Database Research at NTHU and ITRI

Arbee L.P. Chen
Department of Computer Science
National Tsing Hua University
Hsinchu, Taiwan 300, ROC
Email: alpchen@cs.nthu.edu.tw

1 Introduction

National Tsing Hua University (NTHU) was founded in 1911 and is located in a suburb of the city of Hsinchu, Taiwan, about 50 miles southwest of Taipei, the capital city. Its Computer Science Department was established in 1977, and currently has 23 faculty members (including two IEEE Fellows) and 160 graduate students, among them 60 are Ph.D. students. With National Chiao Tung University, the Industrial Technology Research Institute (ITRI) and the Science-based Industrial Park in the vicinities, the entire milieu of the University has been congenial to scientific pursuits. The University is organizing the 11th IEEE Data Engineering Conference to be held from March 6 to 10, 1995 at the Grand Hotel in Taipei.

The Computer and Communication Research Laboratories (CCL) at ITRI are dedicated to the development of the computer, communication and consumer electronics technologies and intend to become a world-class research organization in these fields. The main areas of database research and development at CCL/ITRI include multidatabase systems, real-time database systems, and multimedia database systems. Related projects are supported by the Ministry of Economic Affairs, with collaboration from the Database Research Laboratory at NTHU.

This report presents database research and development activities at NTHU and ITRI. Most of the papers cited in this report are available via anonymous ftp to ftp.cs.nthu.edu.tw under /pub/db-papers.

2 WAW (We Are the World)

Principal Investigator: Prof. Arbee L.P. Chen at NTHU
Support: National Science Council and Institute for Information Industry

This project constructs a heterogeneous database system which contains relational and object databases. An object view and a semantics-enriched relational view are provided to the users such that users can treat the heterogeneous database system as either a relational or object one depending on their preference. Research topics include a schema definition language for constructing relational concept hierarchies, a probabilistic relational algebra, a probabilistic approach to identifying the same real-world entities from instances in different databases, *entity join* and *neighborhood join* optimization, extending semijoin tactics in heterogeneous database environment, foreign function optimization, deriving object views from relational schemas, object view integration, and distributed object query optimization. They will be briefly described in the following.

2.1 Relational Frontend

In heterogeneous database systems, partial values have been used to resolve schema integration problems. Performing operations on partial values may produce *maybe tuples* in the query result which cannot be compared. Thus, users have no way to distinguish which maybe tuple is the most possible answer. To solve this problem, we propose an approach to resolve the schema integration problems using probabilistic partial values and develop a full set of extended relational operators for manipulating relations containing probabilistic partial values. With this approach, the uncertain answer tuples of a query are associated with degrees of uncertainty (represented by probabilities). That provides users a comparison among maybe tuples and a better understanding on the query results. Besides, extended selection and join are generalized to α-selection and α-join, respectively, which can be used to filter out maybe tuples with low probabilities — those which have probabilities smaller than α.

For the same real-world entity, it can be represented as tuples in different databases. By joining the tuples on the keys all the information for the same entity can be obtained. However, due to database heterogeneities various conflicts can exist in the tuples, including incompatible keys. These conflicts need to be resolved before they can be joined. We extend the aforementioned probabilistic techniques to estimate the degree for two tuples to represent the same entity (when their keys are incompatible), to process queries involving joins between relations in different databases on their keys (named entity join), and to estimate the degree of uncertainty for the query results.

Because inconsistent data may exist in different database relations, traditional query optimization methods need to be extended. We propose an approach to correctly transform a global query into local subqueries to preprocess entity join queries in multiple sites with an attempt to lower the cost of data transmission. Besides, an extension of the traditional semijoin tactics is explored to further optimize entity join query processing in a wide area multidatabase environment.

The conversion function needed for schema integration in a multidatabase system can be considered as one kind of foreign functions which is usually time-consuming to process. We consider the order of executing the joins and conversion functions and find an optimal algorithm with polynomial complexity for a special case. However, finding the optimal execution plan for the general case is a hard combinatorial problem with exponential complexity. An efficient heuristic algorithm is developed to solve the problem and the simulation result shows its good quality.

