THE ONLINE CATALOGUE AT THE UNIVERSITY OF GUELPH LIBRARY: THE PROPOSED SYSTEM ARCHITECTURE

Michael Ridley
Systems Division, Library
University of Guelph
Guelph, Ontario N1G 2W1

ABSTRACT

The online catalogue at the University of Guelph Library incorporates a system architecture based on the use of microcomputers, new database hardware, and the principle of distributed processing. The resulting configuration attempts to address the key problems in online catalogue design and provide a flexible and powerful tool for the library's user community.

LE CATALOGUE EN-LIGNE A LA BIBLIOTHEQUE DE L'UNIVERSITE DE GUELPH: L'ARCHITECTURE SYSTEMIQUE PROPOSEE

RESUME

Le catalogue en-ligne de la bibliothèque de l'Université de Guelph possède une architecture systémique basée sur l'utilisation de microordinateurs, sur un nouvel équipement de bases de données et sur le principe du traitement décentralisé. La configuration proposée pour ce système tient compte des principaux problèmes ordinairement rattachés aux catalogues en-ligne. On espère ainsi offrir un puissant outil de travail à la clientèle de la bibliothèque.

THE ONLINE CATALOGUE AT THE UNIVERSITY OF GUELPH LIBRARY: THE PROPOSED SYSTEM ARCHITECTURE

Michael Ridley
Systems Division
University of Guelph Library

INTRODUCTION

The development of online public access catalogues represents the lastest, and possibly the most important, application of automation to library services. The online catalogue will radically alter the interregation of library holdings, freeing the user from much of the baggage of past practice which has threatened to impede library use.

The University of Guelph Library has provided a public access system since 1977. A recent user study confirmed the overwhelming success of the system inspite of its lack of subject access (Pawley, 1982). When the library began development of a full public access online catalogue many critical design and development decisions had to be made.

In order to ensure that the new online catalogue responded to the real needs of the user community, input was gathered from a wide range of library staff and clientele. This involvement concentrated on solving traditional problems in catalogue use and proposing innovative solutions. Extensive use was made of dynamic prototyping tools which permitted effective and rapid establishment of the system's user interface as well as the system's operational methods.

This paper discusses the system configuration upon which the online catalogue will be implemented. Much of the discussion surrounding online catalogue development has centered on the user interface or the procedural design. There is, however, a major revolution in computing hardware and system design which should be applied to the online catalogue environment. The University of Guelph Library, in creating its system, was anxious to utilize the best possible techniques and adopt the most beneficial systems. The architecture described represents a departure from traditional designs, but it reflects the best current thought in the use of database management systems (DBMS).

ONLINE CATALOGUES

Online catalogues pose new problems for libraries in terms of system support and performance. With the online catalogue, libraries are progressing beyond the rather primitive information systems characterized by current circulation

packages, and are moving towards a sophisticated and complex system capable of responding to the full range of information retrieval inquiries. The nature of online catalogues can be characterized and understood by examining their standard operational features. The most prominent of these are:

1) heavy transaction loads
2) very large databases
3) fast response time requirement

The short history of system design for online catalogues is dominated by one particular configuration: a single general purpose CPU running the applications code which comprises the system and mediates between user terminals and the disc files. A variation on this structure utilizes a pool of large mainframes operating on the same principle. Quite often the machine is a mainframe serving as the principle computer resource for a large organization and its many applications. This star configuration suffers from a number of severe restrictions. Because the online catalogue is simply one of many application programs run by the system, it suffers from having to share system resources. Many other factors can combine to reduce the catalogue performance and thereby erratically effect its capability. The single data access path does nothing to alleviate the relentless I/O generated by the catalogue. The typical result is that the library is forced to support or buy a very large and powerful processor to accomplish the rather primitive data access routines required by the system. It is an extremely expensive option, and one which misplaces resources. It is apparent that the structures and procedures implemented for data processing are not adequate for the essentially different demands of information processing.

The University of Guelph Library was instrumental in the development of library systems based on cheaper, library-controlled minicomputer technology. While this permits the dedication of computer resources specifically to library applications and places the technology under library management, it does not fully address the growing traffic patterns and the need to support increasingly complex searching, particularly Boolean operations. The solution required a break with traditional library systems design.

THE SYSTEM ARCHITECTURE

The architecture devised for the catalogue is a distributed system utilizing microcomputer hosts linked to specialized database processing hardware operating a relational database (figure 1).

