of the remote system where possible. For bibliographic data this is relatively simple, because a group of generally accepted static data standards (the MARC standards) are already widely accepted. The approach we propose will lead to rather complex mediating programs that must apply artificial intelligence techniques to remote systems, but it provides the generality necessary to permit almost unlimited growth of networked information systems.

The international linkage of online catalogs presents additional problems because of the necessity of linguistic translation in the search keys and in the supporting authority control structure (to say nothing of the more basic problems of international computer networking, such as incompatible data communications standards and legal/political constraints on transborder data flow). To a certain extent, it is impossible to allow a user to search a foreign-language database in another language—particularly by subject—since concepts in one language may have no precise analog in another language. To solve the multilingual problem (as far as is possible without a case-by-case intellectual analysis), it may well be possible probabilistically to create cross-linguistic concordances in authority files by computer matching—finding items common to both databases and then linking the authority control entities if they apply to many of the same items. Canadian libraries have considerable experience in dealing with multilingual databases, and their efforts deserve closer scrutiny.

The economic issues in this model of catalog interconnection cannot be overlooked. There are several possible scenarios: one is simply to ignore any form of cost recharging, since it is obviously a very complex issue, and to declare that all members of the catalog internet will exchange services with other members at no cost. Another possibility would be to treat the record of each library's holdings as one of the library's major assets and charge other libraries for searching it. Should this occur, the entire matter could be carried out electronically through a system of credits and debits, and a form of electronic scrip would come into being as each library examined its balance of payments at the close of every billing period. Unfortunately, just as some nations must import almost everything, producing a negative balance of payments, the smaller libraries may spend substantial time searching the holdings of the major research libraries.

AUTOMATING THE REFERENCE LIBRARIAN

Often the kind of information provided by today's online catalogs really fails to address the patron's question. Up till now the reference librarian has been viewed as the source of these answers. With today's information deluge, it may be unreasonable to expect the reference librarian to continue this function without some automated assistance. This breaks new ground in library automation; for the first time we are looking at systems to supplement and support human beings rather than simply to mechanize well-defined manual activities.

One demand for answers that today's online catalogs cannot supply comes from questions like, "I need a good book on the Fall of Rome," or "What is the definitive biography of Lord Byron?" In a very real sense, the reference librarian can't answer these questions either; the patron is looking for opinion, or perhaps a summary or consensus of opinions.

Interestingly enough, these are not particularly difficult to provide, at least in subject areas where professional societies compile bibliographies and reading lists, or publish review journals such as *Computing Reviews* or *Mathematical Reviews*. These reviews could be made available as electronic documents over computer networks, allowing users to examine them before selecting books or articles. Similarly, annotated bibliographies find a natural place in an online catalog as electronic documents for browsing. Such aids would assist not only the patron but the reference librarian, who could both exploit them directly and guide patrons in their use.

Reference departments in major research libraries receive many queries each day, often by telephone. It will be many years (if ever) before such queries can be dealt with by machine, but we believe that analysis would indicate that many of them could be resolved simply by putting certain reference books (such as dictionaries and encyclopedias) online with high-quality user interfaces.

Beyond these, the evolving technology of expert systems and knowledge engineering offers interesting prospects.[4] Computer programs have been developed to try to emulate expert behavior in certain areas—for example, medical diagnosis, computer problem diagnosis and repair, and VLSI layout. These efforts are starting to show a rather respectable degree of success. It would be very interesting to try to develop an expert system that could answer (or at least direct users to material containing answers for) queries typically made to reference librarians.

Some sense of perspective is required here, however. While it would be a worthwhile research project to try to build such an expert system, and while the results might be useful, it is hard to believe that any computer program could match the reference librarian who, when approached with a request for "books about Africa," finally determines that the patron really wants a picture of an elephant.

One can classify public access library automation systems into a partial hierarchy of difficulty based on what they provide to the patron as follows:

1. Places to look for an answer (online catalog).
2. Places to look for an answer with evaluation (extended online catalogs that integrate bibliographies and reviews).
3. Material containing an answer (extended online catalog with document delivery/electronic publishing).
4. Answers.
5. Answers to poorly posed questions.

Library automation is still dealing primarily with level 1 today. Levels 2 and 3 are technologically within our grasp. As they are achieved they will create major economic and sociological changes as well as changes in quality of service. Levels 4 and 5 remain, at least for the present, fertile fields for speculation and research rather than realities.

CONCLUSIONS

Online catalogs have given the information community the beginnings of a powerful tool for finding and delivering information. Among the principal consequences have been new perspectives on the cataloging process as a result of keyword access by the catalog user. Another is that no longer might it be useful to think in terms of distinguishing cataloging of monographs from indexing of journal articles for the purpose of searching catalogs. Perhaps the most radical consequence of online cataloging is the potential for change in the locus of cataloging, first from original cataloging by individual libraries, then to the system in which libraries copy cataloging, finally toward a future potential for publishers taking over much of the cataloging load as a form of advertising. Additionally, the locus of library use is leaking out of libraries and into homes and offices. All of these changes are redefining the roles and functions of libraries as we know them today. As this process continues, and one no longer "goes to the library," the user may come to perceive a small number of information utilities instead of many specific libraries.