[1] A.L.P. Chen, "Outerjoin Optimization in Multidatabase Systems," Proc. IEEE International Symposium on Databases in Parallel and Distributed Systems (DPDS), 1990.

[2] C. Wang, A.L.P. Chen, S.C. Shyu, "A Parallel Execution Method to Minimizing Distributed Query Response Time," IEEE Transactions on Parallel and Distributed Systems, pp. 325-333, May 1992.

[3] F.S.T. Tseng, A.L.P. Chen, W.P. Yang, "Answering Heterogeneous Database Queries with Degrees of Uncertainty," Distributed and Parallel Databases: An International Journal, Kluwer Academic Publishers, pp. 281-302, 1993.

[4] P.S.M. Tsai and A.L.P. Chen, "A Localized Approach to Query Optimization in Heterogeneous Database Systems," Journal of Information Science and Engineering (to appear).

[5] P.S.M. Tsai and A.L.P. Chen, "Querying Uncertain Data in Heterogeneous Databases," Proc. IEEE International Workshop on Research Issues on Data Engineering: Interoperability in Multidatabase Systems, 1993.

[6] C.S. Chang and A.L.P. Chen, "Determining Probabilities for Probabilistic Partial Values," Proc. International Conference on Data and Knowledge Systems for Manufacturing and Engineering, 1994.

[7] P.S.M. Tsai and A.L.P. Chen, "Optimizing Entity Join Queries in a Wide Area Multidatabase Environment," NTHU-TR-CS-94-001.

2.2 Object Frontend

In a multidatabase system which consists of object databases, the same real-world entity can be stored as objects in different databases with incompatible object identifiers. How to identify and integrate these objects representing the same entities such that (a) object duplication in the query result can be avoided, (b) information for the entity can be gathered, and (c) the specialization of multiple classes can be built is an important issue to provide a well structured global object view and a more informative query result. We define *isomeric objects* as objects in different databases, which represent the same real-world entity. A technique similar to that for joining relations in different databases on incompatible keys has been applied to investigate the problem of object isomerism in the integration of multiple object databases.

To integrate object schemas, we define corresponding assertions for the database administrators to specify the semantic correspondences among component object schemas. Based on these assertions, integration rules are designed, which use a set of primitive integration operators to restructure the component schemas for resolving the conflicts and do the integration. The principle of our integration strategy is to keep the data of component databases retrievable from the global schema without losing information.

The strategies for processing global queries are proposed, which use the mapping information between the global schema and component schemas to decompose a global query into a set of subqueries. A form of equation is defined to denote the mappings for attributes and object instances among a virtual class and its constituent classes. A Flow Control Language is defined to specify the execution flow of the subqueries as well as the integration of the partial results. Some query optimization techniques are considered in the specification of the execution flow. The results returned from several component databases need to be integrated, which may include certain results and maybe results. Isomeric objects will be identified and integrated. Therefore, more informative query answers may be derived from the databases.

The object frontend enables a relational database system to act like an object database system. It requires two functions. One is to derive object views from relational schemas; and the other to transform an object query into its corresponding relational query and to transform the query result from the relational database system into the corresponding object format. We develop a set of derivation rules to construct an object view from the relational schema. These rules can derive the object features of class hierarchies, complex attributes, set attributes, and complex set attributes. The query transformation strategies are also studied.

[1] A.L.P. Chen, J.L. Koh, C.T. Kuo, C.C. Liu, "Schema Integration and Query Processing for Multiple Object Databases," Journal of Integrated Computer-Aided Engineering: Special Issue on Multidatabase and Interoperable Systems, Wiley Interscience (to appear).

[2] J.L. Koh and A.L.P. Chen, "Integration of Heterogeneous Object Schemas," Proc. Entity-Relationship Approaches, 1993.

