ASSOCIATIVE PROCESSING
A CHALLENGE FOR INFORMATION SCIENTISTS

Ernst J. Schuegraf
Department of Mathematics
St. Francis Xavier University
Antigonish, Nova Scotia
B2G 1CO

ABSTRACT

The concept of associative processing as a special processing method suitable for character handling tasks is introduced. A short survey of the history of associative processing is given together with a description of the basic machine structure. Differences between conventional and associative processors are mentioned. Problems of implementing these processors are outlined and the possibility of an associative file store is discussed. The challenge of associative processing to information scientists is illustrated on the example of indexing.

REPRODUCTION ASSOCIATIVE
UN DEFI POUR LE SCIENTISTE EN INFORMATION

RESUME

Cette étude introduit le concept de reproduction associative comme méthode pour la reproduction ayant rapport au relevé de certains caractères. L'histoire derrière cette méthode de reproduction est étudiée et une description de la structure de la machine est présentée. Les différences entre les méthodes de reproduction conventionnelle et les méthodes de reproduction associative sont énoncées. Les problèmes associés avec l'implantation de cette méthode de reproduction est discuté et la possibilité de créer un centre de dossiers d'association est envisagée. Le défi qui est lancé aux scientistes en information par cette méthode de reproduction associative, est illustré avec un exemple d'indexation.

ASSOCIATIVE PROCESSING

INTRODUCTION

Many computer based information systems operate with just the bare minimum of resources in order to keep costs as low as possible. However, we all know that systems managers always ask for more: computing power, memory, and peripherals. They not only ask for more, but for faster and more expensive equipment. Management invariably asks: How much does it cost; how cost-effective will this new equipment be? Systems managers normally don't like to think about these aspects.

A computer which is installed mainly for the support of an information system can have quite different hardware from that found in a general purpose computing environment. It needs good character handling capabilities, efficient file access methods, good input/output facilities, but very little special hardware for doing "fancy" arithmetic. Large computers for commercial applications generally have these properties, and the main chores for such a computer are those of sorting and searching (Knuth, 1974). Examples from the library environment are the preparation of a KWIC index which involves a sort on the keywords, or the search of a file for all the records which contain a given keyword.

It has been suggested that a considerable increase in performance can be expected with computers which utilize parallel processing (Carroll, 1976). However, even if a parallel processor is used in information system support, the specialized tasks of such a system are still carried out by a computer designed for a variety of applications. Would it not be better to design a new computer just to handle these specialized operations more efficiently?

Fortunately, it is not necessary to design such a new computer, since the basic concept of a suitable machine, namely, the idea of "associative processing," has been around for 23 years. It had been described the first time in 1956 (Slade, McMahon, 1956). Machines based on that concept are called "associative processors." The main difference between a conventional and an associative processor is the method of accessing memory. Associative memory is accessed by content rather than by an address as in a conventional coordinate addressed memory. In addition to content addressing, a group, or all words, of memory are accessed in parallel, a fact which is responsible for the name of "associative parallel processor" (APP).

In its beginning, associative processing was closely connected with cryotron technology, but it became quickly a popular subject for research. The suitability of associative processing for information retrieval applications was recognized very early (Petersen, 1962; Reich, 1969). Special significance was a paper, published in 1962, by Lee and Paull (Lee, Paull, 1962), which carries the title: "Intercommunicating Cells--Basis for a Distributed Logic Computer." They suggested a new type of associative processor, which, since then, has been named a "distributed logic associative processor" (Yau, Fung, 1977). In this processor each word in the associative memory array has comparison logic attached to it. Many variations to this design have been proposed, but the basic structure remains.

ASSOCIATIVE PROCESSING

To understand the operation of an APP, its speed, and the implications for information retrieval, it seems necessary to look closer at its structure.

ASSOCIATIVE PROCESSOR STRUCTURE

To illustrate structure and operation of a distributed logic APP the example of searching for a given keyword in a list of keyword-document-number pairs is given. The implementation in BASIC and the memory outlay on a sequential computer is given in Figure 1. A somewhat simplified structure of an

