

NON-VON: A PARALLEL MACHINE ARCHITECTURE FOR KNOWLEDGE-BASED INFORMATION PROCESSING

David Elliot Shaw

Department of Computer Science
Columbia University
New York, New York 10027

ABSTRACT

NON-VON is a highly parallel machine designed to support the efficient implementation of very large scale knowledge-based systems. The utility of such a machine has been demonstrated analytically [1, 2] and through implementation of a working knowledge-based information retrieval system in which the NON-VON machine instructions were emulated in software [3]. In cooperation with the Stanford Computer Science Department, we have recently begun to implement the most important components of the machine as custom VLSI circuits.

I INTRODUCTION

The central focus of our recent research has been the investigation of highly parallel non-von Neumann machine architectures adapted to the kinds of operations that appear central to the operation of a broad class of large-scale knowledge-based systems. Our approach is based on an architecture [1] that supports the highly efficient evaluation of the most "difficult" set theoretic and relational algebraic operators. This research spans an unusual "vertical distance"—extending from the level of AI description languages, through the arena of logical formula manipulation, and down to the domain of hardware parallelism. Our initial efforts, reported in a recent Stanford doctoral dissertation [3], yielded:

1. The "top-level" architectural specification of a highly parallel machine—portions of which are to be implemented using custom nMOS VLSI circuits— which we have since come to call NON-VON.
2. An analysis of the time complexity of the essential parallel hardware algorithms to be executed on the proposed machine in the course of large-scale, meaning-based data manipulation.
3. The implementation of an operational *knowledge-based information retrieval system*, demonstrat-

ing the use of NON-VON (emulated in software) in support of a very high level descriptive formalism based on the language KRL [4].

Our recent work at Columbia and Stanford has focused on:

1. Progressive refinement and elaboration of the NON-VON architecture, including a preliminary functional design and high-level layout for two of the most important VLSI chips.
2. The exploration of applications of the NON-VON class of architectures to practical problems in such areas as artificial intelligence, database management, and "very-high-level" programming languages.

In this short communication, we will briefly sketch our progress to date.

II A KNOWLEDGE-BASED RETRIEVAL SYSTEM

A surprisingly large share of the kinds of information processing activities with which both human and automated data processors are charged may be viewed as involving various kinds of knowledge-based matching or retrieval problems. For our dissertation research, we chose a concrete task of manageable scope that captured many of the most important characteristics of the general "conceptual matching" problem: the *knowledge-based document retrieval* application.

The user of a knowledge-based document retrieval system formulates a high-level pattern *description* which is then matched against all descriptions contained in a large *target collection* (the set of books in a computer science library, for example). In general, the success of a match may depend on domain-specific entities and relationships. In the case of the computer science library, for example, the system might "know about" such entities as *computers*, *algorithms*, *programmers* and *storage devices*. Certain characteristic attributes of these entities (the storage *medium* attribute, for example, whose values differ for different kinds

of *storage devices*) might also be included in this domain-specific knowledge.

Typical of the kinds of relationships which might be embodied in the *knowledge* base of such a system is the fact that a *tape drive* is a particular kind of *storage device* whose *storage medium* is always magnetic tape. A simple deductive inference based on this relationship might establish that a pattern description in which the subject is described as a *storage device* with *magnetic* tape as its *medium* would be satisfied by a target description with subject described as a *tape drive*. The description language used in our experimental system is based closely on KRL; the document description shown below may suggest some of its most important features:

```

a Document
  authors
    set-with-all-of
      Thompson
      Walters
    set-of
      an Engineer
  coun tries-of- publication
    set-with-any-of
      USA
      Great-Britain
  subject
    involves
      an Invention with
        purpose
        or
          Power-generation
          Power- transmission
  printing-dates
    set-with-exactly
      1959
      1962
a Textbook

