ICOT Technical Report: TR-012

TR-012

A Relational Database Machine: First Step to Knonwledge Base Mahine

by

Kunio Murakami, Takeo Kakuta,
Nobuyoshi Miyazaki, Shigeki Shibayama
and Haruo Yokota

May, 1983

ICOT

Mita Kokusai Bldg. 21F
4-28 Mita 1-Chome
Minato-ku Tokyo 108 Japan

(03) 456-3191~5
Telex ICOT J32964

Institute for New Generation Computer Technology

Technical Report of ICOT Research Center : TR - 012 May, 1983

A Relational Database Machine:
First Step to Knowledge Base Mahine

Kunio Murakami, Takeo Kakuta,
Nobuyoshi Miyazaki, Shigeki Shibayama
Haruo Yokota

ICOT Research Center
Institute for New Generation Computer Technology

Mita Kokusai Bldg. 21F, 1-4-28 Mita,
Minato-ku, Tokyo, 108
Tel:03-456-3195, Telex:ICOT J32964

/

A Relational Data Base Machine: First Step to Knowledge Base Machine

by Kunio Murakami, Takeo Kakuta, Nobuyoshi Miyazaki
Shigeki Shibayama and Haruo Yokota

Research Center
Institute for New Generation Computer Technology

Mita Kokusai Bldg. 21F
4-28 Mita 1-Chome, Minato-ku, Tokyo 108 Japan

Abstract

The Japan's Fifth Generation Computer System project is divided into three stages. In the first three-year stage, a working relational data base machine is developed for a software development support system to be used in the second stage and also for an experimental system which provides a research tool for the knowledge base machine.

The paper briefly describes the concepts and architecture of the ional data base machine named "Delta" which is currently under development at ICOT.

1. Introduction[1]

The Fifth Generation Computer System is expected to realize the knowledge information processing systems which provide a very high level and flexible man-machine interface, based on huge volumes of knowledge data.

The functions of the Fifth Generation Computer System are classified as follows.

(1) Problem-solving and inference.
(2) Knowledge-base management.
(3) Intelligent man-machine interface.

It is one of the most important research subjects to develop a knowledge base machine which stores and manages a large amount of knowledge data with high performance.

In the Fifth Generation Computer System, the knowledge-base machine is expected;

(1) to hold a large amount of knowledge data which are structured well for accessing each knowledge item effectively,
(2) to search and retrieve knowledge items effectively and hand them to the inference machine which requires them,
(3) to compile and integrate knowledge data, which are sent from the inference machine, into the knowledge base.

At ICOT, the knowledge base machine is developed in three stages. In the first stage, a data base machine is developed based on a conventional (von Neumann) multiprocessor architecture with certain enhancements by specialized hardwares. Next, the knowledge base machine is developed based on a highly parallel (non von Neumann) machine architecture. Finally, in the third stage, the knowledge base machine is integrated with the parallel inference machine into the ICOT's prototype of the Fifth Generation Computer System.

2. Design Objectives[2]

As shown in Fig. 1, the data base machine, named "Delta", developed in the first stage is connected with multiple sequential inference machines (SIMs) via a local area network (LAN) and a shared common memory on a multibus. Delta plays the role of an external data base machine for each sequential inference machine. Relational model was chosen as the data base model for Delta because of its affinity to logic programming languages such as Prolog and ICOT's Fifth Generation Kernel Language, that is SIM's machine language. Delta provides a high level command set based on relational algebra such as union, intersection, projection and join.

The main objectives of developing Delta are;

(1) to form a software development support system to be used in the middle research and development stage of the fifth generation computer system,
(2) to collect experimental data for evaluation on the interface between the sequential inference machine and the relational data base machine,
(3) to provide a research tool for integration of knowledge base and inference mechanisms into the knowledge base machine,
(4) to provide a research tool for a distributed relational data base system.

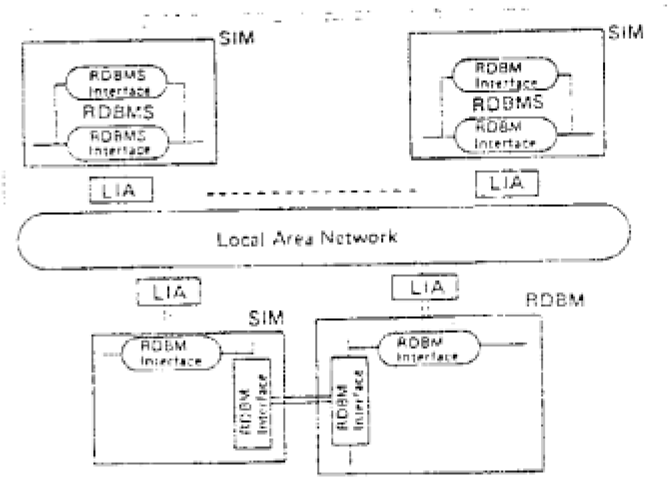

Fig. 1. Software Development Support System

3. Architectural Overview[3]

The basic Delta's concepts are;