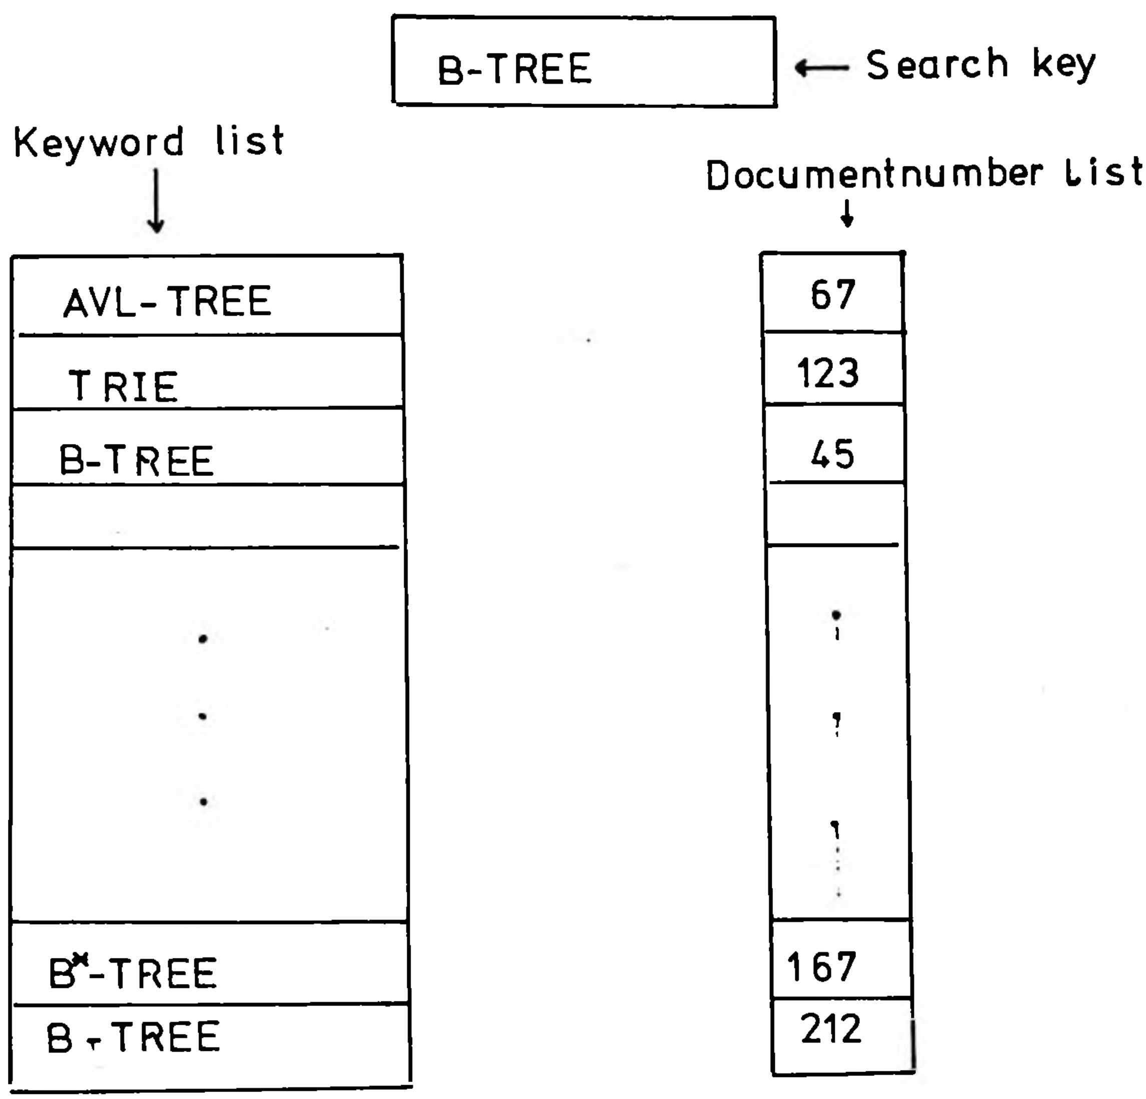


```
10 FOR I = 1 TO N
20    IF Searchkey ≠ Keyword(I) THEN 40
30    PRINT Docno(I)
40 NEXT I
```

FIGURE 1   SEQUENTIAL PROCESSOR MEMORY OUTLAY associative processor is shown in Figure 2 (Lea, 1976). The list of keyword-

