

# Recent Progress on Selected Topics in Database Research — A Report by Nine Young Chinese Researchers Working in the United States

Zhiyuan Chen<sup>1</sup>, Chen Li<sup>2</sup>, Jian Pei<sup>3\*</sup>, Yufei Tao<sup>4</sup>, Haixun Wang<sup>5</sup>, Wei Wang<sup>6</sup>, Jiong Yang<sup>7</sup>, Jun Yang<sup>8</sup>  
and Donghui Zhang<sup>9</sup>

<sup>1</sup>*Microsoft Research*

<sup>2</sup>*University of California, Irvine*

<sup>3</sup>*State University of New York at Buffalo*

<sup>4</sup>*Carnegie Mellon University*

<sup>5</sup>*IBM T.J. Watson Research Center*

<sup>6</sup>*University of North Carolina at Chapel Hill*

<sup>7</sup>*University of Illinois at Urbana Champaign*

<sup>8</sup>*Duke University*

<sup>9</sup>*Northeastern University*

E-mail: zhchen@microsoft.com; chenli@ics.uci.edu; jianpei@cse.buffalo.edu; taoyf@cs.cmu.edu;  
haixun@us.ibm.com; weiwang@cs.unc.edu; jioyang@cs.uiuc.edu; junyang@cs.duke.edu; donghui@cs.neu.edu

Received May 21, 2003; revised July 9, 2003.

**Abstract** The study on database technologies, or more generally, the technologies of data and information management, is an important and active research field. Recently, many exciting results have been reported. In this fast growing field, Chinese researchers play more and more active roles. Research papers from Chinese scholars, both in China and abroad, appear in prestigious academic forums.

In this paper, we, nine young Chinese researchers working in the United States, present concise surveys and report our recent progress on the selected fields that we are working on. Although the paper covers only a small number of topics and the selection of the topics is far from balanced, we hope that such an effort would attract more and more researchers, especially those in China, to enter the frontiers of database research and promote collaborations. For the obvious reason, the authors are listed alphabetically, while the sections are arranged in the order of the author list.

## 1 Introduction

The study on database technologies, or more generally, the technologies of data and information management, is an important and active research field. Recently, many exciting results have been reported. In this fast growing field, Chinese researchers play more and more active roles. Research papers from Chinese scholars, both in China and abroad, appear in prestigious academic forums, such as ACM SIGMOD, VLDB, IEEE ICDE, and EDBT conferences.

In this paper, we, nine young Chinese researchers working in the United States, present concise surveys and report our recent progress on the selected fields that we are working on. Although the paper covers only a small number of

topics and the selection of the topics is far from balanced, we hope that such an effort would attract more and more researchers, especially those in China, to enter the frontiers of database research. Moreover, this paper should be also read as a call for future collaborations. Interested readers are encouraged to contact the authors of the corresponding sections for further discussion. We believe that the extensive collaborations will benefit the Chinese database research community.

The writing of this paper was initialized by the two key officers of Journal of Computer Science and Technology, Prof. Guojie Li, the Chief-Editor, and Ms. Xiaoxian Wan, the executive manager who is acting as a coordinator of the paper.

For the obvious reason, the authors are listed alphabetically, while the sections are arranged in the

---

\* Dr. Jian Pei, the corresponding author of the paper, is a guest editor of J. Comput. Sci & Technol for the database-related fields.

order of the author list. The topics can be roughly divided into the following categories. Section 2 is on database system performance tuning. Section 3 is on data integration. Sections 4, 7 and 8 are on mining various kinds of data. Sections 5, 6, 9 and 10 are on indexing and query answering.

## 2 Automatic Database Tuning and Administration<sup>①</sup>

After decades of development, today's database systems all have numerous features, making it very difficult to choose these features toward the need of the specific applications using them. For example, building indexes and materialized views often dramatically improve the performance on a given query workload, but it is very difficult to select the necessary indexes and views because such a decision depends on how these queries are executed. On the other hand, the cost of hardware has dropped dramatically. Thus the cost for human to tune and manage the database systems often dominates the cost of ownership. To reduce such cost, it is desirable to automate database tuning and administration.

Database tuning and administration includes physical database design and tuning system parameters. Physical database design includes selecting indexes, views, vertical partitioning and horizontal partitioning, parallel database design, etc. Tuning system parameters includes selecting the serializability level, locking granularity, placement of log files, buffer pool size, RAID levels, cache sizes and placement, etc.

There has been much work in the area of physical database design. Earlier work<sup>[1]</sup> uses a stand-alone cost model to evaluate the goodness of a configuration. However, all such work has the drawback that it is difficult to keep the stand-alone cost model consistent with the cost model used by the query optimizer.

More recent work proposed to use the query optimizer for cost estimation, and to use a search algorithm to enumerate possible configurations. Such work includes [2–4] in the area of index selection, [5] in the area of view selection, and [6] in the area of parallel database design.

In the area of system parameter tuning, there is plenty of work giving practical guidelines<sup>[7]</sup>. However, there is relatively less attempt to automate this<sup>[8]</sup>.

There are many research problems unsolved in this area. First, very little work has been done in automatically tuning system parameters, and it is challenging to predict the system performance after changing such parameters. Second, little is known on how to adjust the system to changes of the workload. Ideally, the database system shall be able to automatically adjust to such changes. Third, given the numerous features to tune, it remains challenging to identify the system bottleneck as well as to tune all these together.