REFERENCES

1. Division of Library Automation and Library Research and Analysis Group, *Results of a National Survey of Users and Non-Users of Online Public Access Catalogs*, Final report to the Council on Library Resources (Berkeley, Calif.: Univ. of California, Office of the Assistant Vice President—Library Plans and Policies, Nov. 1982), 186p.
2. Anne Field, "The Space Mirror," *Forbes* 130:189–90 (Oct. 25, 1982).
3. Frederick G. Kilgour, "Problems Related to Interlinking Online Interactive Library Networks" (Paper presented at International Federation of Library Associations (IFLA) general conference, 1982), 10p.
4. William B. Gevarter, *An Overview of Expert Systems* (Washington, D.C.: Department of Commerce, National Bureau of Standards, #NBSIR 82-2502, May 1982). (NTIS Document Report #PB82-227547.) ■■

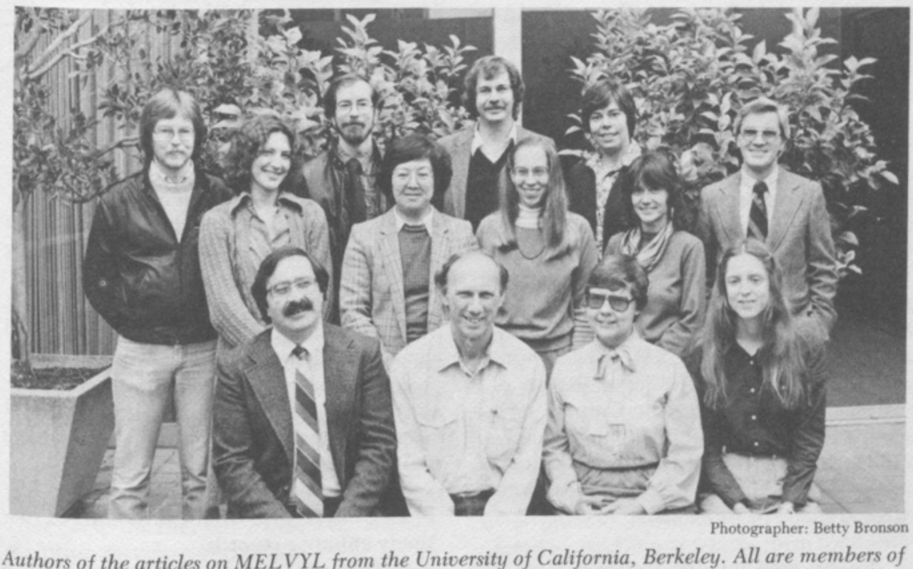

Photographer: Betty Bronson

Authors of the articles on MELVYL from the University of California, Berkeley. All are members of the Division of Library Automation (DLA) staff, except for Stephen Salmon, who heads Library Plans and Policies. From left rear: Ray Larson, David Shaughnessy, Dr. Edwin Brownrigg (director), Kathy Klemperer, Stephen Salmon; middle row: Mary Engle, Barbara Radke, Teresa Montgomery, Karen Coyle; front row: Dr. Bennett Price, Michael Berger, Dorothy McPherson, Vicki Graham; not shown: Clifford Lynch and Mark Needleman.

Reports and Working Papers

Summary Recommendations from Subject Access Meeting

C. Lee Jones

In June 1982, the Council on Library Resources brought together twenty-three people to explore issues related to subject access and to recommend actions that could help improve subject access to monographic literature in bibliographic databases. The meeting was held in Dublin, Ohio, and is one initiative of the Council's Bibliographic Service Development Program. Participants (Alan Benenfeld, David Bishop, Charles Bourne, Brett Butler, Pauline Cochrane, Tamas Doszkocs, Douglas Ferguson, Jeff Griffith, Donald Hawkins, Judith Herschman, Neal Kaske, Tina Kass, Carol Mandel, Davis McCarn, Jim McDonald, William Mischo, Pat Molholt, Mary K. Pietris, Lucia Rather, Martin Runkle, Elaine Svenonius, Velma Veneziano, and Jennifer Younger) included researchers, policymakers, and practitioners representing national and research libraries, library schools, database vendors, index and abstract services, bibliographic service agencies (utilities), and other organizations. This report contains the recommendations made at that meeting. A full report is available from the Council (1785 Massachusetts Ave., NW, Washington, D.C. 20036): Subject Access: Report of a Meeting Sponsored by the Council on Library Resources, Dublin, Ohio, June 7–9, 1982. *Compiled and edited by Keith W. Russell. Washington, D.C.: The Council, Dec. 1982. $10 (prepaid).*

Recommendations flowing from the Subject Access Meeting are divided into three sets. Two sets are equivalent in importance, the distinction being that one set is labeled Short-Term and the other set Long-Term Projects/Issues. The distinction between short-term and long-term is arbitrary but agreed upon by all participants. Short-term projects are those that can be started and substantially completed within three years. Long-term projects are those that may be started in the near term but hold little prospect for substantial progress within three years and will require effort well beyond the three year period.

The recommendations in the third set are those of lower priority, but which may nonetheless lead to useful projects for enhancing subject access for the users of online public access catalogs and other bibliographic reference systems. It should be understood that the subject access strategies discussed in this document are meant to incorporate not only the catalogs now being introduced in libraries, but also are intended to cover other databases of bibliographic citations regardless of their source or location.

ASSUMPTIONS

All recommended projects and noted issues that were identified during the Dublin meetings assume several things. Among those assumptions are the following:

1. The *Library of Congress Subject Headings* will be the basis for the controlled vocabulary in online public access catalogs.
2. It is too early to look for, or seek, standardization of features (including subject access strategies) among the many online public access catalogs available.
3. Local libraries are likely to use LC-MARC "as is" and will be unwilling to routinely enhance records.
4. The optimum subject search tool is the online public access catalog equipped with sophisticated search capabilities including natural language and controlled vocabulary searches.