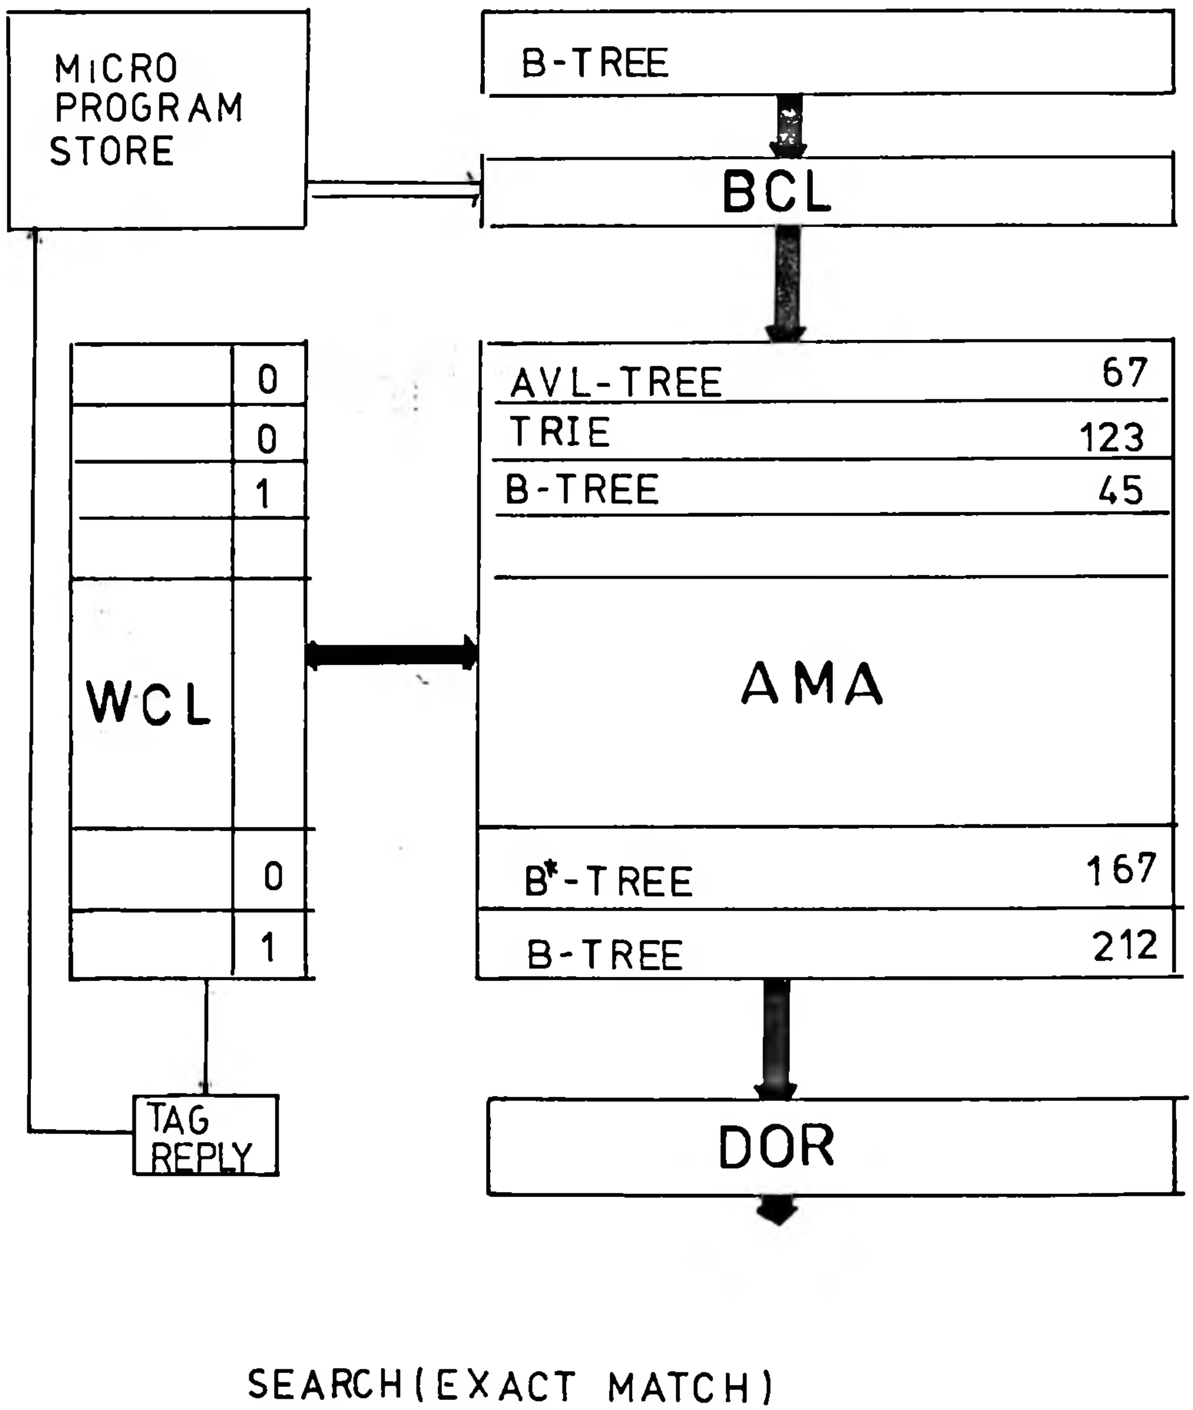

FIGURE 2 ASSOCIATIVE PROCESSOR STRUTURE

document-number pairs is loaded into the associative memory array (AMA) and the search key into the data input register (DIR). The search on the entire memory is just one step, one instruction, as the DIR is compared with all the words in the AMA simultaneously. If the search key matches a key in the AMA a bit in the same row is set in the tag register (TR). Multiple responses can be resolved using these tags and a row of the AMA can be transferred to the data output register (DOR). The bit-control logic (BCL) allows masking part of the DIR so that a search for a "truncated" key is possible. The word-control logic (WCL) contains the tag register and logic necessary to resolve multiple responses. It also facilitates more complicated searches for less than, greater than, maximum and minimum, between limits, etc. It is obvious from this example what the advantages of an APP are. If the AMA can store N keys, the associative processor is at least N-times faster than a sequential computer.

Now is the time to ask the question: "Why don't we all have associative processors, if they are that fast?"

ASSOCIATIVE PROCESSORS AND ASSOCIATIVE FILE STORES

The answer to this question can be found as early as 1966 in a paper by Hanlon (Hanlon, 1966): "The major drawback to development of content addressable memory is the lack of a suitable associative cell in a practical technology at an acceptable price." Associative memories have been implemented with various technologies, cryotrons, magnetic cores, semiconductors and integrated circuits. However, the size of memory was always small, approximately 1024 words of 100 bits. Large scale associative processors were not economical until LSI technology reduced the cost, and the construction of PEPE (Crane et al., 1972) and STARAN (Batcher, 1974) was possible. Another of these large scale associative processors is the Linear Array Processor developed by Finnila (Finnila, 1977). Customers for these machines are the "affluent" computer users, such as the U.S. Air Force and Navy. Experiments have been carried out with these machines to test their "promised" suitability to information retrieval and data base management. It has been reported (DeFiore, Berra, 1973), that they performed as expected, even though some problems remain (Berra, 1974).