## 3 Integrating Information from Heterogeneous Data Sources<sup>②</sup>

The purpose of data integration (a.k.a., information integration, data mediation) is to support seamless access to autonomous, heterogeneous information sources, such as legacy databases, corporate databases connected by intranets, and sources on the Web. Many research systems (e.g., [9–11]) have been developed to achieve this goal. These systems adopt a mediation architecture<sup>[12]</sup>, in which a user poses a query to a *mediator* that retrieves data from underlying sources to answer the query. A *wrapper* on a source is used to perform data translation and local query processing.

Intensive research has been conducted on challenges that arise in data integration. The first challenge is how to support interoperability of sources, which have different data models (relational, XML, etc.), schemas, data representations, and querying interfaces. Wrapper techniques have been developed to solve these issues. The second challenge is how to model source contents and user queries, and two approaches have been widely adopted. In the local-as-view (LAV) approach, a collection of *global predicates* are used to describe source contents as views and formulate user queries. Given a user query, the mediation system decides how to answer the query by synthesizing source views, called *answering queries using views*. Many techniques have been developed to solve this problem<sup>[13,14]</sup>, and these techniques can also be used in other database applications such as data warehousing and query optimization. Another approach to data integration, called the global-as-view (GAV), assumes that user queries are posed directly on global views that are defined on source relations. In this approach, a query plan can be generated using a

<sup>①</sup> This section is contributed by Dr. Zhiyuan Chen, Microsoft Research, zhchen@microsoft.com.

<sup>②</sup> This section is contributed by Dr. Chen Li, University of California, Irvine, chenli@ics.uci.edu.

view-expansion process. Researchers mainly focus on efficient query processing in this case. The third challenge is how to process and optimize queries when sources have limited query capabilities. For instance, the Amazon.com source can be viewed as a database that provides book information. However, we cannot easily download all its books. Instead, we can query the source by filling out Web search forms and retrieving the results. Studies have been conducted on how to model and compute source capabilities, how to generate plans for queries, and how to optimize queries in the presence of limited capabilities<sup>[15,16]</sup>.

Here we give three of the many open problems in data integration that need more research investigation. First, most previous mediation systems adopt a *centralized* architecture. Recently, database applications are seeing the emerging need to support data integration and sharing in *distributed, peer-based* environments. In such an environment, autonomous peers (sources) connected by a network are willing to exchange data and services with each other. Thus each peer is both a data source and a “mediator.” Several projects have been proposed to study issues in this new architecture<sup>[17,18]</sup>. Second, researchers are paying more attention on how to deal with source heterogeneity by cleansing data. One particular problem is called *data linkage*, i.e., identifying and linking duplicate records efficiently and effectively. Sources frequently contain approximately duplicate fields and records that refer to the same real-world entity, but are not identical, such as “Tom M. Franz” versus “Franz, T.”, and “Barranca Blvd., Irvine, CA” versus “Barranca Boulevard, Irvine, California”. Variations in representations may arise from typographical errors, misspellings, abbreviations, and other reasons. This problem is especially severe when data is automatically extracted from unstructured or semi-structured documents or Web pages. In order to integrate information from sources, we need to “clean” the data before doing a higher-level processing. Recently many results are developed on data cleansing and linkage<sup>[19,20]</sup>. Third, the advent of XML introduces many new problems that need to be solved to support data integration.

#### 4 Mining Changes from Data Streams <sup>③</sup>

A growing number of emerging applications, such as sensor networks, networking flow analysis, and e-business and stock market online analysis,

have to handle various data streams. It is demanding to conduct advanced analysis and data mining over fast and large data streams to capture the trends, patterns, and exceptions. Recently, some interesting results have been reported of modelling and handling data streams (see [21] for a comprehensive overview), such as monitoring statistics over streams and queries answering (e.g., [22–24]). Furthermore, conventional OLAP and data mining models have been extended to tackle data streams, such as multi-dimensional analysis (e.g., [25]), clustering (e.g., [26]) and classification (e.g., [27, 28]).

While extending the existing data mining models to tackle data streams may provide valuable insights into the streaming data, it is high time we considered the following fundamental question: *Compared to the previous studies on mining various kinds of data, what are the distinct features/core problems of mining data streams?* In other words, *from mining data streams, do we expect something different than mining other kinds of data?*

Previous studies (e.g., [21, 29]) argue that mining data streams is challenging in the following two respects. On the one hand, random access to fast and large data streams may be impossible. Thus, multi-pass algorithms (i.e., ones that load data items into main memory multiple times) are often infeasible. On the other hand, the exact answers from data streams are often too expensive. Thus, approximate answers are acceptable. While the above two issues are critical, they are not unique to data streams. For example, online mining very large databases also requires ideally one-pass algorithms and may also accept approximations.

We argue that one of the keys to mining data streams is *online mining of changes*. For example, consider a stream of regular updates of various aircrafts’ positions. An air traffic controller may be interested in the clusters of the aircrafts at each moment. However, instead of checking details for “normal” clusters, she/he may be more interested in those “abnormal” clusters, e.g., fast growing clusters indicating the forming of a traffic jam. In general, while the patterns in snapshots of data streams are important and interesting, the *changes to the patterns* may be more critical and informative. With data streams, people are often interested in mining queries like “*compared to the history, what are the distinct features of the current status?*” and “*what are the relatively stable factors*”

<sup>③</sup> This section is contributed by Dr. Jian Pei, State University of New York at Buffalo, jianpei@cse.buffalo.edu.

over time?” Clearly, to answer the above queries, we have to examine the changes.

Some previous works also involve change detection. For example, the emerging patterns<sup>[30]</sup> characterize the changes from one data set to another. In [31], Ganti *et al.* propose methods of measuring the differences of the induced models in data sets. Incremental mining studies how to update the models/patterns by factoring in the incremental part of data. However, mining data streams requires online and dynamic detection and summarization of interesting changes.