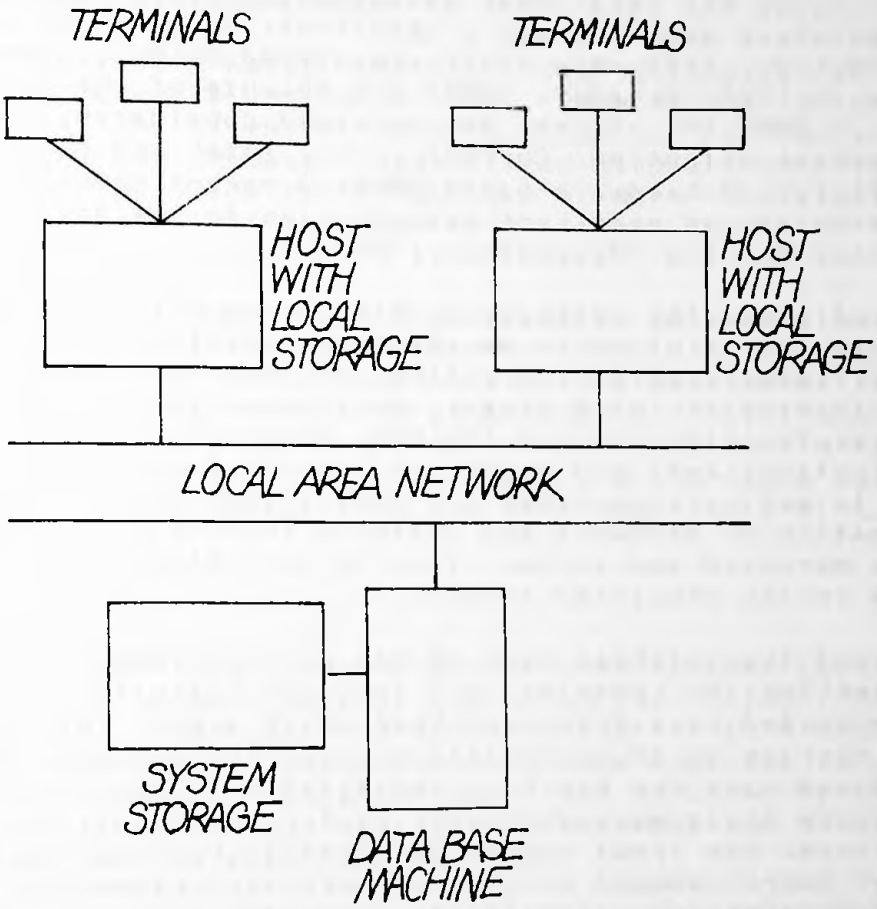

Figure 1: The System Configuration

There are a number of obvious advantages to this particular approach:

1) distributed functions allowing multiprocessing and specialization of operations
2) permits horizontal integration in the case of system expansion
3) increases data access paths relieving I/O bottlenecks
4) component/modular approach allows for increased reliability and system availability
5) cost

1) The Back End

The system is centred around a database machine (DBM) which maintains the relational database and facilitates data access. Database machines are a relatively recent innovation in DBMS technology. They were first prototyped by R.E. Canaday of Bell Labs in 1974 (Canaday, 1974) and because of the large interest in DBMS the concept has received considerable research and commerical attention. Currently only Intel and Britton Lee are manufacturing hardware-based DBMs. A recent Computer Surveys article provides an excellent introduction to the development of the database machine (Maryanski, 1980).

Essentially, the database machine is an optimized hardware and/or software solution to database manipulation. The DBM is exclusively dedicated to the control and management of the data base, performing all data access, updating, storage management and similar functions. Since the DBM, or back end, is free of all applications code and resultant responsibilities, it can be designed to optimally perform its restricted duties. This specialization of hardware and software reduces much of the machine's overhead and allows it to be very efficient and effective in its specified tasks.

Because the database machine has no involvement with running application programs, all terminal activity and access must be provided by a front end host which simply instructs the database machine in an appropriate manner. The dialogue between the front and back end has been optimized with the use of macros. Since these macros specify particular functions the DBM is to perform, the front end is required to pass to the DBM only the proper macro command and the appropriate arguments. The front end handles I/O, formatting, query and response etc. while the back end, the database machine, concentrates on the problem of data access. An extremely valuable aspect of this construction is that the isolation of the data from the application code allows for independent alteration of the database.

The use of a back end database processor responds to the typical library environment which places considerable importance on data access. The creation of specialized software and hardware optimizes this function and relieves the back end of other duties. Such dedication of purpose is appropriate in an application where disc access will predictably continue as the most prominent feature of the system use.

The relational database model supported by the majority of database machine manufacturers has additional advantages in an information retrieval application. One principle feature is that it permits data independence and flexibility of data in a

simple and well structured context. It does not, however, rigidly encase the database, making it is very easy to extract data from multiple relations without the need to maintain elaborate and mystifying pointer structures or by developing special applications code.

It has been noted that relational databases can suffer from poor response (Bridges, 1982). Some of the software implementations of the database machine are particularly prone to this problem. However, specialized hardware and firmware in conjunction with the restrained use of more traditional pointer structured files can overcome this concern and allow the relational model to operate effectively.

Certain trade-offs are apparent in this configuration. Since the database machine is not a general purpose CPU, the very node specialization which permits speed and efficiency also disallows the load-balancing advantages of traditional multi-processing structures. One very significant disadvantage is the fact that the product is new, and therefore largely untested, in the marketplace and in production applications.