C. Lee Jones is program officer, Council on Library Resources, Washington, D.C.

5. Within ten years all online public access catalogs will be equipped with a global change capability that permits, for instance, a change of a single term in a controlling vocabulary to be reflected in every record using that term.

6. Less comprehensive online public access catalogs must have the capacity, when necessary, to search more comprehensive ones, through an appropriately designed procedure.

SHORT-TERM PROJECTS/ISSUES

1A. The Council on Library Resources should accept a leadership role in developing an effective means of communication among several sets of people involved in issues related to online public access catalogs generally, and subject access strategies specifically. The sets of people that might be brought together for a variety of reasons include systems designers, those responsible for the selection of online public access catalog systems, those interested in development of innovative subject access strategies, etc. The techniques that might be used to accomplish these goals include topic oriented conferences with printed summaries and prepared papers to key the discussions, support of efforts to evaluate various online public access catalog systems and products, identification of appropriate factors in the selection of systems, cost considerations, both developmental and operating, etc.

1B. Create and distribute in machine-readable form the *Library of Congress Subject Headings (LCSH)* in the LC-MARC authority format providing for *current* and *regular* updates. The new edition of the *LCSH* in machine-readable form has just recently been released and covers headings added through 1980 only. If other organizations were to help LC implement a subject authority file service with current and periodic updates, it would be necessary to rekey all of the 1981 and 1982 additions and changes to the file. Once that is done a subscription service could begin.

Items 1A and 1B were considered to be the absolutely top priority actions to be taken. In order to achieve maximum effect, they should both be started simultaneously. Online catalogs are just now developing as important library tools. The fact that there is much development going on now underscores the time-sensitive nature of these two recommendations.

2. Develop and establish a mechanism for a set of libraries to contribute new subject headings to the Library of Congress. For a larger set of libraries, establish a similar mechanism for reference librarians and catalogers to suggest see references for inclusion in *LCSH*. In both cases, there should be a rapid review and reporting mechanism for all suggestions. In the case of cross references, transaction logs from online public access catalogs should be used as a source of suggestions.

3. Funding agencies should continue to support transaction log analysis and subsequent studies of the way different categories of users use online public access catalogs and reference databases, including those provided by the commercial sector. Because of the fundamental nature of the results of such studies, efforts must be made to see that this information flows steadily to those developing new online access systems.

4. Edit the *LCSH* see-also structure so that true hierarchical relationships are made explicit and both broader terms and narrower terms can be distinguished and retrieved. Evaluate the LC subject headings themselves to determine if they can be rearranged and displayed hierarchically and whether such a change would be useful to users of online catalogs and catalogers. If such a display is not feasible, major improvements can and should be made in the syndetic structure of *LCSH*.

5. Encourage Forest Press to authorize the design of a machine-readable format for the *Dewey Decimal Classification Schedules*, including schedule, text and indexes, and to distribute, with periodic updates, for search and display only, the resulting database. One aspect of the project to create the format is the need to define what portion of the system should be converted to machine-readable form for public access purposes—the abridged edition or the full version—and, generally, how the database might be used for public access purposes.

6. Explore, evaluate and promote a va-

riety of ways to train people with different backgrounds and experience to use online public access catalogs and other online bibliographic reference systems. Attention will have to be paid to the fact that increasingly such access will be available outside the library environment, a circumstance which calls for different training and updating strategies.

LONG-TERM PROJECTS/ISSUES

1. Identify and evaluate ways to augment and enhance subject access in newly created bibliographic records. These strategies might include additional effort on the part of the Library of Congress as well as efforts contributed by other institutions and special interest groups. The techniques that should be explored include but are not limited to the following:

- provide additional *LCSH* headings in bibliographic records (they may be considered secondary headings and not generate printed cards).
- add terms from special theasauri and see that these enhanced records are made available through LC, though the terms would most likely be added by other institutions or special interest groups.
- add table-of-contents data to the content note field, or as unformatted material appended to the record. Explore how this data, along with other possible data like index information, might come in machine-readable form from cooperating publishers.
- identify special groups of material to receive special in-depth indexing, for example, reference material or conference proceedings.

Any alternative record enhancement strategy should be pursued only after assessing the cost/benefit to be expected from such enhancements.

2. Establish the utility of the *Library of Congress Classification Schedules* (with scope notes)—or edited portions thereof—in machine-readable form for users of online access catalogs and catalogers. If they prove useful, design a format for update, search and display and organize regular, periodic updates. It is expected that incorporating some portion of the schedules with

notes and indexes will provide a more powerful subject access infrastructure for online public access catalogs.

3. Design a basic and transferable strategy for moving from a user's language through the system's language for the retrieval of bibliographic records. A fundamental design criterion is that users need not be aware of the translation, some arguing for an indication of the nature of that translation so the user becomes more aware of the power of the system being used.

OTHER PROJECTS/ISSUES

There were several other projects recommended by the participants during the Subject Access Meeting, but these projects failed to achieve a high priority rating in the consensus process. The group did, however, find them sufficiently commanding that they survived a group discussion as viable and useful projects. One should take note of these in the context of what needs to be done to enhance and improve subject access for the user of online public access catalogs.

The following six projects are not represented in any special order.