Interesting research problems on mining changes in data streams can be divided into three categories: modelling and representation of changes, mining methods, and interactive exploration of changes.

First, while the term “changes” sounds general and intuitive, it is far from trivial to define and describe changes in data streams. First, it is essential to propose concise query language constructs for describing the mining queries on changes in data streams. There can be many kinds of changes in data streams, and different users may be interested in different kinds. The user should be able to specify the changes she/he wants to see. Moreover, the system should be able to rank changes based on interestingness. The operations should be integrable into the existing data mining models and languages. An “algebra” for change mining may be essential. Second, methods of summarizing and representing the changes need to be developed. In particular, effective visualization of changes is very important. Third, while the model for mining “first order” changes is common and useful, the model for mining “higher order” changes can be an important kind of knowledge in some dynamic environments. For example, a stock market analyst may feel particularly interested in the changes in the ranges of price vibration, while the range of price vibration itself is a description of changes.

Second, efficient and scalable algorithms are needed for mining changes in data streams, at various levels. First, specific algorithms can be developed for specific change mining queries. While such query-specific approaches may not be systematic, it will provide valuable insights into the inherent properties, challenges and basic methods of change mining. Second, general evaluation methods for “change mining queries” should be developed based on the general model/query

language/algebra. Third, facilities for various aspects of change mining, such as quality management, should be considered. For example, algorithms should be able to meet user’s specification on level/granularities/approximation error bound of change mining.

Third, the results from change mining per second form data streams, which can sometimes be large and fast. It is important to develop effective approaches to support user’s interactive exploration of the changes. For example, a user may want to monitor the changes at an appropriate level. Once some interesting changes are detected, she/he can closely inspect the related details.

To the best of our knowledge, the above problems have not been researched systematically so far. By no means is the above list complete. We believe that thorough studies on these issues will bring about many challenges, opportunities, and benefits to stream data processing, management, and analysis. Our recent studies<sup>[32,33]</sup> are concrete steps towards this direction.

## 5 Spatio-Temporal Indexing and Querying<sup>④</sup>

The *spatio-temporal database* (STDB) has received considerable attention during the past few years, due to the emergence of numerous applications (e.g., flight control systems, weather forecast, mobile computing, etc.) that demand efficient management of moving objects. These applications record objects’ geographical locations (sometimes also shapes) at various timestamps, and support queries that explore their historical and future (predictive) behaviors. The STDB significantly extends the traditional *spatial database*, which deals with only stationary data and hence is inapplicable to moving objects, whose dynamic behavior requires re-investigation of numerous topics including data modeling, indexes, and the related query algorithms. In this section, we survey the existing solutions for these issues.

Depending on the temporal aspects of data, an STDB aims at either *historical* or *predictive* retrieval. Specifically, given a set of objects  $o_1, o_2, \dots, o_N$  (where  $N$  is termed the *cardinality*), a historical STDB stores, for each object  $o_i$  ( $1 \leq i \leq N$ ), its extent  $o_i.E(t)$  at all the timestamps  $t$  in the history. Following the convention of spatial databases, each extent  $o_i.E(t)$  can be a polygon describing the

<sup>④</sup> This section is contributed by Dr. Yufei Tao, Carnegie Mellon University, taoyf@cs.cmu.edu.

object’s actual shape at time  $t$  (e.g., the contour of a moving typhoon). Specially, if the shape is not important (e.g., cars, flights, etc.),  $o_i.E(t)$  degenerates to a point describing the location of  $o_i$  at time  $t$ . In practice, the extents of the same object at the successive timestamps can be compressed using various methods (e.g., if the object remains stationary at several continuous timestamps, its extent is stored only once during this period). A predictive STDB, on the other hand, stores, for each (usual point) object  $o_i$ , its most recent updated location  $o_i.L(t_{upd})$  (where  $t_{upd}$  is the time of the object’s last update), and the motion function describing its current movement. The most popular motion function is the linear function<sup>[34–36]</sup>, because it (1) can approximate any trajectory, and (2) requires the fewest number of parameters. Specifically, in addition to  $o_i.L(t_{upd})$ , the system only needs to record the object’s velocity  $o_i.vel$ , such that the object’s location at any future time  $t > t_{upd}$  can be calculated as  $o_i.L(t) = o_i.L(t_{upd}) + o_i.vel \cdot (t - t_{upd})$ . Using such modeling, an object needs to issue an update to the database only if the parameters of its motion function (e.g.,  $o_i.vel$  for linear movement) change.

Since the spatial database can be regarded as a special type of STDB where all the objects have zero velocities, all the spatial query types naturally find their counterparts in STDB, except that they are augmented with additional temporal predicates. Specifically, a *window query* (WQ) specifies a query region  $q_R$  and time interval  $q_T$  (consisting of continuous timestamps), and finds all the objects whose extents (or locations for point data) intersect  $q_R$  during  $q_T$ . Particularly, the *selectivity* of a WQ equals the number of retrieved objects divided by the dataset cardinality, and its accurate estimation<sup>[37–39]</sup> is imperative to query optimization. A *k nearest neighbor* ( $kNN$ ) query specifies a query point  $q_P$  and time interval  $q_T$ , and finds the  $k$  objects whose distances to  $q_P$  during  $q_T$  are the shortest. These problems become even more complex if query regions/points (in WQ/ $kNN$ ) are also moving. While the above queries involve only one dataset, the *within-distance* join (WDJ), given two datasets  $S_1, S_2$ , reports all the object pairs ( $o_1, o_2$ ) in the cartesian product  $S_1 \times S_2$ , such that the distance between  $o_1, o_2$  during a query time interval  $q_T$  is shorter than that of a certain threshold  $d$ . The *selectivity* of a join is the number of retrieved pairs divided by the size of  $S_1 \times S_2$ . Similarly, the *k closest pair* ( $kCP$ ) query retrieves the  $k$  object pairs ( $o_1, o_2$ ) such that the distance of  $o_1, o_2$  dur-

ing  $q_T$  is the shortest, among all the pairs in  $S_1 \times S_2$ . Note that the above queries can be defined in both historical and predictive STDB.

