DATABASE RESEARCH AT COLUMBIA UNIVERSITY

Shih-Fu Chang*, Luis Gravano, Gail E. Kaiser, Kenneth A. Ross, Salvatore J. Stolfo
Dept. of Computer Science, Columbia University, New York, NY 10027. http://www.cs.columbia.edu

1 Introduction

Columbia University has a number of projects that touch on database systems issues. In this report, we describe the Columbia Fast Query Project (Section 2), the JAM project (Section 3), the CARDGIS project (Section 4), the Columbia Internet Information Searching Project (Section 5), the Columbia Content-Based Visual Query project (Section 6), and projects associated with Columbia's Programming Systems Laboratory (Section 7).

2 The Columbia Fast Query Project[1]

Faculty: Ross.

The focus of the Columbia Fast Query Project is to process *complex* queries efficiently. Ideally, we aim for interactive query response. However, we also aim to improve the performance of noninteractive queries over huge datasets.

2.1 Complex OLAP Query Processing

In [1] we present the notion of "multi-feature queries." Multi-feature queries succinctly express complex queries such as "Find the total sales among minimum-price suppliers of each item." Such queries need multiple views and/or subqueries in standard SQL. We demonstrate significant performance improvements over a commercial system of a specialized evaluation algorithm for multi-feature queries.

In [2] we have developed techniques for recognizing when an arbitrary relational query is amenable to the following kind of evaluation strategy: (a) Partition the data according to some attributes, (b) apply a (simpler) query to each partition, and (c) union the results. Such evaluation strategies are particularly effective when the partitioned data is small enough to fit in memory. Our criteria for recognizing such programs are *syntactic*, and, surprisingly, are *nonrestrictive* in the sense that *every* query that can be evaluated in this partitioned fashion can be expressed in a way that satisfies our criterion.

In [3], we investigate techniques for evaluating correlated subqueries in SQL. Our techniques apply to all nested subqueries, unlike techniques such as query-rewriting that apply only to a limited class of subqueries. The basic idea is to cache the invariant part of the nested subquery between iterations, and to reevaluate just the variant part on each iteration. Integrating this technique into a cost-based optimizer required careful design. Our work was implemented in the Sybase IQ commercial database system, *and will be present in their next commercial release.*

In [4] we propose techniques for performing a join of two large relations using a join index. Our techniques are novel in that they require only a single pass of each participating relation, even if both tables are much larger than main memory, with all intermediate I/O performed on tuple identifiers. This work is extended in [5] to perform a self-join of a large relation with a single pass through the relation.

A datacube operation computes an aggregate over many different sets of grouping attributes. In particular, for d attributes chosen as possible "dimensions," 2^d aggregates are computed, one for each possible set of dimension attributes. In practice, data is often *sparse* in the d-dimensional space. We have developed techniques that are particularly efficient for computing the datacube of large, sparse datasets [6].

We have also developed novel algorithms to evaluate complex datacube queries involving multiple dependent aggregates at each granularity [7]. Previous techniques asked only simple aggregates (such as sum, max, etc.) at each granularity.

2.2 Materialized Views

Materialized views precompute and store expressions that can be used as subexpressions of complex queries. By doing work in advance, one can speed up the process of answering interactive queries.

An important practical question is choosing which views to materialize. In [8] we address this problem, demonstrating that it sometimes pays to materialize additional views just to support the maintenance of a given materialized view.

In [9] we look at the problem of *view adaptation*, namely how can one incrementally modify a material-

*Dept. of Electrical Engineering
[1] http://www.cs.columbia.edu/~kar/fastqueryproj.html

ized view after a change in the view definition.

When one queries multiple materialized views and base tables one would like to see a single consistent database state. A more general notion of serializability for accessing both base data and materialized views is presented in [10].

Different maintenance policies might be used for performance reasons. For example, deferred maintenance leads to fast base updates but adds overhead to queries; immediate maintenance does the opposite. In [11] we describe how to combine various policies within a single system, and measure the performance of various alternatives.

Maintenance of nested relations within materialized views is described in [12].