[3] C.C. Liu and A.L.P. Chen, "Object View Derivation and Object Query Transformation," Proc. IEEE COMPSAC, 1994.

[4] A.L.P. Chen, P.S.M. Tsai, J.L. Koh, "Identifying Object Isomerism in Multidatabase Systems," NTHU-TR-CS-93-001.

[5] C.C. Liu and A.L.P. Chen, "Parallel Query Processing in Distributed Object Database Systems by Query Packets," NTHU-TR-CS-94-002.

3 Incomplete Information

Principal Investigator: Prof. Arbee L.P. Chen at NTHU
Support: National Science Council

3.1 Partial Values

Imprecise data exist in databases due to their unavailability or data/schema incompatibilities in a multidatabase system. The notion of partial values has been employed for representing imprecise data. We study the problem of eliminating redundant partial values which may result from a projection on an attribute containing partial values. The redundancy of partial values is defined through the interpretation of a set of partial values. This problem is equivalent to searching a minimal semantically-equivalent subset of a set of partial values. A semantically-equivalent subset contains exactly the same information as the original set. We derive a set of useful properties and apply a graph matching technique to develop an efficient algorithm for searching such a minimal subset and therefore eliminating redundant partial values. By this process, we not only provide a concise answer to the user, but also reduce the communication cost when partial values are requested to be transmitted from one site to another in a distributed environment. Moreover, further manipulation of the partial values can be simplified.

We also derive properties to refine partial values into more informative ones. In some cases, they can even be refined into definite values. Such a refinement is possible when there exist range constraints on attribute domains, or referential integrities, functional dependencies, or multivalued dependencies among attributes. Our work actually eliminates redundant elements in a partial value.

Finally, we define a set of extended aggregate operations, namely sum, average, count, maximum, and minimum, which can be applied to an attribute containing partial values. Two types of aggregate operators are considered: scalar aggregates and aggregate functions. We study the properties of the aggregate operations and develop efficient algorithms for count, maximum and minimum. For sum and average, in general it takes exponential time complexity to do the computations.

[1] F.S.T. Tseng, A.L.P. Chen, W.P. Yang, "Searching a Minimal Semantically-Equivalent Subset of a Set of Partial Values," The VLDB Journal, pp. 489-512, October 1993.

[2] F.S.T. Tseng, A.L.P. Chen, W.P. Yang, "Refining Imprecise Data by Integrity Constraints," Data & Knowledge Engineering, North-Holland, pp.299-316, 1993.

[3] C.S. Chang and A.L.P. Chen, "Aggregate Functions over Probabilistic Data" (in revision for Information Sciences).

[4] A.L.P. Chen, J.S. Chiu and F.S.T. Tseng, "Evaluating Aggregate Operations over Imprecise Data" (in revision for IEEE Transactions on Knowledge and Data Engineering).

3.2 PV-Table

We propose an extended relational model, named Pv-table, for representing exclusive disjunctive data (EDD). EDD in Pv-tables are represented by associating partial values and variable sets with attributes. Pv-table is different from other existing extended models with partial values in the following two aspects: (1) the values which do not satisfy the query predicates are retained and signed as unqualified. By this representation, the information about both the data and the query can be preserved. Also, the attribute which contributes to maybe information can be identified. Consequently, indefinite and maybe information can be properly handled in Pv-tables; and (2) each partial value is coupled with a variable set. In evaluating a join operation, all variables in the two variable sets are copied to the resultant variable set. Hence, the relationship between the join result and the original values is kept.

The relational operators are extended to manipulate Pv-tables. We show that queries consisting of extended selection, union, intersection, Cartesian product and join satisfy sound and complete conditions. Moreover, we show that useful algebraic properties are preserved for the Pv-table model. A disjunctive query $Q_1 \vee Q_2$ can be simply replaced by the extended union of Q_1 and Q_2. Similarly, a conjunctive query $Q_1 \wedge Q_2$ can be replaced by the extended intersection of Q_1 and Q_2. These properties not only provide a natural interpretation of queries but also facilitate query decomposition in distributed databases. In contrast, these properties are not held for other models extended with only partial values.