In addition to queries inherited from conventional spatial databases, the dynamic nature of STDB also leads to several novel query types. For historical databases, [40] introduces the *navigational* WQ, which specifies two query regions  $q_{R1}, q_{R2}$  and timestamps  $q_{T1}, q_{T2}$ , and retrieves all the objects that intersect  $q_{R1}$  at  $q_{T1}$ , and also intersect  $q_{R2}$  at  $q_{T2}$  (e.g., “find all the vehicles that appeared in Harvard at 5pm yesterday and then appeared in MIT 10 minutes later”). In predictive STDB, [41] points out that the results of traditional queries (i.e., WQ,  $kNN$ , WDJ,  $kCP$ ) are usually inadequate because they may change (sometimes almost immediately) due to the movements of objects and/or queries (e.g., a user’s nearest gas station may change as she/he drives on the highway). Motivated by this, [41] proposes the *time-parameterized* (TP) queries, which applies to any traditional query, and returns, in addition to the result  $R$ , also (1) an expiry time  $T$  of  $R$ , and (2) the change  $C$  of the result after  $T$ . An example of TPNN is to report (1) the nearest station  $s$ , (2) when  $s$  will cease to be the nearest (given the user’s moving direction and speed), and (3) the new nearest station after the expiry of  $s$ . The concept of TP is extended to the *continuous query* in [42–45], which is another general concept applicable to all traditional queries, and aims at continuously tracking the result changes until certain conditions are satisfied. A continuous WQ, for instance, may “return the aircrafts within 10 miles from flight UA183 now, and continuously update this information until its arrival”. In TP and continuous processing, the moving direction of the query can be clearly specified, which is not true in some applications (e.g., a tourist wandering around casually). The concept useful in such scenarios is the *location-based* (LB) query<sup>[46]</sup>, which applies to WQ and  $kNN$ , and finds the query result as well as its validity region such that, as long as the query is in this region, its result will remain the same. For example, a LB NN may return the nearest restaurant of a tourist, as well as a validity region in which the restaurant will remain the nearest.

Numerous access methods have been proposed for efficient spatio-temporal query processing. A straight-forward approach to index historical STDB is to create a spatial index (the most common ones including the R-tree<sup>[47,48]</sup>) at each timestamp in history, managing objects’ extents at

that timestamp. This is the idea behind the so-called *partially persistent structures*<sup>[49,50]</sup>, which, in order to reduce the space consumption, allow the R-trees at consecutive timestamps to share common nodes, if the objects in these nodes do not incur extent changes. The first partially persistent structure, the *historical R-tree* (HR-tree)<sup>[51]</sup>, however, still involves considerable data redundancy as analyzed in [52], which led to the development of the *multi-version R-tree* (MVR-tree)<sup>[53]</sup>, and its subsequent versions<sup>[54–56]</sup>. Besides the partially persistent methodology, historical STDB can also be indexed using a 3D R-tree by treating time just as an extra dimension (in addition to the two spatial dimensions). Specifically, each record in the 3D R-tree<sup>[57]</sup> represents a 3D box, whose spatial projection corresponds to the extent of a stationary object, and whose temporal projection denotes the time interval during which the object is stationary. Similar ideas are used in the *trajectory bundle tree* (TB-tree)<sup>[40]</sup>, a structure optimized for navigational WQ queries. In practical STDB, [58] adapts the Quadtree<sup>[59]</sup> (a spatial index) for indexing the movements of 1D objects, while the *time-parameterized R-tree*<sup>[60]</sup> (TPR-tree) and its improved versions<sup>[61,62]</sup> support objects of arbitrary dimensionality. Finally, indexing moving objects has also been studied in theory<sup>[63–66]</sup>, which develop numerous interesting structures with provable worst-case performance bounds. These bounds, however, usually involve large hidden constants, rendering these “theoretical” solutions to be outperformed by the “practical” solutions introduced earlier.

## 6 Indexing XML by Tree Structures<sup>Ⓢ</sup>

With the growing importance of XML in data exchange, much research has been done in providing flexible query mechanisms to extract data from XML documents<sup>[67–72]</sup>. We propose a novel index structure that addresses a wide range of challenges in indexing semi-structured data<sup>[73]</sup>. Unlike state-of-the-art approaches that disassemble a structured query into multiple sub-queries, and then *join* the results of these sub-queries to provide the final answers, our method uses tree structures as the basic unit of query to avoid expensive join operations.

XML provides a flexible way to define semi-structured data. Fig.1 shows an XML document representing a purchase record. Many state-of-the-art approaches create indexes on paths or nodes in DTD trees. However, queries involving branching structures usually have to be disassembled into multiple sub-queries, each corresponding to a single path in the graph. The results of these sub-queries are then combined by expensive *join* operations to produce final answers. For the same reason, these methods are also inefficient in handling ‘\*’ or ‘//’ queries, which too, correspond to multiple paths. Moreover, to retrieve semi-structured data efficiently, it is essential to have index on both structure and content of the XML data. Nevertheless, many algorithms index on structure only, or index on structure and content separately, which means attribute values, are not used for filtering in the most effective way.

