

The Indiana Center for Database Systems at Purdue University

Mourad Ouzzani Walid G. Aref Elisa Bertino Ann Christine Catlin
Christopher W. Clifton Wing-Kai Hon Ahmed K. Elmagarmid Arif Ghafoor
Susanne E. Hambrusch Sunil Prabhakar Jeffrey S. Vitter Xiang Zhang
Web Site: <http://www.cs.purdue.edu/icds> Purdue University, West Lafayette IN

1 Introduction

The Indiana Center for Database Systems (ICDS) at Purdue University has embarked in an ambitious endeavor to become a premiere world-class database research center. This goal is substantiated by the diversity of its research topics, the large and diverse funding base, and the steady publication trend in top conferences and journals. ICDS was founded with an initial grant from the State of Indiana Corporation of Science and Technology in 1990. Since then it has grown to now have 9 faculty members and about 30 total researchers. This report describes the major research projects underway at ICDS as well as efforts to move research toward practice.

2 Multimedia

Several projects have been launched to overcome challenges in organizing, storing, mining, and delivering multimedia data in a secure and timely manner.

2.1 Video Database Management

The vdbms Project [9, 3] has developed a video-enhanced database system that supports comprehensive and efficient database management, including: feature-based preprocessing for video content representation and indexing, video and meta-data storage, video query processing, buffer and storage management, and continuous video streaming. In this project, we view video as a well-defined data type with its own description, parameters, and applicable methods. Supporting video operations (storing, searching by content, and streaming) and new query types (query by example and multi-feature similarity search) requires major changes in many of the traditional system components. Specifically, the storage and buffer manager will have to deal with huge volumes of data with real-time constraints. Query processing has to consider the video methods and operators in generating, optimizing and executing query plans.

2.2 Semantic Video Databases

Several content-based video retrieval systems have been proposed in the past, but they still suffer from

the following challenging problems: semantic gap, semantic video concept modeling, semantic video classification, and concept-oriented video database indexing and access. We propose a framework [8] that includes semantic-sensitive video content representation, semantic video concept interpretation, a novel semantic video-classifier training framework, and a concept-oriented video database organization technique to enable semantic-sensitive video retrieval and browsing.

2.3 Video Mining

To achieve efficient video indexing and access, we introduce strategies for video content structure and event mining [26]. The video shot segmentation and representative frame selection strategies are first utilized to parse the continuous video streams into physical units. Video shot grouping, group merging, and scene clustering schemes are then proposed to organize the video shots into a hierarchical structure using clustered scenes, scenes, groups, and shots, in increasing granularity from top to bottom. Audio and video processing techniques are then integrated to mine event information from the detected scenes. Finally, the acquired video content structure and events are integrated to construct a scalable video skimming tool.

2.4 Distributed Multimedia Systems

A large number of emerging Web-based applications require Distributed Multimedia Systems (DMIS) infrastructures. Examples of such applications abound in the domains of medicine, manufacturing, and e-commerce. Development of DMIS needs a broad range of technological solutions for organizing, storing, and delivering multimedia information in an integrated, secure and timely manner with guaranteed end-to-end quality of presentation (QoP). Management of integrated end-to-end QoP and ensuring information security in DMIS present formidable research and implementation challenges. These challenges encompass all the sub-system components of a DMIS. The ultimate objective relies on the performance and allocation of resources of each of the DMIS sub-system components including networks, databases, and end-systems. Several projects addressing these challenges are being pursued.

3 Security and Privacy

We have initiated several projects to address security and privacy issues in database systems, data mining, and data sharing.

3.1 Privacy-preserving DBMSs

This project investigates ways to extend current DBMS architectures, models and languages to support high-assurance privacy and achieve fine-grained access control. A first result has been the development of a model for labeling data with purpose information [4]. Our approach has several key features: (1) hierarchical organization of purposes according to an ontology, (2) support for positive and negative purposes, (3) support for the association of sets of purposes with a single data item, and (4) support for a variable range of purpose labeling granularities. The model has been implemented through the use of query modification techniques. The approach, initially developed for relational data, has since been extended to complex data, such as XML, and integrated with a role-based access control model [25]. We are currently investigating innovative access control techniques for relational database systems, to be able to enforce accesses to data at a very fine granularity level. In this context, we have proposed the innovative notion of micro-views, which allows one to fine-tune the values returned by queries.