Query processing on Pv-tables is based on set operations. It avoids solving the satisfiability problem which is known to be computationally expensive.

Three kinds of Pv-tuple sets are classified according to relationships among Pv-tuples in a Pv-table. Each set possesses an important property. Based on these properties, we show that the interpretation of Pv-tables can be formalized in a semantically meaningful way. In addition, the lower and upper bounds of interpretation and the range of cardinalities of possible relations can also be characterized. Moreover, we show how to reduce a Pv-table as a query result into a more concise form. The reduction process preserves the same information, and thus is commutative with relational operations. The *condensation process* can further reduce Pv-tables. However, it may lose some relationships among Pv-tuples, and thus is not commutative with relational operations. By examining relationships among Pv-tuples, *redundant* and *superfluous* Pv-tuples can be identified and removed. The conditions for determining *uniteable* and *mergeable* Pv-tuples, and the processes for uniting and merging Pv-tuples are also given.

Pv-tables are further generalized to handle queries, updates and constraint enforcement in databases with EDD. A set of generalized relational operators are defined, on which queries and updates are based. The chase algorithm are generalized to enforce constraints consisting of a set of functional and join dependencies.

The semantic correctness of the generalized relational operators, update operations and constraint enforcement are established. Since the result of difference between generalized Pv-tables may be cumbersome and difficult to interpret, we also show how to adequately approximate a generalized Pv-table as the query result.

[1] J.S. Chiu and A.L.P. Chen, "A Sound and Complete Extended Relational Algebra for Exclusive Disjunctive Data," to be presented at VLDB'94 as a poster paper.

[2] J.S. Chiu and A.L.P. Chen, "An Exploration of Relationships among Exclusive Disjunctive Data" (in revision for IEEE Transactions on Knowledge and Data Engineering).

[3] J.S. Chiu and A.L.P. Chen, "Queries, Updates and Constraint Enforcement in Databases with Exclusive Disjunctive Data," NTHU-TR-CS-94-003.

[4] J.S. Chiu and A.L.P. Chen, "A Note on "Incomplete Relational Database Models Based on Intervals"" (in revision for IEEE Transactions on Knowledge and Data Engineering).

4 User-Friendly Interface

Principal Investigator: Prof. Arbee L.P. Chen at NTHU
Support: National Science Council

In a database system, data are retrieved via a well-defined query language. Although some query languages such as SQL can be very powerful, users may suffer greatly from their complex usage. Research on accessing databases in natural language usually employs an intermediate form for the mapping process from natural language to database languages. However, much effort is needed to bridge the gap between the existing intermediate forms and the database languages. We propose a methodology to map natural language constructs into relational algebra through E-R representation. This methodology employs a logical form to represent the natural language queries. The logical form has the merits that it can be mapped from natural language constructs by referring to the Entity-Relationship conceptual schema and can be efficiently transformed into relational algebra for query execution. The whole process provides a clear and natural framework for processing natural language queries to retrieve data from database systems.

The evolution of database systems tends to the development of higher degree of user-friendliness such that the system can be directly handled by nonprofessionals. In order to reach this goal, the database system needs to provide a query language by which queries can be specified conceptually. Also, the query condition may be relaxed such that information within a certain semantic distance to the exact answer can be obtained. Moreover, the real-world information is usually imprecise and incomplete. It is therefore important to store imprecise and incomplete information in a database, and to manipulate this information accordingly. We design a conceptual language by which *fuzzy query conditions* and *neighborhood query condi-*

tions can be specified. Query processing strategies for these two query conditions are proposed considering imprecise and incomplete information. A domain concept hierarchy is constructed on top of a numerical domain to handle imprecise data, while dependencies between database attributes are derived for incomplete information. These techniques can also be applied to process distributed queries under network partitioning.