We transform XML data and queries into *structure-encoded sequences* by a depth-first traversal

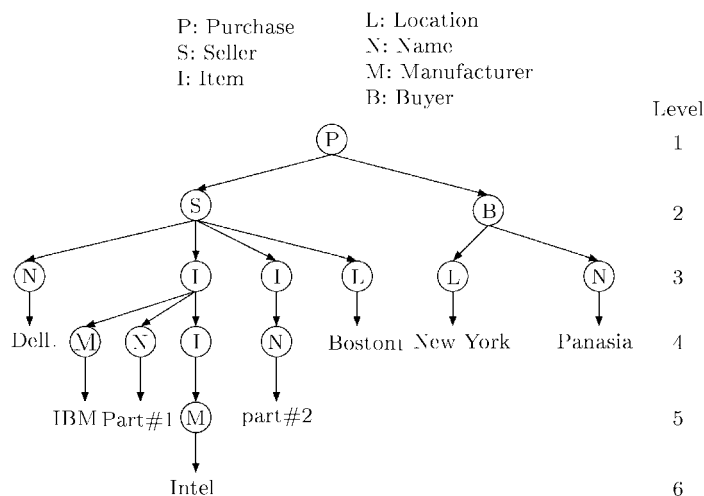


Fig.1. A single purchase record.

<sup>Ⓢ</sup>This section is contributed by Dr. Haixun Wang, IBM T.J. Watson Research Center, haixun@us.ibm.com.

sal of their tree structure, and represent each node by (Symbol, Prefix) where a symbol encodes the node and a prefix is the path from the root node. We also convert XML queries to sequences. Most structural XML queries can be performed through (non-contiguous) subsequence matching. The only exception occurs when a branch has multiple identical child nodes<sup>[73]</sup>. The benefits of modeling XML queries through sequence matching is that structural queries can be processed as a whole instead of being broken down to smaller query units (paths or nodes of XML document trees). For example, Fig.2 is a sequential representation of the purchase record, where the underlining part matches query “find purchase records of Boston sellers and NY buyers”.

The index structure is built in three steps: 1) adding all structure-encoded sequences into a suffix tree; 2) labeling each node in the suffix tree by  $\langle n, size \rangle$  through a preorder traversal where  $n$  is the preorder number of the node, and  $size$  is the number of its descendants; and 3) for each node (Symbol, Prefix) labeled  $\langle n, size \rangle$ , inserting it to the **D-Anccestor B+Tree** using (Symbol, Prefix) as the key, and then into the **S-Anccestor B+Tree** using  $n$  as the key.

$(P, \epsilon), (S, P), (N, PS), (v_1, PSN), (I, PS), (M, PSI),$   
 $(v_2, PSIM), (N, PSI), (v_3, PSIN), (I, PSI), (M, PSII),$   
 $(v_4, PSIIIM), (I, PS), (N, PSI), (v_5, PSIN), (\underline{L, PS}),$   
 $(v_6, PSL), (\underline{B, P}), (\underline{L, PB}), (v_7, PBL), (N, PB), (v_8, PBN)$

Fig.2. Sequential representation of Fig.1.

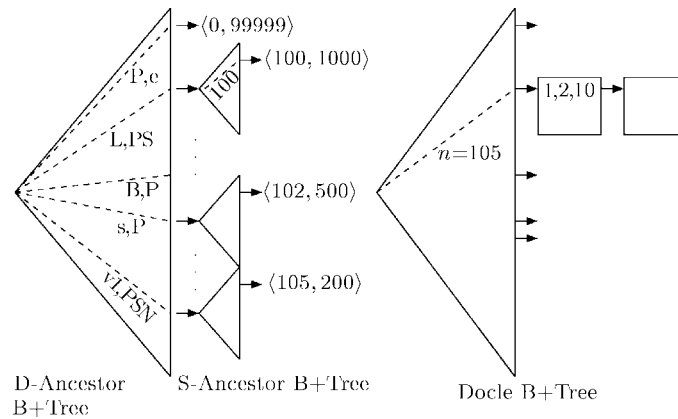


Fig.3. The index structure.

Suppose node  $x$ , labeled with  $\langle n_x, size_x \rangle$ , is one of the nodes matching a query prefix  $q_1, \dots, q_{i-1}$ . To match the next element  $q_i$  in the query, we consult the **D-Ancursor B+Tree** using  $q_i$  as a key. The **D-Ancursor B+Tree** returns the root of an **S-Ancursor B+Tree**. We then issue a range query  $n_x < n \leq n_x + size_x$  on the **S-Ancursor B+Tree** to find the descendants of  $x$  immediately. For each descendant, we use the same process to match symbol  $q_{i+1}$ , until we reach the last element of the query.

If node  $y$  is one of the nodes that matches the last element in the query, then the document IDs associated with  $y$  or any descendant node of  $y$  are answers to the query. Based on  $y$ 's label, say  $\langle n_y, size_y \rangle$ , we know  $y$ 's descendants are in the range of  $(n_y, n_y + size_y]$ . Thus, we perform a range query  $[n_y, n_y + size_y]$  on the **DocId B+Tree** to retrieve all the document IDs for  $y$  and  $y$ 's descendants.

Representing XML documents by sequences and performing XML queries by subsequence matching create many research issues. We are currently improving the index structure to support 1) dynamic index, as the labeling scheme described above prevents the index structure from supporting dynamic insertion and deletion<sup>[73]</sup>; 2) better selectivity, as different (Symbol, Prefix) pairs have different selectivity, hence their order in the query sequence has an impact on the query performance; 3) approximate queries, where approximate tree matching can be simulated by approximate subsequence matching; and 4) top-K queries.