3.2 Privacy-preserving Data Mining

Data mining relies on the collection of massive amounts of data - but this often collides with privacy considerations [18]. We propose to address this challenge in the distributed case. The solutions involve algorithms that share some information to calculate correct results where the shared information can be shown not to disclose private data. Our contribution in this context has been manifold: (1) We propose a toolkit of components that can be combined for specific privacy-preserving data mining applications. (2) We propose secure mining of association rules over horizontally partitioned data [17]. The methods incorporate cryptographic techniques to minimize the information shared, while adding little overhead to the mining task. (3) We propose a method for k-means clustering [23] when different sites contain different attributes for a common set of entities. (4) We propose a method to apply classification rules without revealing either the data or the rules. In addition, the rules can be verified not to use any “forbidden” criteria. A key strength of ICDS in this area is its collaboration with the Krannert School of Management to identify problems and develop solutions where incentives discourage cheating.

3.3 Privacy-preserving Data Sharing

Integrating and sharing data from multiple sources has been a long-standing challenge. Data integration is seriously hampered by an inability to ensure privacy.

Problems include fear of disclosing confidential information as well as regulations protecting individual privacy. This project proposes to develop the technology needed to create and manage federated databases while controlling the disclosure of private data [7]. To address the above problems, we plan to develop solutions to several fundamental problems in privacy-preserving data integration and sharing, including (1) a privacy framework that is flexible and clear to the end users, (2) schema-matching techniques that expose the source data and schemas, and (3) querying across sources while ensuring that results do not violate privacy policy, disclosure of sources, and preventing private information leakage from answering a set of queries.

3.4 Watermarking

By enabling relatively cost-free, fast, and accurate access channels to information in digital form, computers have radically changed the way we think and express ideas. As increasingly more information is produced, packaged and delivered in digital form, one of its main features threatens to become its worst enemy: zero-cost verbatim copies. Digital Watermarking deploys Information Hiding as a method of Rights Protection to conceal an indelible “rights witness” (watermark) within the digital Work to be protected. We analyze digital watermarking from a higher level, domain-independent perspective [22]. We propose a theoretical model and ask: are there any limitations to what watermarking can do? What are these and when can they be reached? We then propose, design and analyze watermarking solutions for (1) numeric sets, (2) numeric relational data, (3) categorical data, (4) streams and (5) semi-structured data.

3.5 Untrusted Private Databases

Entities such as corporations often use databases to store confidential operational data. There is, however, a need to share this data with outside entities in a limited fashion. External entities may not fully trust data owners to share private data without malicious modification. We therefore need to enforce constraints on a private database without having access to the database. We are developing protocols to help a third party enforce constraints over an untrusted private database and obtain a guarantee that the constraints are enforced. One particular application of this research is privacy-preserving query result verification. In this model, we assume that the database periodically commits itself without revealing its contents. Subsequent queries are guaranteed to return correct results with respect to the committed state of the database. In addition to the correctness of the protocol, we are concerned about the overhead of the solution and the degree of exposure of private data in order to guarantee correctness.

4 Top-k Query Processing

Ranking queries produce results that are ordered on some computed score. Typically, in these queries, users are usually interested only in the top-k results. Current relational query processors do not handle ranking queries efficiently, especially when joins are involved. We developed rank-join algorithms [15, 14] that use the individual orders of its inputs to produce join results ordered on a user-specified scoring function. The idea is to rank the join results progressively during the join operation. We introduce physical query operators based on variants of ripple join.

Rank-join operators progressively rank the join results while performing the join operation. We propose a rank-aware query optimization framework [16] that fully integrates rank-join operators into relational query engines. The framework is based on extending the System R dynamic programming algorithm in both enumeration and pruning. Unlike traditional join operators, optimizing for rank-join operators depends on estimating the input cardinality of these operators. We introduce a probabilistic model for estimating the input cardinality, and hence the cost of a rank-join operator.