1. Develop a way to preserve the provenance, over time, of each entry in the machine-readable *LCSH*. This means that a history of all headings changed or deleted from the file would be preserved and noted in the file itself.

2. Eliminate obsolete terminology and heading style from *LCSH*. It is possible that this specific objective would become part of any one of several other projects included under short- and long-term projects/issues. This work might be accomplished by agencies outside of LC but approved by LC.

3. Assess the impact, in terms of cost and service enhancement, of imposing consistency of form and language on *LCSH*.

4. Develop the basic strategy for integrating *LSCH* with other thesauri and with classification schedules and ways to switch among them all.

5. Prepare an analysis of the command and record structures and the indexes maintained in a wide variety of online public access catalogs.

6. Attempt to evaluate whether or not online public access catalogs need to be restricted to the holdings of one or a set of libraries. Could/should they contain the full MARC record set as well as other records, like one or more article citation or reference databases? What would the local impact be of providing access to records for which holdings information was not available? ■■

News and Announcements

British Columbia Library Network

The British Columbia Union Catalogue (BCUC) is proceeding with a six-month pilot implementation of an online computer system for its proposed British Columbia Library Network (BCLN), despite its failure to obtain a special start-up grant during a time of cutbacks in provincial government spending. Approvals to proceed were obtained through the Management Advisory Council of B.C. Colleges and Institutes and the Universities Council of B.C. via the re-allocation of funds granted to the B.C. Union Catalogue Office for 1982–83, and from a special grant made to the three B.C. university libraries for the retrospective conversion of their catalog records. In order to minimize the initial cash requirement, staff and data communications facilities are being volunteered by the three B.C. university libraries.

It is planned that five BCUC member libraries that are participating in phase one of implementation will be using the new system for cataloging support functions by March 31, 1983. The five participating libraries consist of the following: the British Columbia Institute of Technology, Simon Fraser University, the University of British Columbia, the University of Victoria, and Vancouver Community College. It is expected that other B.C. libraries will implement the BCLN system when phase one has been successfully completed.

The system chosen for implementation is BLIS, a turnkey version of the Washington Library Network System marketed by Biblio-Techniques, Inc., of Olympia, Washington. R. W. MacDonald, assistant librarian, University of British Columbia, has been appointed BCLN project manager for phase-one implementation. BLIS and the BCLN database are being loaded onto an IBM 4341 computer at the University of Victoria Computing Centre. It is expected that BCLN-BLIS will be able to operate on a cost-recovery basis from the outset. User fees for services will be established during phase one, which will enable the project to continue on a "pay as you go" basis.

A number of interesting solutions are being accomplished in the building of the BCLN database and the design of the data communications network. The BCLN database is being built from bibliographic records supplied from the participating libraries' separate catalog files at UTLAS and authority records obtained from the Washington Library Network. These separate files are being merged into one shared file using software developed by the systems staff at the University of British Columbia Library. BCLN plans to retain online links to UTLAS and to establish online access to WLN as well. A user sitting at a BCLN terminal in his/her own library will be able to access the BCLN host, UTLAS, WLN, and other bibliographic databases such as CANOLE, DIALOG, BRS, etc. All of the access facilities are being provided through a data communications front-end processor developed on DEC equipment by the systems staff at the University of British Columbia Library. Both synchronous and asynchronous data communications will be supported. The BCLN system will be implemented over a three-month period after taking over three years to obtain approvals to proceed. ■■

Electronic Novel

The nation's first electronic novel was written in two and one-half days, published in three hours, and distributed within sixteen minutes each to a nationwide network of personal computer users.

The novel, *Blind Pharaoh*, was created by writer Burke Campbell at the ArtCulture Resource Centre in Toronto, Canada.

Campbell began writing on an Apple III personal computer, equipped with Apple Writer word processing software at 9:37 p.m. on November 14, 1982. He completed the novel 61½ hours later, at 11:07 a.m. on November 17.

As Campbell wrote, each chapter was proofread on a second Apple III equipped with a printer, and reformatted for electronic transmission by a representative from Apple Canada. Corrections to the manuscript were made the afternoon of November 17, and at 5:30 p.m. that day the novel was transmitted from Toronto to Source Telecomputing Corporation in McLean, Virginia.

The transmission occurred using Access III communications software and a Ven-Tel 1200 baud modem with the Apple III computer in Toronto. The novel was carried on the DataPac communications network in Canada and Tymnet in the United States.

In McLean, the novel was received chapter by chapter by members of the STC Editorial Department. Each chapter was proofread a second time, separated into page lengths, and assigned the appropriate chapter heading to enable subscribers to STC's service, THE SOURCE, to read the novel with ease. The entire nineteen chapters of *Blind Pharaoh* were successfully transmitted, formatted, and filed on THE SOURCE by 8:30 p.m. on November 17, at which point the novel could be received by any subscriber who typed the command NOVEL on the system.

When any of its 24,000 subscribers request the novel from THE SOURCE, they are given a menu selection that enables them to select any of the nineteen chapters for reading from their terminals; downloading onto a personal computer disk; or printing out onto a printer.

The novel, which is approximately 20,000 words in length, can be downloaded or printed in sixteen minutes using a 1200 baud modem. The only cost to subscribers for receiving the novel is their usage time on THE SOURCE, which costs $20.75 per-hour weekdays, $7.75 evenings and weekends, and $5.85 after midnight (prorated by the minute). If *Blind Pharaoh* were printed out at 1200 baud in the evening, the cost to a subscriber would be $2.03.