```

III PREDICATE LOGIC FOR SPECIFICATION AND EXECUTION

The rules defining the semantics of matching within our knowledge-based description language were not embedded inextricably within the code of the retrieval system, but were instead formulated as an independent, separable set of twenty-two axioms expressed in a restricted first-order predicate calculus. One of these axioms (which embodies the basis on which a target description of the form "a *Storage-device* with *medium* = *Magnetic-tape*" would successfully match a less restrictive description of the form "a *Storage-device*" is presented below:

```

Per-imp-per (tar-per, pat-per) =
3 proto .
(Per-proto (pat-per, proto) A
Per-proto (tar-per, proto)) A
V slot .
((3 pat-Gil .
Obj-slot-fill (pat-per, slot, pat-fill))
3 3 pat-fill, tar-fill .
(Obj-slot-fill (pat-per, slot, pat-fill) A
Obj-slot-fill (tar-per, slot, tar-fill) A
Dtion-imp-dtion (tar-fill, pat-fill)))

```

The set of matching axioms serves not only as a modular, perspicuous and easily formulated *specification* of the semantics of the knowledge-based description language, but as an *executable* "program" for actually carrying out the matching process which it describes. More precisely, the matching process is carried out by a procedure called *LSEC* (for Logical Satisfaction by Extensional Constraint), which interprets the set of axioms to identify all target descriptions that match the user-specified pattern description according to the matching rules for the description language, making reference to the domain-specific knowledge base. (Details of the LSEC algorithm are presented in (3).)

IV NON-VON AS A LOGIC ENGINE

Unfortunately, the LSEC algorithm relies very heavily on the execution of several operations which, on a von Neumann machine, are quite expensive when their operands comprise a large amount of data (as is the case in the sort of application with which we are concerned). Specifically, the algorithm repeatedly evaluates the most expensive operators of a *relational algebra* [5] — in particular, the equi-join operation. NON-VON evaluates the most time-consuming relational algebraic operators in a highly efficient manner, particularly when the argument relations are large. The architecture in fact achieves an asymptotic improvement of $O(\log n)$, with very favorable constant factors, over the best evaluation methods known for these operators on a conventional system, without the use of redundant storage, and using currently available and potentially competitive technology.

V THE NON-VON ARCHITECTURE

At the top level, the NON-VON machine architecture comprises a *Secondary Processing Subsystem* (SPS), based on a bank of "intelligent" rotating storage devices and designed to provide very high accessing and processing bandwidth, along with a smaller, but faster *Primary Processing Subsystem* (PPS), again utilizing a high degree

of parallelism, in which the relational algebraic operators may be very quickly evaluated. Transfer between the two devices is based on a process of hash partitioning, which is performed entirely in hardware by logic associated with the individual disk heads, and which divides the argument relations into *key-disjoint* buckets suitable for "internal" evaluation.

In our demonstration system, the operation of NON-VON was emulated sequentially (in MACLISP on the DEC KL-10 at the Stanford Artificial Intelligence Laboratory), and hence, so inefficiently that testing using a large sample document collection would have been impractical. On a physical implementation of NON-VON, though, our knowledge-based retrieval system would in a sense be able to look at *all* documents in the library "at once", applying a considerable amount of physical hardware concurrency to a much more complex (inferential) matching task than has previously been attacked using parallel hardware.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the substantial contributions of Bob Floyd, Hussein Ibrahim, Don Knuth, Gio Wiederhold and Terry Winograd to the work reported in this communication. This research was supported in part by the Defense Advanced Research Projects Agency under contract MDA903-77-C-0322.

REFERENCES

- [1] Shaw, David Elliot, "A Hierarchical Associative Architecture for the Parallel Evaluation of Relational Algebraic Database Primitives", Stanford Computer Science Department Report STAN-CS-79-778, October, 1979.
- [2] Shaw, David Elliot, "A Relational Database Machine Architecture", *Proceedings of the 1980 Workshop on Computer Architecture for Non-Numeric Processing*, Asilomar, California, March, 1980.
- [3] Shaw, David Elliot, *Knowledge-Based Retrieval on a Relational Database Machine*, Stanford Ph.D. Dissertation and Stanford Computer Science Department Report STAN-CS-80-823, August, 1980.
- [4] Bobrow, D. G., and Winograd, Terry, "An Overview of KRL-0, a Knowledge Representation Language", *Cognitive Science*, vol. 1, no. 1, 1977.
- [5] Codd, E. F., "A relational model of data for large shared data banks", *Communications of the ACM*, vol. 13, no. 6, June, 1970.