2.3 Data Reduction and Visualization

In huge data warehouses it often makes sense to summarize a dataset in order to reduce its size. An approximate answer to a query computed using the summary may be feasible when an exact answer using the full dataset may be infeasible. A survey of such data reduction techniques is presented in [13].

Techniques for visualizing large multidimensional datasets are presented in [14]. Our techniques enable one to visually identify, for any pair of dimensions, regions where the two-dimensional distribution is not explainable as the independent combination of one-dimensional distributions.

3 JAM[2]

Faculty: Stolfo.

We address the performance problem associated with Knowledge Discovery in Databases (or data mining) over inherently distributed databases. We study approaches to substantially increase the amount of data a knowledge discovery system can handle effectively. Meta-learning is a general technique we have developed to integrate a number of distinct learning processes executed in a parallel and distributed computing environment. Many other researchers have likewise invented ingenious methods for integrating multiple learned models as exemplified by stacking, boosting, bagging, weighted voting of mixtures of experts, etc. Several meta-learning strategies we have proposed have been shown to process massive amounts of data that main-memory-based learning algorithms cannot efficiently handle.

We collaborate with the Financial Services Technology Consortium in applying our prototype systems to large problems of interest in the financial and banking industry. One such problem is fraud detection.

Fraudulent electronic transactions are a significant problem, one that will grow in importance as the number of access points in the nation's financial information system grows. Financial institutions today typically develop custom fraud detection systems targeted to their own asset bases. Recently though, banks have come to realize that a unified, global approach is required, involving the periodic sharing with each other of information about attacks.

We propose another wall to protect the nation's financial systems from threats. This new wall of protection consists of pattern-directed inference systems using models of anomalous or errant transaction behaviors to forewarn of impending threats. This approach requires analysis of large distributed databases of information about transaction behaviors to produce models of "probably fraudulent" transactions.

The key difficulties in this approach are: financial companies don't share their data for a number of (competitive and legal) reasons; the databases that companies maintain on transaction behavior are huge and growing rapidly; real-time analysis is desirable to update models when new events are detected and distribution of models in a networked environment is essential to maintain up to date detection capability.

We propose a novel system to address these difficulties and thereby protect against electronic fraud. Our system has two key component technologies: local fraud detection agents that learn how to detect fraud and provide intrusion detection services within a single corporate information system, and a secure, integrated meta-learning system that combines the collective knowledge acquired by individual local agents.

Once derived local classifier agents or models are produced at some site(s), two or more such agents may be composed into a new classifier agent by a meta-learning agent. The meta-learning system proposed will allow financial institutions to share their models of fraudulent transactions by exchanging classifier agents in a secured agent infrastructure. But they will not need to disclose their proprietary data.

Our publications include [15, 16, 17, 18, 19, 20, 21].

4 CARDGIS[3]

Faculty: Gravano, Kaiser, Ross, Stolfo, and others.

CARDGIS is the acronym for a newly-formed center: The USC/ISI and Columbia University Center for Applied Research in Digital Government Information Systems. The center's mission is research in the

[2] `http://www.cs.columbia.edu/~sal/JAM/PROJECT/`

[3] `http://www.cs.columbia.edu/~sal/DIGGOV/census/index.htm`

design and development of advanced information systems with capabilities for information workers and the general public to individually and cooperatively generate, share, interact with and effectively utilize knowledge in a national-scale statistical digital library. The center's goals include demonstrating the value of such knowledge sharing when applied to vast and continually growing amounts of statistical data. This will be achieved through the development, deployment and evolution of pilot systems in collaboration with several Federal Statistical agencies, among them the U.S. Bureau of the Census and the Bureau of Labor Statistics, and their broad user communities. Both the Census Bureau and the Bureau of Labor Statistics have committed to participate in and collaborate with CARDGIS research projects.