BLIND PHARAOH

TABLE OF CONTENTS

1. The Beginning
2. A Knock at the Door
3. Only A Miracle
4. In the Middle of Nowhere
5. Hard Times
6. The Missing Fact
7. Fiction
8. Lunch at Emilio's
9. The Dive
10. Crisis
11. The Crystal Fan
12. The End of the Day
13. In the Desert
14. The Toy
15. Possibilities
16. No Help Anywhere
17. In Case of Danger
18. The City by Night
19. Later

Enter item number or Help

Among the first readers of *Blind Pharaoh* over THE SOURCE was Henry Kisor, book editor of the *Chicago Sun-Times*. In his review of the novel, Mr. Kisor wrote:

> Writing a 20,000 word novel on a computer in 72 hours is a stunt, but Burke Campbell pulled it off handsomely. The result, "Blind Pharaoh," is rough-edged as one would expect, but its splendid narrative drive makes up for that. I enjoyed it very much.
>
> More important, perhaps, is its proof that videotex is a going concern as an alternative to the printed book. The technology is here already; all that is needed is a larger audience. And that will come; the only question is when. ■■

OCLC's Nine Millionth Record

The University of Denver input the nine millionth bibliographic record into OCLC's online database on November 29, 1982.

A team of seven terminal operators headed up by Elaine Henjum, catalog coordinator at the University of Denver's Penrose Library, captured OCLC control no. 9,000,000 at 7:45 p.m., eastern standard time. Was it an exciting moment for them? "To put it mildly," said Ms. Henjum. "Everyone started screaming."

The Ninth Solid Gold OCLC Record was *Inventory of Health Services in the Denver Metropolitan Area*, published by the Metropolitan Council for Community Service.

Members of the victorious University of Denver team were: Juanita Carter, library assistant IV; Lisa Cook, library assistant II; Carol John, library assistant IV; Robin Ottoson, library assistant IV: Barbara Sudol, library assistant IV; and Carolyn Wilmoth, library assistant II.

Each bibliographic record that enters OCLC's database is assigned an OCLC control number. When the OCLC number approaches the million mark, member libraries compete to see which institution will claim an OCLC Gold Record.

The steady growth of OCLC's online union catalog since it became operational on August 26, 1971, is evident from the dates that "Millionth Records" were input:

OCLC #9,000,000: November 29, 1982, University of Denver (BCR).
OCLC #8,000,000: December 16, 1981, Southern Illinois University (ILLINET).
OCLC #7,000,000: December 4, 1980, Northeastern Illinois University (ILLINET).
OCLC #6,000,000: February 19, 1980, Upstate Medical Center (SUNY).
OCLC #5,000,000: May 24, 1979, University of Louisville (SOLINET).
OCLC #4,000,000: June 24, 1978, University of North Carolina (SOLINET).
OCLC #3,000,000: May 27, 1977, SUNY-Potsdam (SUNY).
OCLC #2,000,000: February 18, 1976, Boston University School of Theology (NELINET).
OCLC #1,000,000: September 6, 1974, Northeastern University (NELINET). ■■

Automated Bibliographic System for Manuscripts and Archives

The U.S. Office of Education has awarded a Title II-C grant of $400,000 to Yale University to undertake a joint project with Cornell and Stanford University libraries, the Hoover Institution, and the Research Libraries Group, Inc. The purpose of this project is to design and implement

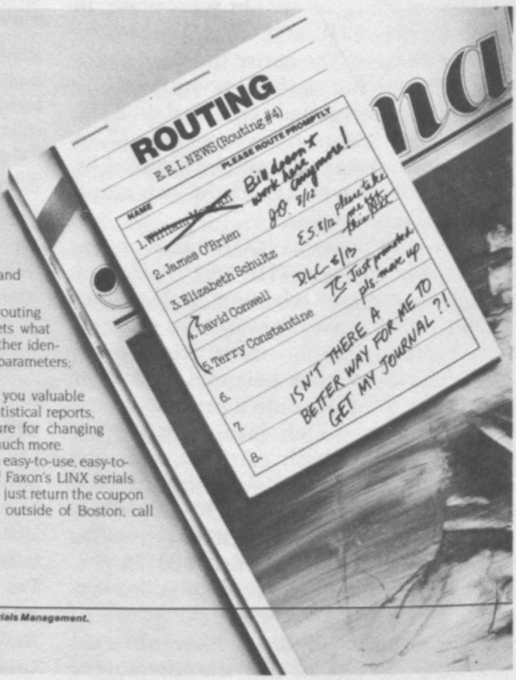

enhancements to the Research Libraries Information Network (RLIN), which will facilitate the acquisition, cataloging, and management of manuscript and archival materials. The development of a standard bibliographic exchange format will enable Cornell, Stanford, and Yale to integrate their manuscript and archival holdings into RLIN, thereby forming the foundation of a national database.

Project activities in 1983 will include determining cataloging standards; establishing guidelines for authority control; entering records into the RLIN database; and producing user documentation for dissemination to other RLG institutions. These activities build upon previous work done with Title II-C funding at Yale, which involved the development of functional specifications for enhancements to the RLIN bibliographic system. ■■

Tonkery Joins Faxon

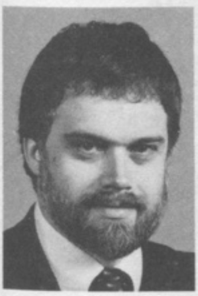

Faxon is pleased to announce the appointment of Dan Tonkery as vice-president and managing director for North America. Beginning January 1, 1983, he is responsible for directing the marketing, sales, and operations of Faxon's library services in North America.