i) High-level interface based on relational algebra applicable to logic programming classes of languages
ii) Efficient query processing capability by use of dedicated relational data base operation engine hardwares
iii) large capacity hierarchical-structured memory.

Delta is a functionally distributed multiprocessor system which is composed of three main subsystems as shown in Fig. 2:

(1) Control Processor (CP)
(2) Relational Data Base Engine (RDBE)
(3) Hierarchical Memory (HM).

Other subsystems provided to form Delta are:

(1) Interface Processor, a front-end processor which interfaces Delta to a LAN environment.
(2) Maintenance Processor, a supervisory processor which monitors other subsystems.
(3) Multibus Interface, a general-purpose bus interface which provides Delta another interface besides LAN.

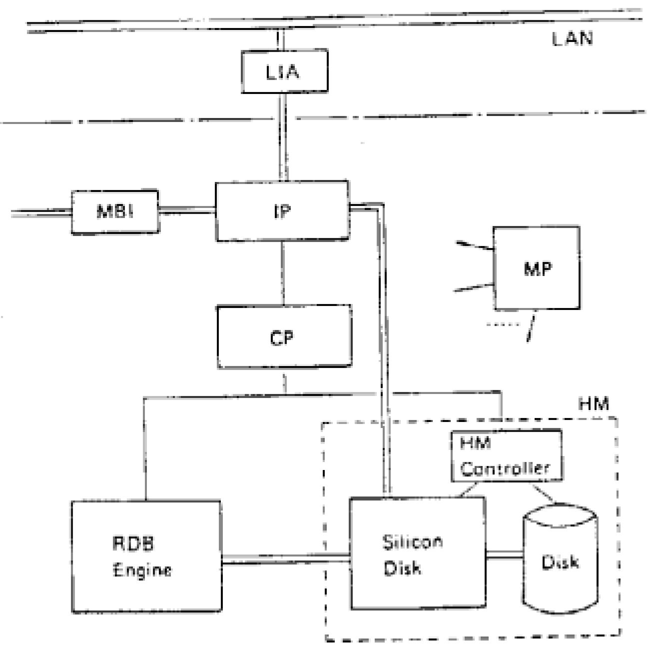

Fig. 2. Block Diagram of Delta

3.1 Control Processor (CP)

The main role of Control Processor is general control of the whole RDBM. CP receives a command-tree composed of a series of RDBM commands structured to a logical tree which represents a DB transaction, then translates a command-tree into a set of internal subcommands to activate RDBE and HM, which actually performs query processing. CP's other important role is to provide a mutl-user environment for interactive use via LAN. This transaction concurrency is one of the most complex functions among levels of software in CP to implement a working Data Base Machine.

CP's software is divided into two levels of hierarchical structure: DB Management level and Unit-Command Process level.

This DB Management level is similar to that of an existent DBMS with which the multi-user service along with DBMS manipulating facilities should be provided.

The CP's Unit-Command Process software controls hardware resources to complete a command in the RDBM command set which is determined in conformity to relational algebra.

3.2 RDB Engine (RDBE)

An RDBE is a piece of dedicated hardware for executing RDBM commands at the lowest level. RDBE has its own command set associated with the RDBM command set.

The RDBE command set not only matches the higher-level one, but also has efficient processing algorithm done at the hardware level.

RDBE is tightly connected to HM in the sense that a data stream path is provided between RDBE and HM to flow data strem in both directions. When a query is translated in CP and RDBE receives a command from CP, it waits for arrival of data stream from HM, which also received a command specifying a particular data streaming into RDBE. Receiving the first data from HM, RDBE starts the operations on the data stream to complete most of the command processing overlapped with data transfer. The basic RDBE command processing mechanism is the sorting as the preprocessing and the merging for command executing. The input data stream is first sorted using pipeline merge-sorting algorithm. A sorted stream is then sent to a Relational Operation Unit (ROU) where a relational algebra-like command is executed on a sorted data stream. The basic processing scheme in ROU is merging or comparing an input stream against another data stream stored in backtraceable first-in-first-out (FIFO) memory and then controlling the resulting output. Heavy relational algebra operations such as join and projection can be efficiently processed if the relation or attribute is sorted by its values.