5 Data Streaming

The number of applications dealing with streams of data produced in an unpredictable and bursty fashion is growing in several areas including network traffic management, high-throughput instruments, and sensors. Nile [13] is a data stream management system under development at Purdue. Nile supports extended SQL operators that handle sliding-window execution. Specifically, Nile supports the efficient and correct pipelined execution of sliding window queries over multiple data streams. The correct execution is enforced via the Negative Tuple Approach [13]. Negative tuples are tuples that are generated whenever a tuple expires out of a sliding window. Each operator in the pipeline needs to react differently whenever it receives a negative tuple to counteract the effect of the corresponding positive tuple that just expired out of the window. Although negative tuples guarantee correct execution of query pipelines, they induce performance overhead. We are looking into optimization techniques to reduce the performance overhead induced by negative tuples. Another interesting feature of Nile is its predicate windows. In contrast to sliding windows that limit the focus of queries on streams to the most recent tuples, predicate windows can select tuples of interest that meet a certain select or join predicate.

We are currently working on several other extensions to Nile: (1) Shared execution of concurrent continuous queries over data streams [12]. (2) A summary manager based on the concept of *promising tuples* – we define shared summaries over data streams enabling the

query engine to avoid looking at the original streams. Summaries are built using promising tuples which represent tuples that are “judged” important. (3) Context and profile-aware query processing – we are investigating context and profile awareness at the engine level with a generic extensible interface to define new contexts and modules for context-aware processing within the engine. (4) Native online data mining – Nile will have embedded online data mining operators that would operate on incoming streams to discover “interesting phenomena”. This capability raises several issues like change detection in the streams and the need to cluster streams based on their behavior.

6 Spatio-Temporal Databases

The tremendous increase of cellular phones, GPS-like devices, and RFIDs results in highly dynamic environments where objects and queries are continuously moving. We are investigating issues related to spatio-temporal queries and location-aware services. The PLACE (Pervasive Location Aware Computing Environments) project focuses on the efficient support of location-aware services. The PLACE server extends data streaming management systems to support location-aware environments. The query processor includes: (1) new incremental spatio-temporal operators, (2) extended semantics for sliding window queries, and (3) a shared execution framework for scalable execution.

We propose two algorithms: SINA [20] for evaluating concurrent continuous spatio-temporal queries and SEA-CNN [24] for answering a collection of continuous concurrent k-nearest neighbor queries (CKNN). SINA achieves scalability by employing a shared execution paradigm where the execution of continuous spatio-temporal queries is abstracted as a spatial join between a set of moving objects and a set of moving queries. SINA employs an incremental evaluation that is achieved by computing only the updates of the previously reported answer. Experimental results show that SINA is scalable and is more efficient than other indexed spatio-temporal algorithms. SEA-CNN has two important features: incremental evaluation and shared execution. SEA-CNN achieves both efficiency and scalability in the presence of a set of concurrent queries. Furthermore, SEA-CNN does not make any assumption about the movement of objects or the mutability of the objects or the queries. Experimentation shows that SEA-CNN is highly scalable and is more efficient than other R-tree-based CKNN techniques.

7 Managing Data Uncertainty

Due to limited bandwidth and battery power, it is infeasible for a system to keep track of the actual values continuously produced by sensors. Database queries, in this case, may produce incorrect answers. We model

the uncertainty inherent to dynamic sensor data with a range of possible values together with the probability distribution of the values within that range [6]. We also propose probabilistic queries, which evaluate uncertain data and produce answers with probabilistic guarantees. A query classification scheme is proposed, where for each class, query evaluation algorithms and answer quality are presented. We also study the semantics of probabilistic comparison operators that return imprecise comparison results. We illustrate how uncertainty management techniques for sensor data can be extended to location data in moving-object databases. We investigate the efficiency issues of probabilistic query evaluation. In particular, we propose efficient algorithms to enhance the performance of nearest-neighbor queries, range queries, and joins. We are currently developing a prototype for handling uncertain data and probabilistic queries. Uncertain data types are treated as a first-class data type and an extended version of SQL allows probabilistic queries to be specified.

8 Indexing Techniques

8.1 Space Partitioning Trees

Emerging database applications require the use of new indexing structures beyond B-trees and R-trees. Examples are the k-D tree, the trie, the quadtree, and their variants. A common feature of all these indexes is that they recursively divide the space into partitions. We developed an extensible index structure, called SP-GiST [1, 2, 10], to support the class of space partitioning un-balanced trees. Simple method implementations are provided demonstrating how SP-GiST can behave as a k-D tree, a trie, a quadtree, or any of their variants. Issues related to clustering tree nodes into pages as well as concurrency control for SP-GiST are addressed. A dynamic minimum-height clustering technique is applied to minimize disk accesses and to make using such trees in database systems possible and efficient. SP-GiST is supported both in PostgreSQL and in a standalone version.