Tonkery was associate university librarian at the University of California at Los Angeles, where he was responsible for Technical Services and Bibliographic Products, which include library automation. Prior to U.C.L.A., he was chief of Technical Services at the National Library of Medicine. He is nationally recognized as an expert in the field of library automation. ■■

ASCLA Offers Nonmember Subscriptions to *Interface*

The Association of Specialized and Cooperative Library Agencies (ASCLA), a division of the American Library Association, has announced that its quarterly publication, *Interface*, is now available to nonmembers on subscription. ASCLA members receive *Interface* as a benefit of their division membership.

Interface includes the latest developments and news in networks and cooperatives, state library agencies, and library services to special clientele. During the 1981–82 year, it featured topics as varied as the future of state libraries, information brokerage, a filmography for the International Year of Disabled Persons, ideas for innovative programs for the elderly, consumer information on teletypewriters for the deaf, and a series of articles on libraries and the new Federalism.

Individual issues include features, resources, news items, special columns and previews, and wrap-ups of ALA Midwinter Meeting and Annual Conference activities. *Interface* also reports on ASCLA activities and business.

To subscribe, send $10 for a one-year subscription (four issues) to *Interface* Subscriptions, ASCLA/ALA, 50 E. Huron St., Chicago, IL 60611. ■■

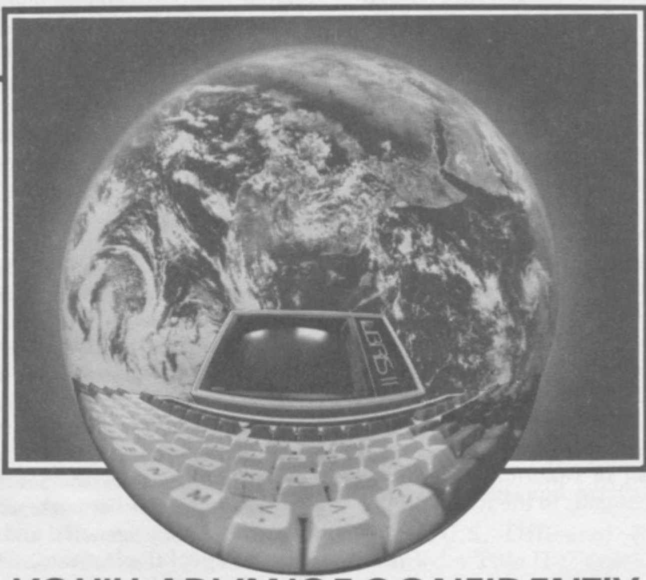

Recent Publications

Bibliographic citations were produced by María Clark, Yale University Library, New Haven, Connecticut, in accordance with the American National Standard for Bibliographic References. *New York: American National Standards Institute, 1977. 92p.* (American National Standards on Library Work and Documentation; *ANSI Z39.29-1977).*

Reviews

Microcomputers in libraries: option or necessity? [audiocassette recording]: Proceedings of the CLASS Workshop; 1982 May 20–21; San Jose, Calif., [n.p.:1982]. 10 cassettes. $50 CLASS members; $60 non-members; from CLASS, 1450 Koll Circle, San Jose, CA 95112.

If the number of articles appearing in library journals on the subject is any indication, the hottest topic in librarianship this year must be microcomputers. To help satisfy the demand for information on this topic, California Library Authority for Systems and Services (CLASS) has made available on audiocassette the proceedings of its workshop for librarians either involved with, or thinking about becoming involved with, microcomputers. Included with the ten cassettes are a conference program, glossary, and twenty-five-page bibliography.

The speakers, without exception, are enthusiastic advocates of microcomputers in the library and pass quickly over the question of whether they belong, to why, where, and how they can be utilized in libraries. They also caution the uninitiated about some of the pitfalls (hardware, software, and human) to beware of, and indulge in a little stargazing into the constantly changing universe of microcomputer technology. The speakers vary greatly in their knowledge of, and experience with, microcomputers, from those with highly technical backgrounds to others who acquired their experience concurrent with the acquisition of their library's first microcomputer.

Consultant Susan Baerg Epstein, Mark Rorvig (author of *Microcomputers and Libraries*), Ralph Shoffner of Ringgold Management Systems, Debbie Christian (editor of *Access: Microcomputers in Libraries*), and Dr. Koichiro Isshiki (author of *Small Business Computers: A Guide to Evaluation and Selection*) are just a few of those presenting papers. They and the other dozen or so speakers discuss such topics as: systems analysis and planning; selection of databases; public access microcomputers; system costs and methods of financing; and statistical applications.

Epstein's paper on systems analysis and planning focuses on the process of assessing the library's need for a microcomputer, including the identification of appropriate hardware and software. She discusses in some depth the preparation of a Request for Proposal (RFP). Isshiki also addresses the planning process that precedes library automation: how to decide what type and size of computer is needed; how to finance the system; and how to draw up a contractual agreement between the library and the vendor. Shoffner, on the other hand, puts little stock in the feasibility study as a tool for deciding whether or not to automate—he rejects the myth that such decisions are made rationally. The important questions, he says, in deciding whether to get a computer are: Do you want it? and Can you sell it to your administration?