This project brings together researchers in databases, collaboration technologies and tools, human-computer interaction, knowledge representation, knowledge discovery and data mining, WWW technology, security, and software engineering, drawn from Columbia and USC/ISI, along with social and statistical science experts from Federal, State and Local Statistical agencies. These agencies routinely collect large amounts of statistical data. The project will develop technology that will build upon and incorporate existing systems, including the integration of legacy databases. It will support the development of new information systems with features that provide for sophisticated viewing, analysis and manipulation of data by statisticians, sociologists, policy makers, teachers, students, and the general public, while maintaining data security; secure collaboration among dynamic groups of users intending to achieve purposeful goals such as policy development; and dynamic integration of statistical data with more traditional digital libraries and other information resources in secured application- specific contexts.

CARDGIS research will proceed initially along six key coordinated thrusts: Ontology development and merging; integration of information distributed over multiple, heterogeneous legacy sources; human-computer interaction; collaboration technologies to assist dynamic groups of distributed users; secured infrastructures to maintain confidentiality; and distributed data mining and analysis for large statistical sources. A report describing the challenges of moving towards a digital government is available [22].

5 Internet Information Searching[4]

Faculty: Gravano.

[4] `http://www.cs.columbia.edu/~gravano`

Increasingly, users want to issue complex queries across Internet sources to obtain the data they require. Because of the size of the Internet, it is not possible to process such queries in naive ways, e.g., by accessing all the available sources. Thus, we must process queries in a way that scales with the number of sources. Also, sources vary in the type of information objects they contain and in the interface they present to their users. Some sources contain text documents and support simple query models where a query is just a list of keywords. Other sources contain more structured data and provide query interfaces in the style of relational model interfaces. User queries might require accessing sources supporting radically different interfaces and query models. Thus, we must process queries in a way that deals with heterogeneous sources.

Our past research focused on processing queries over static document sources. As part of this research, we implemented *GlOSS*, a scalable system that helps users choose the best sources for their queries [23, 24, 25]. One key goal of *GlOSS* is to scale with large numbers of document sources, for which it keeps succinct descriptions of the source contents. *GlOSS* needs to extract these content summaries from the sources. Unfortunately, the source contents are often hidden behind search interfaces that vary from source to source. In some cases, even the same organization uses search engines from different vendors to index different internal document collections. To facilitate accessing multiple document sources by making their interfaces more uniform, we coordinated an informal standards effort involving the major search engine vendors and users of their technology. The result of this effort, *STARTS* (Stanford Protocol Proposal for Internet Retrieval and Search), includes a description of the content summaries that each source should export so that systems like *GlOSS* can operate [26].

Sources on the Internet offer many types of information: dynamic documents generated on the fly, structured, non-text information with sophisticated query interfaces, resource directories specialized in a certain topic, images, video. To complicate matters further, many Internet sources provide little more than a query interface, and do not export any content summaries or information about their capabilities. Searching for information across general heterogeneous Internet sources thus presents many challenging new problems: from characterizing sources with novel information types, for efficient query processing, to indexing non-traditional resource properties, for helping users get what they really want. We are developing tools to allow users to exploit all the information available on the Internet in meaningful ways, by building on work from several fields, most notably from the database, information retrieval, and natural language processing fields. Combining technologies from different areas is crucial to achieve the goal of letting users access the information that they want easily, without forcing them to be aware of the different interfaces, formats, and models used to represent this information at the originating sources.

6 Content-Based Visual Queries[5]

Faculty: Chang.

The amount of visual information on-line, including images, videos, and graphics, is increasing at a rapid pace. Textual annotations provide direct means for visual searching and retrieval, but are not sufficient. Visual features of the images and video provide additional descriptions of their content. By exploring the synergy between textual and visual features, image search systems can be greatly improved.

There has been substantial progress in developing powerful tools which allow users to specify image queries by giving examples, drawing sketches, selecting visual features (e.g., color and texture), and arranging spatial-temporal structure of features [27]. Much success has been achieved, particularly in specific domains, such as sports, remote sensing, and medical applications.

New challenges remain in applying the above content-based image search tools to meet real user needs. Our experience indicates that use of the image search systems varies greatly, depending on the user background and the task. One unique aspect of image search systems is the active role played by users. By modeling the users and learning from them in the search process, we can better adapt to the users' subjectivity. We describe our current research directions in the following.