8.2 Indexing in High-Update Environments

Emerging applications such as sensor networks and location-based services require monitoring continuous changing data. In high update environments, traditional database indexes suffer from the need for frequent updates and result in poor performance. We propose new index structures that are tolerant to constant data updates and optimized to achieve the best overall performance for both querying and updating [6, 5].

For moving object data indexing, we proposed the Change-tolerant Rtree (CTRtree). We observe that objects often stay in a region (e.g., building) for an extended amount of time, and exploit this phenomenon

to optimize an index for both updates and queries. We have developed the algorithms for creation and use of change tolerant Rtree. Experimental results establish the superior performance of the proposed index structures.

For sensor data indexing, we proposed the Mean Variance Tree (MVTtree), which is built based on the mean and variance of the data instead of the actual data values that are in constant flux. We also proposed to take each constantly evolving data as a time series and use the time series methodology to analyze and model it. The time series model enables effective forecasts for the constantly evolving data. Based on the forecasted intervals, we developed the Forecasted Interval Index (FI-index). Experiments show that, compared to traditional indexing schemes, both the MVTtree and the FI-Index substantially improve index updating performance while maintaining satisfactory query performance.

9 Efficient I/O for Data Intensive Applications

Recent trends in hardware development and data-intensive applications have resulted in a performance bottleneck for storing and retrieving data. The I/O communication between the fast internal memory and the slow external memory (such as disk) can be a bottleneck when processing massive amounts of data. The goal of this project is to develop techniques to alleviate the I/O bottleneck. One of the activities of the project has centered on the development of data placement, migration, and indexing techniques for video and multi-dimensional spatio-temporal data.

Problems involving massive amounts of geometric data are ubiquitous in spatial databases, geographic information systems, and virtual reality systems. For example, the NASA's Earth Observing System project produces petabytes of raster data per year. A major challenge is to develop mechanisms for processing the data, or else much of the data will be useless. In our work, we concentrate on the design and analysis of external memory data structures for batched and online problems involving geometric and spatial data. Many problems on geometric objects can be reduced to a small core of problems, such as computing intersections, convex hulls, and nearest neighbor search, for which we discuss useful paradigms for solving them in external memory.

10 Compressed Data Structures

10.1 Compressed Indexes for Text

This project targets the efficient indexing of massive documents. When a text T (of length n) is compressible, a natural question to ask is whether the index

can be compressed by a similar amount and still support fast searching of any substring in T . In addition, we would like to create indexes that implicitly encode the text as part of the indexing structure. Achieving this goal means that the text is no longer needed alongside the index. We also intend to support fast decompression starting from any position in the text and to support searching for arbitrary patterns in the text *without* scanning the entire compressed text. We develop a data structure that operates in nearly optimal space in terms of the entropy of the text T , while still supporting searches in $O(m \log |\Sigma| + \text{polylog}(n))$ time, where m is the length of the pattern being searched and $|\Sigma|$ is the size of the alphabet for text T . We also show several interesting tradeoffs and an adaptation of our data structuring technique for use only in compression.

10.2 Data-Aware Indexes for Set Data

In this project [11], we studied a way for representing a set S of items from $[1, n]$. Our target is to minimize the space usage, while trying to support ‘nearly-optimal’ time for queries on S . Our queries include: (1) $\text{rank}(i)$: how many items in S that is less than i ? (2) $\text{select}(j)$: what is the j -th smallest item in S ? and (3) $\text{member}(k)$: is k an item in S ? We use a popular data-aware method called gap encoding for space compression. In the practical case when the size of S is small (precisely, $|S| = o(n)$), we show that our data structure is a better alternative to the smallest existing one, matching its space and query time for $\text{rank}(i)$, but supports faster query for $\text{select}(j)$ and $\text{member}(k)$.

11 Bioinformatics

The objective of our efforts in bioinformatics is to address different issues where database support is needed.