## 7 Clustering Biological Data<sup>©</sup>

Currently, bio-informatics has become one of the research areas that receive most of attention. In general, bio-informatics aims to solve complicated

<sup>©</sup>This section is contributed by Dr. Wei Wang, University of North Carolina at Chapel Hill, weiwang@cs.unc.edu.

biological problems, e.g., *gene* regulatory network induction, *motif* discovery, etc, with computation algorithms. Many of the data is in the form of sequences, e.g., DNA sequences, protein sequences, etc. It is widely believed that the functionality of these biological sequences is highly dependent on its structures. Automatically extracting/analyzing the biological structures is an important step in better understanding their functionality.

Clustering has been widely recognized as a powerful data mining technique and has been studied extensively during recent years. The major goal of clustering is to create a partition of objects such that the objects in each group have similar features. The result can potentially reveal unknown object groups/categories that may lead to a better understanding of the nature. In the context we address, each object is a symbol sequence and all potential features of a sequence are encoded (implicitly) in the specific symbol layout in the sequence. It is interesting to notice that these structural characteristics sometimes can uniquely determine the functional properties of the sequence and often play a decisive role in meaningful clustering of the sequence.

Significant challenges exist on how to abstract and represent distinctive structural characteristics (of sequences of vastly different lengths) and to design an effective yet efficient similarity measure based on them. One possible approach to tackling this difficulty is to use the *edit distance*<sup>[74]</sup> to measure the distance between each pair of sequences. This, however, is not an ideal solution because, in addition to its inefficiency in calculation, the edit distance only captures the optimal global alignment between a pair of sequences, but ignores many other local alignments that often represent important features shared by the pair of sequences. These overlooked features may be very crucial to produce meaningful clusters.

Another approach that has been widely used in document clustering is the keyword-based method. Instead of being treated as a sequence, each text document is regarded as a set of keywords or phrases and is usually represented by a weighted word vector. The similarity between two documents is measured based on keywords and phrases shared by them and is often defined in some form of normalized dot-product. A direct extension of this method to generic symbol sequences is to use short segments of fixed length  $q$  (generated using a sliding window through each sequence) as the set of “words” in the similarity measure.

This method is also referred to as the  $q$ -gram-based method in the literature. While the  $q$ -gram-based approach enables significant segments (i.e., keywords/phrases/ $q$ -grams) to be identified and used to measure the similarity between sequences regardless of their relative positions in different sequences, valuable information may be lost as a result of ignoring sequential relationship (e.g., ordering, correlation, dependency, etc.) among these segments, which impacts the quality of clustering.

In [75], a new generative model is devised for measuring the similarity between a cluster and a sequence. Statistics properties of sequence construction are used to assess the similarity. Sequences belonging to one cluster may subsume to the same probability distribution of symbols (conditioning on the preceding segment of a certain length), while different clusters may follow different underlying probability distributions. This feature, typically referred to as *short memory*, which is common to many applications, indicates that, for a certain sequence, the empirical probability distribution of the next symbol given the preceding segment can be accurately approximated by observing no more than the last  $L$  symbols in that segment. Significant features of such probability distribution can be very powerful in distinguishing different clusters. By extracting and maintaining significant patterns characterizing (potential) sequence clusters, one can easily determine whether a sequence should belong to a cluster by calculating the likelihood of (re)producing the sequence under the probability distribution that characterizes the given cluster.

In many applications, neither the number of clusters in the data set nor the percentage of outliers is known in advance. For instance, the biologists usually do not know the number of protein families for a set of newly mapped protein sequences. To solve this problem, a novel clustering algorithm (CLUSEQ) that produces a set of possibly overlapped clusters is proposed and is able to automatically adjust the number of clusters and the boundary to separate clustered sequences from outliers, benefiting from a unique combination of successive new cluster generation and cluster consolidation. CLUSEQ is an iterative algorithm. At each iteration, it first random generates a set of new seed clusters and adds them to the set of existing clusters. Next, CLUSEQ reclusters all sequences according to the set of clusters. Finally, the thresholds are adjusted to fit the current clusters. The algorithm terminates when no improvement on the



clustering is made in the current iteration.

CLUSEQ provides a sequence level generative model (by assuming the sequences belonging to a cluster are generated by a conditional probability distribution). However, many biological data are present in more complicated forms. For instance, the natural state of a protein is a folded 3D structure. This additional complexity poses great challenges to the clustering task and mandates further investigation and development of powerful methods.

## 8 Mining Sequential Patterns <sup>⑦</sup>

With the emergence of various applications, e.g., web log traces, system log traces, bio-informatics, etc., analyzing sequential data becomes one of the urgent problems with great practical implications. Sequential patterns is an important model for analyzing the sequential data. In this section, we are surveying several techniques used for mining various sequential patterns.

A sequence may have several periodic patterns. For instance, an event may occur at 6th, 12th, 18th, ... position in a sequence. Periodic pattern can further our understanding of the behavior of the underlining system that generates the sequence. The periodic sequential patterns were introduced in [76]. In this work, a sequence is partitioned into non-overlap contiguous portions of segments. Each segment has the same length,  $l$ . A pattern of period  $l$  will occur once at each segment. Using this as a pruning condition, the search space for periodic patterns can be reduced.

However, in many situations, the pattern may not exhibit the periodicity in the entire sequence. The pattern may be only present during some portion of the sequence. For example, in a system trace log, the system behavior may change over time. Two parameters, namely *min\_rep* and *max\_dis*, are employed to qualify valid patterns and the event subsequence containing it, where this subsequence in turn can be viewed as a list of **valid segments** of perfect repetitions interleaved by disturbance. Each valid segment is required to be of at least *min\_rep* contiguous repetitions of the pattern and the length of each piece of disturbance is allowed only up to *max\_dis*. The intuition behind this is that a pattern needs to repeat itself at least a certain number of times to demonstrate its significance and periodicity. On the other hand, the distur-

bance between two valid segments has to be within some reasonable bound. Otherwise, it would be more appropriate to treat such disturbance as a signal of “change of system behavior” instead of random noise injected into some persistent behavior. The parameter *max\_dis* acts as the boundary to separate these two phenomena. This type of pattern is called *asynchronous pattern*<sup>[77]</sup>.