Although today's computer vision systems cannot recognize high-level objects in unconstrained images, we are finding that low-level visual features can be used to partially characterize image content. Local region features (such as color, texture, face, contour, motion) and their spatial/temporal relationships are being used in several innovative image searching systems [28, 29]. We argue that the automated segmentation of images/video objects does not need to accurately identify real world objects contained in the imagery. One objective is to extract "salient" visual features and index them with efficient data structures for powerful querying. To link the features to high-level semantics, we are also investigating (semi-)automated systems for extraction and recognition of semantic objects (e.g., people, car) using user-defined models.

Ideal image representations should be amenable to dynamic feature extraction and indexing. Today's compression standards (such as JPEG, MPEG-2), are not suited to this need. The objective in the design of these compression standards was to reduce bandwidth and increase subjective quality. Although many interesting analysis and manipulation tasks can still be achieved in today's compression formats [30], the potential functionalities of image search and retrieval were not considered. Recent trends in compression, such as MPEG-4 and object-based video, have shown interest and promise in this direction. One of our goals is to develop systems in which the video objects are extracted, then encoded, transmitted, manipulated, and indexed flexibly with efficient adaptation to users' preference and system conditions.

Another direction is to break the barrier of decoding image semantic content by using user-interaction and domain knowledge. These systems learn from the users' input as to how the low-level visual features are to be used in the matching of images at the semantic level. For example, the system may model the cases in which low-level feature search tools are successful in finding the images with the desired semantic content. We have developed a unique concept called Semantic Visual Templates [31], which use a two-way interaction system to find a samll subset of graphic icons representing the semantic concepts (e.g., sunsets or high jumpers). In an environment including distributed, federated search engines, a meta-search system may monitor user's preference of different feature models and search tools and then recommend search options to help users improve the search efficiency [32].

Exploring the association of visual features with other multimedia features, such as text, speech, and audio, provides another fruitful direction. Video often has text transcripts and audio that may also be analyzed, indexed, and searched. Images on the World Wide Web typically have text associated with them. In [33], we have shown significant performance improvement in searching Web visual content by using integrated multimedia features.

7 The Programming Systems Lab[6]

Faculty: Kaiser.

The Programming Systems Lab at Columbia University is currently working on two database-oriented projects: Pern and Xanth.

[5] `http://www.ctr.columbia.edu/~sfchang/vis-project`

[6] `http://www.psl.cs.columbia.edu/`

7.1 Pern

The Pern project is concerned with programmable transaction managers that can realize a wide range of human-oriented extended transaction models (ETMs) and, more recently, with application of such transaction management services in conjunction with emerging distributed computing protocol standards such as CORBA and HTTP.

An earlier prototype, called simply Pern [34], operated as a component (written in C) that can be linked into or externally added on to a system that requires either the conventional ACID properties or a special-purpose extended transaction model. Pern's lock conflict table is replaceable and its standard transaction operations (begin, commit, abort, lock, unlock) can be associated with before and after plugins to modify or augment the semantics of the operations, by default implemented through the usual two-phase locking. Pern has been added on externally to Cap Gemini's ProcessWeaver workflow product and linked into our own Oz software development environment, and has also been used in tandem with the conventional transaction manager built into GIE Emeraude's PCTE environment framework product. The Cord [35] extended transaction modeling language and corresponding interpreter provides means for defining rules for resolving concurrency control conflicts as they arise. It has been used to implement altruistic locking and epsilon serializability, and integrated with U. Wisconsin's Exodus as well as with Pern. The current version of Pern, called JPernLite [36], has been redesigned and is now implemented in Java. We are currently developing JCord, which will expand on Cord to support dynamic construction and modification of ETMs on the fly from rules received from clients and negotiation of conflicts among ETMs for different applications coincidentally accessing the same data served by the same or different instances of JPernLite.