[1] F.S.T. Tseng, A.L.P. Chen, W.P. Yang, "On Mapping Natural Language Constructs into Relational Algebra through E-R Representation," Data & Knowledge Engineering, North-Holland, pp. 97-118, October 1992.

[2] S.J. Yen and A.L.P. Chen, "Neighborhood/Conceptual Query Answering with Imprecise/Incomplete Data," Proc. Entity-Relationship Approaches, 1993.

5 Access Method

Principal Investigator: Prof. Y.H. Chin at NTHU
Support: National Science Council

Statistical and scientific databases are widely used in the applications such as decision making, data analyses, and actuarial studies. Compared with traditional databases, queries often inquire the database about descriptive statistics, order statistics, or sampling. TBSAM (Tree Based Statistics Access Method) provides an efficient way to respond to an aggregate query. However, TBSAM is designed for single-user environments. To improve its utilization, we propose a new structure called $TBSAM^{link}$ for processing aggregate queries in a multi-user environment. Several operations such as range reading, range updating, and sequential reading are provided by using the techniques of right-link, lock-coupling, and rechecking. Mathematical formulas for the number of nodes to be accessed and/or locked are derived to evaluate the efficiency of these concurrent operations. The results show that at most two nodes are required to be locked at any time for each update operation.

Similarly, we propose new algorithms to solve the concurrency control problems on R-trees. Without modifying the original data structure of R-trees, the algorithms are designed based on the locking strategy and the techniques of lock-coupling, driving-off, and side-branching. At most one node is required to be exclusively locked at any time for each update operation. The correctness and the deadlock-free property of the algorithms are proven. Besides, the access cost and lock efficiency are analyzed mathematically.

[1] F. Wu, J.K. Chen and Y.H. Chin, "$TBSAM^{link}$: A Structure for Concurrent Operations on Statistical Data," Journal of Information Science and Engineering (to appear).

[2] J.K. Chen, Y.F. Huang and Y.H. Chin, "A Study of Concurrent Operations on R-trees," NTHU-TR-CS-94-004.

6 Multidatabase Systems

Dr. Bao-Shuh P. Lin, Shu-Chin S. Chen, Mike Yu, Yen-Yao Yao, and Dr. San-Yih Hwang at ITRI

The Multidatabase Management System (MDBMS) developed at CCL seeks to provide an effective and efficient way for retrieving information from distributed, autonomous, and possibly heterogeneous databases. The MDBMS behaves like a complete DBMS from the user's/programmer's point of view.

The subgoal of this project is to develop a MDBMS prototype that can be easily commercialized by our industrial partners. Thus, this project covers a broad class of issues in MDBMS. It addresses the problems of transaction management, query processing/optimization, security enforcement, and catalog maintenance. In system catalog maintenance, we have found that the employment of existing storage managers to be very inefficient. Instead, we develop a specialized storage manager purely to meet the requirements of the MDBMS. This storage manager manages the MDBMS meta data and partial query execution results. The storage manager provides special interfaces to the MDBMS so that the MDBMS has closer control over the data managed by the storage manager. For example, the MDBMS is able to "pin" some catalog data to ensure that these data are most likely to stay in memory. Besides, the storage manager has transaction management ability to guarantee the consistency of the MDBMS system catalog.

This project also addresses the problem of MDBMS integrity constraints. Our approach is to better utilize the constraints enforced at each local site. The local constraints are converted to global constraints which are maintained at the MDBMS level. These global constraints give the MDBMS users a clear concept on why a query is rejected, avoid redundant local constraints checking, and can be utilized by the MDBMS to semantically optimize queries.

An MDBMS prototype is currently in operation. Most of the techniques we developed have been implemented in the prototype. To ease future production of our system, we adopt industrial standards in our design, e.g., ODBC for the interface between the MDBMS and its applications, Kerberos Server for security enforcement, and X/Open XA&TX for the interface between the MDBMS and the component database systems.