To efficiently mine the asynchronous patterns, a dynamic programming algorithm is devised. First, a distance-based pruning mechanism is proposed to discover all possible periods and the set of events that are likely to appear in some pattern of each possible period. In order to find the longest valid subsequence for all possible patterns, a level-wise approach is utilized. The apriori property also holds on patterns of the same period. That is, a valid segment of a pattern is also a valid segment of any pattern with fewer events specified in the pattern. For example, a valid segment for  $(d_1, d_2, *)$  will also be one for  $(d_1, *, *)$ . Then, for each likely period, all valid patterns with their longest supporting subsequences can be mined via an iterative process.

In many applications, users may be interested in not only the frequently occurred patterns, but also the surprising patterns (i.e., beyond prior expectation) as well. A large number of occurrences of an “expected” frequent pattern sometimes may not be as interesting as a few occurrences of an “expected” rare pattern. The support model is not ideal for these applications because, in the support model, the occurrence of a pattern carries the same weight (i.e., 1) towards its significance, regardless of its likelihood of occurrence. Intuitively, the assessment of significance of a pattern in a sequence should take into account the expectation of pattern occurrence (according to some prior knowledge). A new model is proposed in [78] to characterize the class of so-called *surprising* patterns (instead of frequent patterns).

The measure of *surprise* should have the following properties: (1) The surprise of a pattern occurrence is anti-monotonic with respect to the likelihood that the pattern may occur by chance (or by prior knowledge); (2) The metric should have some physical meaning, i.e., not arbitrary created. It is fortunate that the *information* metric which is widely studied and used in the communication field can fulfill both requirements. Intuitively, information is a measurement of how likely a pattern

<sup>⑦</sup>This section is contributed by Dr. Jiong Yang, University of Illinois at Urbana Champaign, jioyang@cs.uiuc.edu.

will occur or the amount of “surprise” when a pattern actually occurs. If a pattern is expected to occur frequently based on some prior knowledge or by chance, then an occurrence of that pattern carries less information. Thus, we use information to measure the surprise of an occurrence of a pattern. The *information gain* metric is introduced to represent the accumulated information of a pattern in an event sequence and is used to assess the degree of surprise of the pattern.

Many problems in sequential pattern mining still remain as open problems. When the amount of data is large, we only can afford to scan data once, e.g., the stream environment. In this type of application, it is not only important to discover the new sequential patterns, but also interesting to discover how they evolve over the time. For instance, at the beginning of a stream *abc* is a frequent pattern, but later it becomes *abd*. From this change, we may conjecture that some user/system behavior has changed.

## 9 Derived Data Maintenance<sup>⑨</sup>

The use of *derived data* to facilitate access to *base data* is a recurring technique in many areas of computer science. Used in hardware and software caches, derived data speed up accesses to the base data. Used in replicated systems, it improves reliability and performance of applications in a wide-area network. Used as index structures, it provides fast alternative access paths to the base data. Used as materialized views<sup>[79]</sup> in databases or data warehouses<sup>[80]</sup>, it improves the performance of complex queries over the base data. Used as synopses<sup>[81]</sup>, it provides fast, approximate answer to queries or statistics necessary for cost-based query optimization. Derived data vary in complexity: It can be a simple copy of the base data, in the cases of caching and replication, or it can be the result of complex structural or content transformation of the base data, in the cases of indexes and materialized views. Derived data also vary in precision. Caches and materialized views are usually precise, while synopses are approximate.

Regardless of the form of the derived data, however, a fundamental problem is efficient maintenance of the derived data. If base data change, we need to keep the derived data consistent with the new state of the base data. Derived data maintenance is becoming an increasingly impor-

tant and difficult problem. First, more and more data are being generated at high speeds (e.g., sensor networks, network monitors, Web clickstreams, high-volume transactions from multiple sources), making it extremely expensive to compute queries from scratch. A solution is to store the query results as derived data, and incrementally maintain them as new data arrive, essentially creating a one-pass algorithm. Second, computing is becoming increasingly pervasive, and there are more and more clients accessing dynamically changing data simultaneously over a wide-area network. An approach is to place replicas of data at various locations in the network based on client access patterns. Maintenance of derived data in this case must scale up not only to the size, but also to the number of replicas. Third, as computing resources keep getting cheaper, human resources become more expensive. Traditionally, systems often rely on hand-tuning of derived data to achieve optimal performance, e.g., index tuning in databases. Nowadays, excessive hand-tuning is no longer an option. This requirement calls for more self-tuning approaches to derived data management.

At the first glance, derived data maintenance has been extensively studied and well understood in specific contexts, e.g., index updates and materialized view maintenance in databases, cache coherence and replication protocols in distributed systems. Although they share the same underlying theme, these techniques have been developed largely independently by different research communities. However, an increasing number of newer and more complex applications call for a creative combination of the traditionally separate ideas. A good example is *semantic caching*<sup>[82]</sup>, which has received tremendous interests recently for its application in caching dynamic contents<sup>[83–85]</sup>. Semantic caching incorporates the idea of materialized views into a cache. Traditionally, caching is done at the object level, and hence only supports access to objects by identifiers. If the cache receives a declarative query, it is generally impossible to tell whether the cached data provide a complete answer to the query; therefore, base data must be queried. Semantic caching supports declarative queries by maintaining the cache as a pool of materialized views. From the semantic descriptions of materialized views, it is possible to determine whether the query answers are complete, so base data are only queried when necessary. This feature makes

<sup>⑨</sup>This section is contributed by Dr. Jun Yang, Duke University, junyang@cs.duke.edu.

semantic caching a perfect solution for those applications that require declarative accesses to the data.