7.2 Xanth

Xanth addresses integration of heterogeneous information resources from the viewpoint of accomplishing useful work, particularly in a team context. Our goal is to develop a middleware infrastructure for group information spaces, or *groupspaces*, in which workflow and transaction collaboration services, and the distributed query, information retrieval, data mining, etc. services being developed by other Columbia database projects, can operate.

The Xanth prototype in particular operates as a *middlebase* [37], imposing hypermedia object management on top of external data residing in any combination of conventional databases, websites, proprietary repositories, etc. dispersed over the Internet or organizational intranet. One key concept of a middlebase is that not only are the information repositories heterogeneous, but so are the protocols used to access these repositories. Further, the presentation clients and protocols for communicating with them may be similarly heterogeneous, and legacy, like the data sources constructed without the potential for sharable information spaces in mind.

The Xanth project is currently working on *groupviews*, team-oriented user interface metaphors for collaboration environments like those supported by Xanth. One experiment combines VRML with MUDs, with information subspaces mapped to platforms in the New York City subway system and movement between subspaces mapped to trips on the subway trains.

References

[1] D. Chatziantoniou and K. A. Ross. Querying multiple features of groups in relational databases. In *Proceedings of the International Conference on Very Large Databases*, pages 295–306, 1996.

[2] D. Chatziantoniou and K. A. Ross. Groupwise processing of relational queries. In *VLDB Conference*, pages 476–485, 1997.

[3] J. Rao and K. A. Ross. Reusing invariants: A new strategy for correlated queries. In *ACM SIGMOD Conference on Management of Data*, June 1998.

[4] Z. Li and K. Ross. Fast joins using join indices. *VLDB Journal*, 1998. (to appear).

[5] H. Lei and K. A. Ross. Faster joins, self-joins and multi-way joins using join indices. In *Proceedings of the Workshop on Next Generation Information Technologies and Systems*, pages 89–101, 1997. Full version to appear in Data and Knowledge Engineering.

[6] K. A. Ross and D. Srivastava. Fast computation of sparse datacubes. In *International Conference on Very Large Databases*, pages 116–125, August 1997.

[7] K. A. Ross, D. Srivastava, and D. Chatziantoniou. Complex aggregation at multiple granularities. In *Extending Database Technology*, pages 263–277, 1998.

[8] K. A. Ross, D. Srivastava, and S. Sudarshan. Materialized view maintenance and integrity constraint checking: Trading space for time. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 447–458, June 1996.

[9] A. Gupta, I. S. Mumick, and K. A. Ross. Adapting materialized views after redefinitions. In *Proceedings of the ACM-SIGMOD International Conference on Management of Data*, pages 211–222, 1995.

[10] A. Kawaguchi, D. Lieuwen, I. S. Mumick, D. Quass, and K. A. Ross. Concurrency control theory for deferred materialized views. In *Proceedings of the Inter-*

national Conference on Database Theory, pages 306–320, 1997.

[11] L. Colby, A. Kawaguchi, D. Lieuwen, I. S. Mumick, and K. A. Ross. Supporting multiple view maintenance policies. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 405–416, May 1997.

[12] A. Kawaguchi, D. Lieuwen, I. S. Mumick, and K. A. Ross. Implementing incremental view maintenance in nested data models. In *Proceedings of the Workshop on Database Programming Languages*, 1997.

[13] D. Barbara et al. The New Jersey data reduction report. *Data Engineering Bulletin*, 20(4):3–45, 1997.

[14] S. Berchtold, H. V. Jagadish, and K. A. Ross. Independence diagrams: A technique for visual data mining. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*, 1998.

[15] W. Lee, S. J. Stolfo, and K. Mok. Mining audit data to build intrusion detection models. In *KDD98*, 1998.

[16] P. K. Chan and S. J. Stolfo. Toward scalable learning with non-uniform class and cost distributions: A case study in credit card fraud detection. In *KDD98*, 1998.

[17] A. L. Prodromidis and S. J. Stolfo. Mining databases with different schemas: Integrating incompatible classifiers. In *KDD98*, 1998.

[18] W. Lee and S. Stolfo. Data mining approaches for intrusion detection. In *Proceedings 7th USENIX Security Symposium*, 1998.