As associative machines did not proliferate due to their high cost, other attempts were made to bring the cost down and make it available to the "normal" user. Researchers turned their attention to the disk and suggested to modify a disk into an associative file store by adding special circuitry to the disk heads. This was to make a disk behave like content addressable memory, albeit a slower one than regular associative memory. The large capacity of a disk, coupled with reasonably cheap logic, promised to produce a valuable associative file store. These ideas originated from Slotnick (Slotnick, 1971), and his "logic-per-track" research. Parhami's design (Parhami, 1972) was specifically created for information retrieval applications. A group at the University of Toronto is working in this area. They created an interesting device called RAP and published a considerable amount of literature (Ozkarahan et al., 1975; Schuster et al., 1976; Sadowski, Schuster, 1978).

Unfortunately, the slow speed is considered a serious drawback of such an associative file store. The alternative of using a fixed-head disk and adding the logic to each head is technically feasible but economically not viable, becuase the cost of duplicating the logic is too high.

One of the most recent suggestions is to insert a small associative processor into the disk channel, rather than having separate logic for each head. It has been shown (Lea, 1978), that a small APP with approximately 1K characters of memory inserted into the disk channel can be used successfully for data compression and decompression. Another recent paper (Schuegraf, Lea, 1979) has suggested that an approach which integrates special software techniques with such an associative processor in the disk channel will result in an economically feasible design, and will permit searching of large files "on the fly."

IMPLICATIONS FOR INFORMATION RETRIEVAL

Many associative devices have been proposed in the literature, but only large-scale processors have reached the production stage and are available to the user who can pay the high price. It is anticipated that more and smaller associative processors will become available in the near future as hardware technology progresses.

From the previous sections it seems obvious that a new "processing" philosophy and associated hardware are just around the corner. To be able to work with one instruction on a large number of data items in parallel is a radical change from the conventional sequential computer.

Programming techniques and software will have to change dramatically to accommodate the capabilities of the new hardware. Efficient utilization of such a system, driven by suitable software, can produce maximum benefit from the inherent parallelism and change the economics of information processing by computer. Some problems, up to now too expensive to handle, can suddenly be converted to a parallel executable program and done at a reasonable cost.