11.1 Database Support in Proteomics

High throughput instruments like mass spectrometers (MS) are becoming commonplace in life sciences. Current MS techniques may require repeated experimental runs consuming time and limited biological samples while demanding larger data storage resources. In this project, we propose a novel strategy for managing MS data that acquires data only for interesting molecules. In this strategy, a MS is operated primarily under the single MS mode. The system performs a rapid initial analysis of the full mass spectrum as it is being acquired and, based on the results, can predict the direction of the experiment and determine the data needed for the next step. If a potentially interesting ion is found, the instrument is directed to switch to the tandem MS/MS mode to acquire sequence information on the potentially interesting ion. We are also investigating other issues including prediction of upcoming spectra based on statistical methods and data mining, spectra indexing, and access to external databases.

11.2 Protein Annotation

One of the problems of current biological databases is that computational function annotations are sometimes not clearly distinguished from those with experimental evidence, and consequently, erroneous annotations are propagated in the databases by computational methods. Moreover, computational function annotation is informative and indispensable in the current large and growing biological databases. To manage the quality of computational function annotations, we address the following specific goals in this project: (1) Quantifying annotation reliability and errors. We will establish quantitative values related to annotation reliability for each protein family. (2) Providing benchmark data sets for computational annotation methods, including curated multiple sequence alignments of protein families. (3) Clarifying annotation dependency. (4) Providing re-annotation of genes associated with the reliability values. The system will support the functionality to allow users to generate their own private annotations. With a quantitative measure of the error rate, the system will provide aggressive function annotation, which will be useful to some researchers. (5) Identifying hot spots for experimental function verification from annotation dependencies.

12 Web Services

The Web is evolving from a passive medium for publishing data to a more active platform for conducting business using Web services. Complex applications accessing diverse Web services (e.g., a travel package) require an integrated and efficient way to manipulate and deliver Web services. We propose a query infrastructure [21] that offers complex query facilities over Web services. Users submit declarative queries that are resolved through the combined invocations of different Web service operations. As large numbers of Web services are expected to compete in offering “similar” functionalities, we propose an optimization model based on Quality of Web Service (QoWS). We propose to monitor QoWS to measure the Web services fluctuations and rate them. We are also investigating ways to enable life scientists pose complex queries which will transparently search and dynamically compose different bioinformatics tools represented by Web services.

13 From Research to Practice

ICDS has performed several large-scale, multi-year efforts that brings research from ICDS and elsewhere to bear on real-world problems giving students training in database development techniques in a much greater depth than is possible in the classroom.

13.1 Dynamic Maintenance

Knowledge projection aims to revolutionize maintenance operations for Navy vessels by developing tech-

nologies needed to create high performance knowledge bases that represent, capture, analyze, and deliver maintenance knowledge. In this project, we aim at transforming shipboard maintenance into a dynamic, pro-active process that links real-time shipboard troubleshooting actions, diagnostic data, tacit sailor experience, shore-based expertise, and failure resolution directly into knowledge discovery. The project uses XML, web services, and database technologies for its underlying knowledge base support structure, and the knowledge projection concept guarantees a real-time end-to-end solution that extends from the shipboard maintainers to the shore-based experts.

13.2 Digital Government

The core of our research in digital government (a joint project with Virginia Tech) is the development of techniques to efficiently access government services and databases [19]. We propose a framework that enables uniform access to a large number of government services and databases called *WebDG*. Its main features can be summarized as follows. Databases are organized as *distributed ontologies* to accelerate their discovery. Each ontology contains databases that share the same domain of interest. We wrap applications by *Web services* to cater for the dynamic *discovery* and *composition* of welfare programs. To preserve citizens' privacy, we propose a model based on the citizen's privacy preference. Users are granted *privacy credentials* that define their access scope.