[19] S. Stolfo, A. Prodromidis, S. Tselepis, W. Lee, D. Fan, and P. Chan. Jam: Java agents for meta-learning over distributed databases. In *KDD97 and AAAI97 Workshop on AI Methods in Fraud and Risk Management*, 1997.

[20] P. Chan and S. Stolfo. Sharing learned models among remote database partitions by local meta-learning. In *KDD96*, 1996.

[21] P. Chan and S. Stolfo. On the accuracy of meta-learning for scalable data mining. *Journal of Intelligent Integration of Information*.

[22] H. Schorr and S. Stolfo. Towards the digital government of the 21st century. *Report of the Workshop on Research and Development Opportunities in Federal Information Services*, 1997. `http://www.isi.edu/nsf/final.html`.

[23] L. Gravano, H. García-Molina, and A. Tomasic. The effectiveness of *GlOSS* for the text-database discovery problem. In *Proceedings of the 1994 ACM International Conference on Management of Data (SIGMOD'94)*, May 1994.

[24] L. Gravano and H. García-Molina. Generalizing *GlOSS* for vector-space databases and broker hierarchies. In *Proceedings of the Twenty-third International Conference on Very Large Databases (VLDB'97)*, pages 78–89, September 1995.

[25] L. Gravano, H. García-Molina, and Anthony Tomasic. Precision and recall of *GlOSS* estimators for database discovery. In *Proceedings of the Third International Conference on Parallel and Distributed Information Systems (PDIS'94)*, September 1994.

[26] L. Gravano, C. K. Chang, H. García-Molina, and A. Paepcke. STARTS: Stanford proposal for Internet meta-searching. In *Proceedings of the 1997 ACM International Conference on Management of Data (SIGMOD'97)*, May 1997.

[27] S.-F. Chang, J. R. Smith, M. Beigi, and A. Benitez. Visual information retrieval from large distributed online repositories. *Communications of ACM, Special Issue on Visual Information Management*, 40(12):63–71, Dec. 1997.

[28] S.-F. Chang, W. Chen, H.J. Meng, H. Sundaram, and D. Zhong. Videoq- an automatic content-based video search system using visual cues. In *ACM Multimedia Conference*, 1997. Demo: `http://www.ctr.columbia.edu/VideoQ`.

[29] J. R. Smith and S.-F. Chang. Visualseek: A fully automated content-based image query system. In *ACM Multimedia Conference*, 1996. Demo: `http://www.ctr.columbia.edu/VisualSEEk`.

[30] J. Meng and S.-F. Chang. Cveps: A compressed video editing and parsing system. In *ACM Multimedia Conference*, 1996. Demo: `http://www.ctr.columbia.edu/WebClip`.

[31] S.-F. Chang, W. Chen, and H. Sundaram. Semantic visual templates - linking visual features to semantics. In *IEEE International Conference on Image Processing*, 1998.

[32] M. Beigi, A. Benitez, , and S.-F. Chang. Metaseek: A content-based meta search engine for images. In *SPIE Conference on Storage and Retrieval for Image and Video Databases*, 1998. Demo: `http://www.ctr.columbia.edu/metaseek`.

[33] J. R. Smith and S.-F. Chang. Visually searching the web for content. *IEEE Multimedia Magazine*, 4(3):12–20, 1997. Demo: `http://www.ctr.columbia.edu/webseek`.

[34] G. T. Heineman and G. E. Kaiser. An architecture for integrating concurrency control into environment frameworks. In *17th International Conference on Software Engineering*, 1995.

[35] G. T. Heineman and G. E. Kaiser. The CORD approach to extensible concurrency control. In *13th International Conference on Data Engineering*, 1997.

[36] J. J. Yang and G. E. Kaiser. Jpernlite: An extensible transaction server for the world wide web. In *9th ACM Conference on Hypermedia and Hypertext*, 1998.

[37] G. E. Kaiser and S. E. Dossick. Workgroup middleware for distributed projects. In *IEEE 7th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, 1998.