The MEDUSA Project: Autonomous Data Management in a Shared-Nothing Parallel Database Machine

G.M. Bryan
Department of Computing
University of Western Sydney, Nepean,
PO Box 10 Kingswood NSW 2474 Australia
g.bryan@nepean.uws.edu.au

W. E. Moore, B.J. Curry, K.W. Lodge, J.Geyer
Centre for Parallel Processing and Industrial Control,
Charles Sturt University, Mitchell,
Bathurst, NSW 2795, Australia
{wmoore,bcurry,klodge,jgeyer}@csu.edu.au

The MEDUSA Project

In a 1989 article in Scientific American, David Gelernter speculated about the need and design of "Information Refineries", machines capable of taking massive amounts of data and converting it into knowledge. In order to investigate the design of a database machine which would be capable of acting as the data storage engine of Gelernter's "Information Refineries" the authors initiated the MEDUSA Project as a joint undertaking between the University of Western Sydney, Nepean, and Charles Sturt University, Mitchell.

Hardware Architecture

The current MEDUSA prototype utilises a shared-nothing architecture based on the INMOS Transputer. Each of the three processing nodes used in the prototype consists of two T805 transputers with a T222 SCSI interface to a Maxtor 180 MByte disk unit, as shown in figure 1.

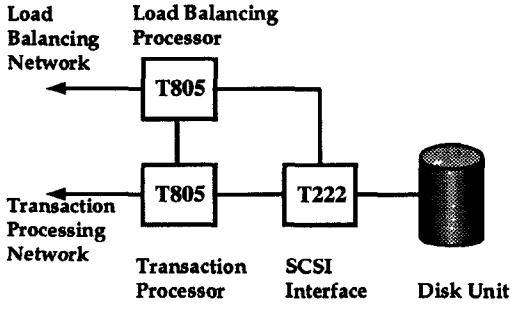

Figure 1, Schematic Diagram of MEDUSA Processing Node

Principle Goals of the Project

The principle goals of the MEDUSA Project are to develop a prototype database machine based on a shared-nothing architecture using low cost "off the shelf " components to support research in the following areas: autonomous data management; user data interfaces; back-up and security systems.

Autonomous Data Management

If the full potential of these machines is to be exploited they must exhibit a level of operational autonomy similar to the existing file server technology used on local area networks. That is, they should be capable of seamless integration into a network without requiring changes to existing software development practices or additional specialised staff to maintain their operational efficiency. Operational autonomy can be achieved by a self-organising or self-tuning database. In a self-organising database environment the database management system (DBMS) is responsible for the system tuning and data re-organisation, to achieve optimal performance. Human intervention in the system's maintenance task is reduced to no more than the changing of backup media or carrying out hardware maintenance. The traditional approach taken to performance tuning is normally heuristically based, with the DBMS providing the administrator with a number of tuning parameters which are adjusted based on the DBA's experience. The lack of an experienced DBA forces most sites to rely on intuition or guesswork, the latter being more common.

SQL Interface

An SQL interface, MedusaSQL, is available for data retrieval on MEDUSA. Currently, a reduced version of SQL92 standard is operational. This version allows retrieval of selected attribute values from tuples which meet criteria specified in the WHERE clause. Simple projections, joins and selections are possible using the features of the SELECT, FROM and WHERE clauses implemented thus far. Full implementation of SQL92 is planned to be completed by mid-1994.

Back-up and Security Systems

Back-up methods for a large database machine are vital for data security. A back-up algorithm for logical files has been developed for the MEDUSA system that allows parallel back-up during machine use. This algorithm provides continuous back-up while still allowing continuous query servicing by the MEDUSA system.

Future Directions

Major research thrusts will be made in the following areas.

- Development of a full SQL92 interface with extensions.
- Implementation of a prototype system in industry (A hospital Records System is our current objective.)
- The development of an interface protocol to allow the connection of MEDUSA to super computers.
- Investigation of redundant data distribution for enchanced query performance.