[1] A.L.P. Chen, K.F. Chieng, C.T. Kuo, J.D. Lee, and S.S. Chen, "Improving the Performance of a Distributed Computing System through Inconsistent Caches," Proc. IEEE Workshop on Future Trends of Distributed Computing Systems, 1992.

[2] J.L. Koh, W.F. Wang, N.J. Wu, S.S. Chen and B.P. Lin, "Accessing Remote Databases through a Uniform Interface," Proc. International Joint Workshop on Computer and Communication, 1993.

[3] S.S. Chen, M. Yu, Y.Y. Yao, and S.Y. Hwang, "MDBMS: Using DCE to Integrate Heterogeneous Databases," Proc. OSF DCE Developers Conference, 1994.

[4] N.J. Wu and A.L.P. Chen, "Improving Execution Concurrency for Long-Duration Database Transactions," Proc. IEEE COMPSAC, 1994.

[5] C.S. Chang, A.L.P. Chen, Y.Y. Yao, M. Yu, and S.S. Chen, "Deriving Global Constraints in Multidatabase Systems," NTHU-TR-CS-94-005.

7 Real-Time Database Systems

Dr. David Su, Shu-Chin S. Chen, Ying-Wei Hsu, and Dr. San-Yih Hwang at ITRI

This project aims to develop a database system for intelligent network applications. Intelligent network is a next generation telecommunication system that provides broader and flexible services to customers. The fundamental requirement of the database system is to be quick in responding to users' inquiries and transactions. To fulfill this requirement, all the data have to be placed in main memory. In addition, different classes of queries and transactions may have different requirements on their response time. How to dynamically schedule a set of requests with different timing constraints is an essential research issue of this project. Besides, as the intelligent network covers a wide area, the real-time database system has to be distributed geographically, with a centralized system maintaining the backup data. Thus algorithms have to be developed to maintain the consistency of data stored at different sites.

Although this project involves several active research areas, it is important to note that the data and operations managed by the real-time database system are highly specialized (i.e., they are used only by telecommunication applications). Thus, the general solutions proposed in research literature may be overkill. The services provided and the storage structure employed by the real-time database system are best specialized. For example, the replica consistency protocol is specially designed as the intelligent network applications only issue transactions with "blind write" (i.e., transactions contain only write operations). The backup disk as required by transaction management algorithms in main memory databases is also not required in our system, as the centralized system already stores the primary copies of data. Several other algorithms are tailor-made just to meet the requirements of intelligent network applications.

In addition to the above, we are currently investigating the design issues and proposing an architecture for a prototyping system that incorporates various techniques we have addressed.

8 Multimedia Database Systems

Dr. Bao-Shuh P. Lin, Dr. David Su, Shu-Chin S. Chen, and Lee-Chuan Chuang at ITRI

The goal of this project is to extend database technologies to support efficient and convenient manipulation of various kinds of non-traditional media data, e.g., image, video and audio. Our approach is to develop a database engine capable of processing multimedia data and to build a user-friendly environment to ease the development of multimedia applications. To support efficient processing of multimedia data, we aim to develop algorithms for efficient indexing on each individual media. Then we seek to propose a framework for expressing the relationships among different media data.

In the area of image indexing, we propose an algorithm to quickly find out a set of images that satisfy some spatial properties. This approach is similar to 2D string with significantly reduced string matching time. For video indexing, we are currently investigating an approach that stores the information about object movement across frames. The data organization adopted by this approach is able to efficiently respond to the queries concerning the movement of a particular object.

This project is still in its early stage. Nevertheless, it has resulted in a class library that provide a convenient way for multimedia applications to manage media data stored in databases. We also built a set of tools that enable users to access media data stored in different database management systems. Further enhancement will be conducted using the proposed approaches.

[1] S. Mao, L. Chuang and S. Wu, "An ODBC based Distributed Multimedia DBMS Architecture," Proc. International Joint Workshop on Computer and Communication, 1993.