3.3 Hierarchical Memory (HM)

Hierarchical Memory subsystem is to store the actual database. To provide fast access-time and large storage space, Delta has a memory system with two layers of different-natured memory devices, incorporating mechanisms to stage and destage bunches of data.

The lowest level storage devices are moving-head magnetic disks, which are still considered to be the most appropriate devices for storing a good amount of data at moderate cost. In actual database query manipulation, the access speed of a moving-head disk is too unsatisfactory if every temporary result goes to the disk and it needs time-consuming disk activation every time to access the temporary result. A large semiconductor disk comprised of random access memory fills this gap. This semiconductor disk works as both cache and buffer, in the sense that transparent data management is carried out in cache part and programmable control can be accommodated in a buffer part, thereby optimizing greater system performance. The control mechanism for cache and buffer is self-contained in HM. Other subsystems only need to specify HM minimum amount of information or parameter to utilize this facility.

To design a database machine, it is of great importance to define an internal schema, or how relations are stored in memory, because it greatly influences system performance and flexibility.

We have chosen the attribute-based schema because it suited for our RDBE hardware in which an attribute of relation is processed on-the-fly to give fast RDB query manipulation. The problem associated with this schema is tuple reconstruction, which is another important HM function. Tuple reconstruction, that is to make up a tuple by gathering the corresponding attribute values by their tuple-identifier as a key, is needed in the later phase of query processing. This problem is solved by the two-level clustering, which is used to reduce the amount of data to be reconstructed. The idea is that in the first-level, the relation is primarily clustered by the attribute values, then the primary clusters are divided secondarily by the tuple-identifier values to form secondary clusters which correspond to the access units of storage.

3.4 Other Subsystems

Interface Processor (IP) is responsible for interfacing between LAN and Delta. It has facilities for handling the levels of network protocols, which make it easy for Delta to communicate with the other nodes (SIM's) connected to LAN. It also has a port to HM, this port acts as the main data path to SIM's. Once the resulting relation is calculated and stored in HM, CP delegates the IP to transfer it to the SIM via LAN which has made the query and requiring the result.

The Roles of Maintenance Processor (MP) are to monitor and maintain other components and to collect the evaluation data. Since Delta is an experimental vehicle to investigate various characteristics in actual operation, statistical data such as resource contention, cache usage, and query processing rate are needed for future improvement.

Multibus Interface (MBI) is a subsystem connected to Interface Processor. The function that this interface provides is to open a general-purpose interface to the outside environment. One thing to note is the shared common memory in this interface which could be used to experiment the effectiveness of faster and closer interface with SIM's.

4. Toward Knowledge Base Machine [2],[4]

Now in the first stage of research on the fifth generation computer systems, the relational data base machine Delta as the first version of a knowledge base machine, and the sequential inference machine as that of an inference machine are separately being developed as independent hardware modules.

Because Delta can be connected to SIMs via LAN and a shared common memory on a multi bus, it will be used to form an experimental system which provides a research tool for the following investigation toward the knowledge base machine.

(1) A research on how to interface between logic programming languages and relational data bases through an evaluation of the compiled approach implemented on the software development support system.

(2) A research on the hardware mechanisms to support facilities of knowledge representation, acquisition and reconstruction in the course of implementation of knowledge representation languages and expert systems.

(3) A research on how to implement distributed relational data bases and protocols to handle them.

(4) A research on how to integrate some elementary reference machanisms (e.g., variable bindings) into a relational data base machine.

5. Conclusion

This paper describes the concepts and the architecture of the ICOT's relational data base machine, which will be implemented by the end of 1985 as the first step toward the knowledge base machine.

Acknowledgments

The authers would like to thank Mr. K. Fuchi, Director, Dr. K. Furukawa and Dr. S. Kunifuji of ICOT for their useful comments.

References

[1] ICOT (ed.) : Outline of Research and Development Plans for Fifth Generation Computer Systems, May 1982.

[2] JIPDEC (ed.): Proceedings of International Conference of Fifth Generation Computer Systems, Oct. 19 ~ 22, 1981.

[3] S. Shibayama, T. Kakuta, N. Miyazaki, H. Yokota and K. Murakami: A Relational Date Base Machine "Delta", ICOT TM-003.

[4] S. Kunifuji and H. Yokota: PROLOG and relational databases for fifth Generation Computer Systems, Proceedings of the CERT82 workshop on "Logical Bases for Data Bases", 1982.