It appears that the new technology has serious implications for information retrieval and that we information scientists have a lot of thinking to do before we can apply the new techniques to our best benefit.

We must make an inventory of all our methods and techniques and analyze them thoroughly, how they can be adopted to the associative environment with its built-in parallelism. We may find the same phenomenon as researchers working in numerical analysis. Algorithms which were little known and performed poorly on sequential computers could be easily adapted to the parallel machines and out-performed parallel versions of the best sequential algorithms. Our basic approach may have to be changed too, a process which may be painful. Let me illustrate this on the example of indexing.

ASSOCIATIVE PROCESSING

One of the main purposes of indexing for computers is to establish a link between the content of a document and its location in a file, so that it may be retrieved quickly. The more thorough the indexing the easier the retrieval. If we have access to an associative file store which allows us to search for any arbitrary character string, or combinations thereof, in a reasonable amount of time, then things look different. It is obvious that the need for high-quality content indexing is reduced and that a more "fuzzy" method of indexing which keeps related items together would suffice. This problem will be discussed in a future paper (Schuegraf, 1979).

These ideas may be difficult for us to handle individually, and we should work together and share all our resources, physical as well as intellectual, in tackling these problems. When the new technology finally arrives on our doorsteps, delivered at reasonable cost, we should be ready for it, rather than being caught by surprise.

REFERENCES

Batcher, K. E. STARAN parallel processor system hardware, Proceed. Natl. Computer Conf., 1974, 405-410.

Berra, P. B. Some problems in associative processor applications to data base management, Proceed. Natl. Computer Conf., 1974, 1-5.

Carroll, J. M. Prospects for parallelism and the computer crunch, J. ASIS, Vol. 27, 1976, 63-69.

Crane, B., et al. PEPE computer architecture, Proceed. IEEE COMPCON., 1972, 57-60.

DeFiore, C. R., Berra, P. B. A data management system utilizing an associative memory, Proceed. Natl. Computer Conf., 1973, 181-185.

Finnila, C. A. The associative linear array processor, IEEE Trans. Computers, Vol.-C26/2, 1977, 112-125.

Hanlon, A. G. Content-addressable and associative memory systems, IEEE Trans. Computers, Vol. EC-15, 1966, 509-521.

Knuth, D. E. The art of computer programming, Vol. 3, Sorting and Searching. Addison-Wesley, 1974.

Lea, R. M. Associative processing of non-numerical information, Proceed. 1976 NATO Summer Institute on Computer Architecture, San Raphael, France, 171-215.

Lea, R. M. Text compression with an associative parallel processor, Computer Journal, Vol. 21, 1978, 45-56.

ASSOCIATIVE PROCESSING

Lee, C. Y., Paull, M. C. Intercommunicating cells--Basis for a distributed logic computer, Proceed. FJCC, 1962, 130-136.

Ozkarahan, E. A., Schuster, S. A., Smith, K. C. RAP--An associative processor for data base management, Proceed. Natl. Computer Conf., 1975, 379-389.

Parhami, B. A highly parallel computing system for information retrieval, Proceed. FJCC, 1972, 681-690.

Petersen, H. E. Content addressing and information retrieval, Proceed. IFIP Congress. 1962, North Holland Publishing Company.

Reich, D. L. Associative memories and information retrieval, in "Some Problems in Information Science." M. Kochen (Ed.), Scarecrow Press, 1969.

Sadowski, P. J., Schuster, S. A. Exploiting parallelism in a relational associative processor, Proceed. Fourth Workshop on Computer Architecture for Non-Numeric Processing, Aug. 1-4, 1978, ACM Press.

Schuegraf, E. J., Lea, R. M. An integrated approach to on-line retrieval from non-numeric data bases: An associative parallel processor using fragment indexing. Paper presented 1979 Conference on Information Science and Systems, Johns Hopkins University, Baltimore, Md.

Schuegraf, E. J. Indexing for associative processing. Paper submitted to Cdn. J. Inf. Science, 1979.

Schuster, S. A., Ozkarahan, E. A., Smith, K. C. A virtual memory system for a relational associative processor. Proceed. Natl. Computer Conf., 1976, 855-862.

Slade, A., McMahon, H. O. A cryotron catalog memory system, Proceed. EJCC, 1956, 115-120.

Slotnick, D. A. Logic-per track devices. In Advances in Computers, Vol. 10, 1971, 291-296.

Yau, S. S., Fung, H. S. Associative processors--A survey, Computing Surveys, Vol. 9/1, 1977, 3-28.