Rorvig and others discuss the technical evolution of the microcomputer and the consequences of this evolution (such as less-expensive micros with greater technical ability). They also describe and compare brands and models. Additionally, commercial vendors make brief presentations about their library-oriented, micro-based products (among them, coin-operated microcomputers for public use).

A current problem facing libraries that want to use microcomputers is the scarcity of software written especially for them. In their papers, Robert Newhard of Torrance, California, Public Library, and Debbie Christian tell how off-the-shelf software designed primarily for business can be adapted to serve library purposes. For example, Newhard adapted Visicalc in order to use it for scheduling, cumbersome monthly statistical reports, and an online file for the library's frequently requested *Consumer Price Index*. Christian used a data management program designed for small business to develop a circulation system for a small public library in Oregon.

The adaptability of microcomputers to a variety of library settings was further demonstrated by personal accounts of librarians in environments as diverse as the children's and adult departments of a public library, a law library, and the corporate library of a manufacturing company. Microcomputers are being used in libraries for such functions as public access, cataloging, acquisitions, dial-up access to bibliographic databases, word processing, and general business functions such as payroll, mailing lists, and accounting as well as those already mentioned. There can be no doubt that microcomputers provide many options for improving library operations, so many, in fact, that one begins to wonder if they will soon be considered as essential to libraries as the librarian.

Overall, the speakers succeed in delivering relevant information that can readily be grasped by the nonspecialist. And the glossary provides a useful reference tool for the technical terminology that is essential for microcomputer literacy.

The format chosen for publication—audiocassette—has both positive and negative aspects. Technically, the sound is good, and, because most of the speakers are dynamic, there is a sense of participation in the workshop that would be absent with printed proceedings. On the negative side, many of the speakers used visual aids that were insufficiently explained to allow the cassette listener to follow these parts of the talks. A minor annoyance is the inaudibility of the questions from the audience, and few speakers thought to repeat questions into the microphone.

On the one hand, for an individual, listening slows down the intake of information. Since one cannot scan, the audio format becomes quite tedious at times. On the other hand, the audio format is well suited for a group training and/or discussion situation. Furthermore, since speed of publication is critical in such a rapidly developing field as microcomputers, publication via audiocassette allows the material to reach the public before it becomes obsolete.

The conference program partially serves as a table of contents to the tapes; however, because the audiocassette format makes indexing infeasible and scanning impossible, the user has little means of selecting from the whole those papers of prime interest to him; he is, thus, made captive to the entire set of tapes. The inclusion of printed abstracts of the papers would help to compensate for these limitations.

The workshop proceedings make available to a broad audience the contents of a well-rounded program that speaks directly and intelligibly to the current needs of librarians interested in microcomputer technology and its application to libraries. Is the set worth the price? Each library must assess its own needs in order to answer this question. For group use, the tapes are a good buy as they simulate the workshop situation, provide information to a number of people simultaneously, and facilitate group discussion; however, if only one or two staff members need the information, the cost is probably not justified.—*Faye Powell, Washington County Cooperative Library Services, Aloha, Oregon.* ■■

Videotex Canada. 1(2,3): 1982 Aug., Nov. "The national videotex magazine and directory." Issued quarterly by and available from: Tele-Direct, Box 8000, Postal Stn. "A," Toronto, Ontario, M5W 9Z9. Annually: Canada $25 (Can.), elsewhere $30 (Can.).

Every field needs its trade journal it seems, even the field of electronic information. *Videotex Canada* is off to a good start as the trade journal of the Canadian videotex industry. Pundits may muse on how long it will stay "in print," as opposed to becoming an electronically distributed product itself.

Now in its third issue (Nov. 1982), *Videotex Canada* is a four-color, glossy-covered journal, black and white inside, with an editorial, seven to ten articles, usually illustrated, news briefs, and a directory of Canadian TELIDON system suppliers in each issue. There is a fair bit of advertising by TELIDON hardware, software, and system suppliers, which adds to the journal's value.

Content takes a decidedly business-oriented viewpoint. Recent articles have covered specific TELIDON projects, the TELIDON Marketing Secretariat of the Government of Canada, the government's Industry Investment Stimulation Program, retail strategy using TELIDON, cost modeling of centralized, decentralized, and distributed architecture for large videotex systems, page creation, technical standards, and privacy issues.

For those interested in seriously investigating the current state and future of videotex using the TELIDON-developed North American standard, this journal is interesting and valuable reading.—*Brian Aveney, Blackwell North America, Inc., Lake Oswego, Oregon.* ■■

Other Recent Receipts

Listed here are books and other publications of potential interest to members of LITA, received for review. Some of these materials may be reviewed in later issues of ITAL.

Cuadra Associates, Inc. *A library and information science research agenda for the 1980's: summary report.* Santa Monica, Calif.: Cuadra Associates, Inc., 1982. 52p. Project conducted for the United States Department of Education, Office of Libraries and Learning Technologies, under contract number 300-81-0022.

Fenichel, Carol Hansen and Harter, Stephen Paul. *Survey of online searching instruction in schools of library and information science.* Dublin, Ohio: OCLC Development Division, 1981. vii, 100p. (Research report series). Bibliography: p.99–100. Report no.: OCLC/DD/RR-81/3. ERIC ED 210 038.

The future of the union catalogue. C. Donald Cook, ed. International Symposium on the Future of the Union Catalogue; May 21–22, 1981; University of Toronto. New York: Haworth Press; 1982. 130p. (Cataloging & classification quarterly; 2(1–2)). ISBN: 0-86656-175-7. $19.95.