2) The Front End

The front end system runs the applications code and provides the general purpose computing facilities typical of most systems. The front end can itself exist in a number of different formats, ranging from a single computer to a large number of computers linked by some form of network. The configuration described here anticipates the use of multiple microcomputers as front end devices.

The terminals, both in public areas and in restricted technical services areas, are clustered around a number of similar microcomputers acting as front end hosts for the database machine. The various application programs which constitute the online system run on these computers and pass data and instructions between the terminal and the database machine. The various hosts have local disc storage and can communicate with each other in addition to the database. The hosts will run the code such as the screen drivers, back end query processing and host-DBM link. An obvious performance feature is the multi-processor environment which allows simultaneous operations. The system will support a larger number of terminals without the resultant reduction in system response time. Additional flexibility is gained since there is no restriction on the type of terminal which can be attached to the front end. In fact, different applications within the library require different terminal characteristics and capabilities. Being able to utilize the most cost-effective terminal for the particular application is a considerable advantage.

Since the local storage of the front end contains the system code, the micro is virtually completely isolated from the back end machine, except for the function of data base accesses. Typically the user-system interface places a substantial demand on a system. Having the front end accomplish this permits two principle features: 1) allows a more complex interface to be maintained without fear of adverse system impact, and 2) it logically splits resource intensive operations over multiple machines. At the times when the host is able to satisfy users with local resources, the database machine is free to service other nodes. Since, in a library environment, a substantial amount of I/O is generated before the data access request is actually made, it is functionally appropriate to divide the workload amongst processors in this manner.

When the manufacturers of database machines began to market their product, it was expected that the front end would be some form of mainframe or minicomputer. For many applications this configuration would be practical, but in a library environment the use of multiple microcomputer hosts has obvious benefits. Library users do not need the processing power of a general purpose minicomputer. The micro will adequately perform the relatively restricted applications code needed to support an online catalogue. In addition, system reliability can be increased since it is cost effective to maintain spare CPUs in the event of a production host failure. Being able to simply and immediately replace the host maximizes the time of system availability.

3) The Network

The host and the back end are connected by a local area network, allowing communication amongst the components of the system. As such both the host and the back end must have communication code enabling the proper protocols and the exchange of data. It is another trade-off, therefore, that the overhead caused by the communications software should not outweigh the advantages of the multiprocessing environment.

From a management perspective, there are two overriding considerations and advantages: 1) flexibility and 2) cost. Because the system is network-based, and not a monolithic unit, it will respond to changes in organizational needs. The addition of word processing and electronic mail facilities for library managers becomes an easily mountable applications program. By utilizing low cost microcomputer hardware and specialized processors, the overall system cost can be held to a minimum. The processing requirements can be tailored to the individual organizational situation.

A considerable advantage of this configuration is the ability to horizontally expand the system. If additional terminals affect response time, then an additional front end can

be easily attached to the system to handle the new terminals. Given the low cost of the individual hosts, this represents a cheap means of system upgrading, thereby maintaining the performance edge afforded by new and faster processors. Just as the front ends can be multiplied, it is also possible to attach multiple database machines to the system. A second or third machine could be accessed via the local area network and thereby introduce certain disc storage advantages. The database could be logically split over the various DBMs or the database could be duplicated, providing a redundant, highly secure configuration. The advantages include multiple paths to the data, simultaneous processing and redundant processing in the event of a back end hardware failure.

CONCLUSION

Because online catalogues are still very much in a state of evolution and will continue to be subject to the rapid changes typical of computing, it is mandatory that any system configuration be able to adjust itself easily to new conditions and new user demands. The architecture described here attempts to respond to these priorities and addresses the problems typical of current and projected online catalogues.

There appears to be an increasingly popular philosophy amongst librarians that to be on the "leading-edge" is to be on the "bleeding-edge", and that we, as an organization, should wait for innovation and most assuredly not anticipate innovation (Graham 1983). At the University of Guelph Library just the opposite views are held. Libraries must assume control of their own destinies, providing the innovation and the direction which will ensure effective service to the user community. We hope that the online catalogue at Guelph will reflect our concerns and permit our users access to a new era in information retrieval.

REFERENCES

BRIDGES, Terry. "Data Base Machines - What and Why", in Data Management. Vol. 20, No. 11 (Nov. 1982), pp. 14-16.

CANADAY, R.E. et al. "A Back-end Computer for Data Base Management", in Communications of the ACM. Vol. 17 No. 10 (Oct. 1974), pp. 575-582.

GRAHAM, Peter S. "Technology and the Online Catalog", in Library Resources and Technical Services. Vol. 27, No. 1 (Jan./March 1983), pp. 18-35.

MARYANSKI, Fred J. "Backend Database Systems", in Computer Surveys. Vol. 12 No. 1 (March 1980), pp. 3-25.

PAWLEY, Carolyn. "Online Access: User Reaction", in College and Research Libraries. Vol. 43 No. 6 (Nov. 1982), pp. 473-477.