References

- [1] W. G. Aref and I. F. Ilyas. An extensible index for spatial databases. In *SSDBM*, 2001.
- [2] W. G. Aref and I. F. Ilyas. Sp-gist: An extensible database index for supporting space partitioning trees. *J. Intell. Inf. Syst.*, 17(2-3), 2001.
- [3] E. Bertino, J. Fan, E. Ferrari, M.-S. Hacid, A. K. Elmagarmid, and X. Zhu. A hierarchical access control model for video database systems. *ACM Trans. Inf. Syst.*, 21(2), 2003.
- [4] J. Byun, E. Bertino, and N. Li. Purpose based access control of complex data for privacy protection. In *Proc. 10th ACM Symposium on Access Control Models and Technologies*, Stockholm, Sweden, June 2005.
- [5] R. Cheng, Y. Xia, S. Prabhakar, and R. Shah. Change tolerant indexing for constantly evolving data. In *ICDE 2005*.
- [6] R. Cheng, Y. Xia, S. Prabhakar, R. Shah, and J. S. Vitter. Efficient indexing methods for probabilistic threshold queries over uncertain data. In *VLDB 2004*.
- [7] C. Clifton, A. Doan, A. Elmagarmid, M. Kantarcioglu, G. Schadow, D. Suci, and J. Vaidya. Privacy preserving data integration and sharing. In *ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, Paris, France, 2004.
- [8] J. Fan, H. Luo, and A. K. Elmagarmid. Concept-oriented indexing of video databases: Towards semantic sensitive retrieval and browsing. *IEEE Trans. on Image Processing*, 13(5), 2004.
- [9] J. Fan, X. Zhu, A. Elmagarmid, W. Aref, and L. Wu. Classview: Hierarchical video shot classification, indexing, and accessing. *IEEE Trans. on Multimedia*, 6(1), 2004.
- [10] T. M. Ghanem, R. Shah, M. F. Mokbel, W. G. Aref, and J. S. Vitter. Bulk operations for space-partitioning trees. In *ICDE 2004*.
- [11] R. Grossi, A. Gupta, and J. S. Vitter. High-order entropy-compressed text indexes. In *SODA*, 2003.
- [12] M. A. Hammad, M. J. Franklin, W. G. Aref, and A. K. Elmagarmid. Scheduling for shared window joins over data streams. In *VLDB 2003*.
- [13] M. A. Hammad, M. F. Mokbel, M. H. Ali, W. G. Aref, A. C. Catlin, A. K. Elmagarmid, M. Eltabakh, M. G. Elfeky, T. M. Ghanem, R. Gwadera, I. F. Ilyas, M. S. Marzouk, and X. Xiong. Nile: A query processing engine for data streams. In *ICDE 2004*.
- [14] I. F. Ilyas, W. G. Aref, and A. K. Elmagarmid. Supporting top-k join queries in relational databases. In *VLDB 2003*.
- [15] I. F. Ilyas, W. G. Aref, and A. K. Elmagarmid. Joining ranked inputs in practice. In *VLDB*, 2002.
- [16] I. F. Ilyas, R. Shah, W. G. Aref, J. S. Vitter, and A. K. Elmagarmid. Rank-aware query optimization. In *SIGMOD 2004*.
- [17] M. Kantarcioglu and C. Clifton. Privacy-preserving distributed mining of association rules on horizontally partitioned data. *IEEE TKDE*, 16(9), 2004.
- [18] M. Kantarcioglu, J. Jin, and C. Clifton. When do data mining results violate privacy? In *KDD*, 2004.
- [19] B. Medjahed, A. Rezgui, A. Bouguettaya, and M. Ouzzani. Infrastructure for E-Government Web Services. *IEEE Internet Computing*, 7(1), 2003.
- [20] M. F. Mokbel, X. Xiong, and W. G. Aref. Sina: Scalable incremental processing of continuous queries in spatio-temporal databases. In *SIGMOD 2004*.
- [21] M. Ouzzani and A. Bouguettaya. Efficient Access to Web Services. *IEEE Internet Computing*, 8(2), 2004.
- [22] R. Sion, M. J. Atallah, and S. Prabhakar. Rights protection for relational data. In *SIGMOD 2003*.
- [23] J. Vaidya and C. Clifton. Privacy-preserving -means clustering over vertically partitioned data. In *KDD*, 2003.
- [24] X. Xiong, M. F. Mokbel, and W. G. Aref. Sea-cnn: Scalable processing of continuous k-nearest neighbor queries in spatio-temporal databases. In *ICDE 2005*.
- [25] Y. Koglin, G. Mella, E. Bertino, and E. Ferrari. An update protocol for xml documents in distributed and cooperative systems. In *ICDCS 2005*.
- [26] X. Zhu, A. E. X. Wu, A. Feng, and L. Wu. Video data mining: Semantic indexing and event detection from the association perspective. *IEEE TKDE*, 17(5), 2005.