Grosch, Audrey N. *The selection of cathode ray tube (CRT) terminals for library applications.* (Library technology reports. 1982 May–June; 18(3); 271–373). Chicago: American Library Assn. ISSN: 0024-2586.

IRCIHE bulletin. 7(1–4): 1981. CODEN: IRBUD5. ISSN: 0351-0123. "IRCIHE is created within the framework of the UNISIST Programme." Available from: International Referral Centre for Information Handling Equipment—IRCIHE, P.O.B. 327, 41001 Zagreb, Yugoslavia. Annual subscription: $24.

The microelectronics revolution: the complete guide to the new technology and its impact on society. Tom Forester, ed. Oxford: Basil Blackwell, 1980. xvii, 589p. ISBN: 0-631-11651-6, hardcover, £16.00; ISBN: 0-631-12572-8, softcover, £4.95.

Norman, Adrian R. D. (team leader, Arthur D. Little Ltd.). *Electronic document delivery: the ARTEMIS concept for document digitalisation and teletransmission.* White Plains, N.Y.; London: Knowledge Industry Publications, 1981. xiv, 226p. (Communications Library). "A study prepared for the Directorate-General Information Market and Innovation, Commission of the European Communities." Includes bibliographies. ISBN: 0-86729-011-0. $45.

Online micro-software guide & directory 1983–84. Nov. 1982. 346p. Accompanied by: Addendum. ISSN: 0734-5097. Available from: Online, Inc., 11 Tannery Lane, Weston, CT 06883. Base publication and addendum: $40. Annual supplements (Nov. 1983 & Nov. 1984): $30 for the pair.

Sager, Donald J. *The American public library.* Dublin, Ohio: OCLC Office of Research, June 15, 1982. vii, 49p. (Research report series). Report no.: OCLC/OPR/RR-82/1. ■■

Letters

To the Editor:

In your review of *Online Hotline* (Sept. '82 p.305), your reference to the "Apple Access III" software program should be corrected.

The software program supporting the *Online Hotline* is VIDACCESS™, a micro-publishing software program for the Apple II microcomputer which is distributed by Information Intelligence Inc. (A-C-C-E-S-S™ is also another software program for microcomputer networking distributed by III.) Apple Access III is a software program distributed by Apple Computer, Inc. for their Apple III microcomputer. (We have previously notified Apple Computer of prior use of the name.)—*Richard S. Huleatt, President & Publisher, Information Intelligence Incorporated, Phoenix, Arizona.* ■■

NOW…

The Concise AACR2

MICHAEL GORMAN

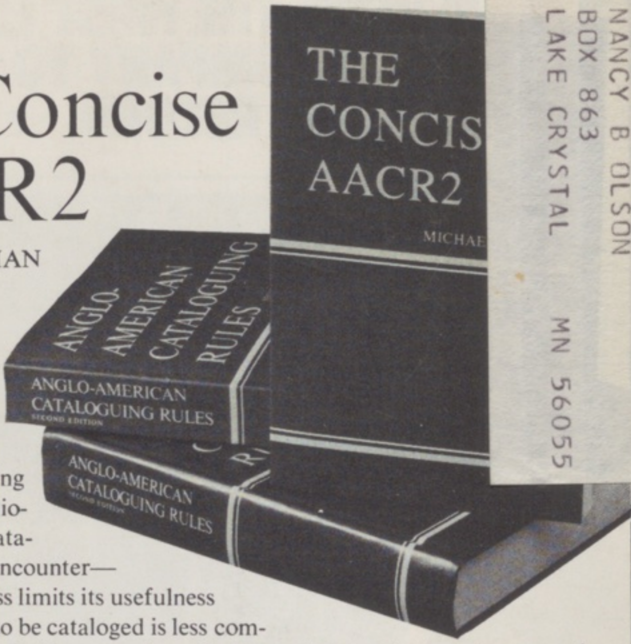

While the original *AACR2*, published in 1979, is a formidable work—covering as it does all the bibliographic situations catalogers are likely to encounter—its very thoroughness limits its usefulness where the material to be cataloged is less complicated or where only the operative principles of the code must be grasped. The latter is the case in teaching, for example.

Without disregarding the importance of properly cataloging the more difficult material, catalogers will recognize that most of their questions can be answered by a simpler version of the *Rules*. *The Concise AACR2* was designed for that purpose. It emphasizes essential principles, dropping separate treatment by medium and the less frequently used rules. In many cases it simplifies the wording of the rules while often giving additional explanation. *The Concise AACR2* should be consulted first when problems occur; it follows the rule numbers of *AACR2* so that if it cannot provide the answer the complete version may be used.

176 pages Paper ISBN 0-8389-0325-8 (1981) $6.50

***Anglo-American Cataloguing Rules*, 2nd ed.**
Cloth ISBN 0-8389-3210-X (1979) $15.00

***Anglo-American Cataloguing Rules*, 2nd ed.**
Paper ISBN 0-8389-3211-8 (1979) $12.00

Sample entries shown in typewriter type to guide users in the exact preparation of their own original copy. ▶

The Family of man . . . / created by Edward Steichen
(a collection of 503 photographs by various people)
Main entry under title
Added entry under the heading for Steichen

Why a duck? : visual and verbal gems from the Marx Brothers movies / edited by Richard J. Anobile
Main entry under title
Added entry under the heading for Anobile

The New York times a

Order Department
American Library Association 50 East Huron St., Chicago, IL 60611