With more thinking of “outside the box,” we can discover more techniques such as semantic caching that combine multiple flavors of derived data to provide better solutions to the problems. For example, a lot of work in data warehousing focuses on storing auxiliary data to make warehouse data self-maintainable<sup>[86–88]</sup>, without ever accessing base data. In many cases, e.g., materialized top- $k$  and join views, we must store huge amounts of auxiliary data to guarantee self-maintenance. Borrowing the caching idea, we could remove the rigid requirement of self-maintenance, and manage auxiliary data as a semantic cache whose size is tunable by the application. The warehouse data are no longer guaranteed to be self-maintainable since the cache may occasionally “miss,” i.e., the auxiliary data required for maintenance are not found in the cache. However, with a good cache management policy, we may be able to achieve a low miss rate that lowers the expected overall cost of maintenance. We have obtained some promising results for top- $k$  view maintenance<sup>[89]</sup>, and the work is under way to generalize this technique to other types of views including joins.

As another example, we can apply indexing techniques to materialized view maintenance. Traditionally, materialized views are simply stored as database tables (possibly with a B+tree primary index). However, updating such views can be very expensive, as clearly illustrated by our earlier work on maintaining materialized temporal aggregate views<sup>[89]</sup>. For a temporal aggregate view, even a slight modification of the base table (e.g., on a single row) may affect a large number of rows in the view, making it prohibitively expensive to maintain incrementally. The solution is to materialize the view, not as a regular table or B+tree, but instead as a novel index structure called the *SB-tree*, which supports efficient range updates to the aggregate view (in time logarithmic in the size of the view).

For yet another example, consider applying view maintenance techniques to data replication. Instead of propagating all updates immediately to each replica, we propagate them only when relevant views at the replica are affected. We could also throw in approximation techniques borrowed from synopses to improve performance further. Repli-

cas can be approximations of the base data within the error bounds controlled by applications. Recent work on TRAPP<sup>[91,92]</sup> represents a step toward this general direction.

## 10 Aggregation Query Processing<sup>⑨</sup>

In this section, we summarize some recent work on *selection-based aggregation*. The selection query has been one of the most widely used queries in databases. For instance, find the gas stations that are within a given spatial range. Such queries select a subset of records from a large collection of data based on a given selection condition. In contrast, the aggregation query aims not at finding the actual records themselves, but at some aggregate value of these records. For instance, find the total number of gas stations that are within a given range. Besides this COUNT query, some other interesting queries are SUM, AVERAGE, MIN, and MAX.

A straight-forward approach to solve the selection-based aggregation problem is: first find all the records that satisfy the selection condition and then perform the aggregation on-the-fly. The problem with this approach is that the query performance is at least linear to the size of the selection result. If many records satisfy the selection condition, the performance is not satisfactory. A better approach is to build some specialized index which can help compute the aggregation result without scanning through the records. Below we summarize some major results of the latter approach.

The first problem we address is called *Range-Temporal Aggregation (RTA)*: “given a set of temporal records, each having a key, a time interval and a value, compute the total value of records whose keys are in a given range and whose intervals intersect with a given time interval”. Previous work on temporal aggregation (the most efficient one being<sup>[93]</sup>) aggregates on the whole key space. To solve the RTA problem using the previous approaches, we need to maintain a separate index for each possible key range, which is prohibitively expensive. [94] proposes a new index structure called *Multi-version SB-tree (MVSB-tree)* to solve the RTA problem in logarithmic time.

Another problem is the *Hierarchical Temporal Aggregation (HTA)*. Here we are interested in computing temporal aggregates (both with and without the key-range predicate) with fixed storage space.

<sup>⑨</sup>This section is contributed by Dr. Donghui Zhang, Northeastern University, donghui@ccs.neu.edu.

Since historical data accumulates over time, while with fixed storage space we cannot keep all of them, we have to live with storing partial information. One approach is to keep just the most recent information. However, it leads to lose of the ability to answer queries for the past. In [95], we propose to maintain the temporal aggregates under multiple granularities, with more recent information being aggregated at finer granularities.

When we consider the aggregation over spatial objects, a related problem is the *Box-Sum Aggregation*, which computes the total value of spatial objects whose regions intersect with a given region. One approach is to extend the R-tree index (state-of-art spatial index used for the selection query) by maintaining summary information in internal nodes of the tree. This approach was proposed by [96, 97] and is called the *aggregation R-tree (aR-tree)*. However, the worst-case performance is still linear to the number of records. In [98] we proposed a new index structure called the BA-tree to this problem. The query performance is pretty much logarithmic to the number of records, which is a big improvement.

A variation of the previous problem is the *Functional Box-Sum Aggregation*: “given a set of objects, each having a box and a value function, and a query box  $q$ , compute the total value of all objects that intersect  $q$ , where the value contributed by an object  $r$  is the integral of the value function of  $r$  over the intersection between  $r$  and  $q$ ”. Compared with the original box-sum problem, in some cases the new problem more accurately captures the application needs. This is because an object contributes to the query result proportional to how large it intersects the query rectangle. In [98], we proposed techniques to solve the functional box-sum problem. Basically, we managed to reduce this problem to the original box-sum problem where no function was involved.

In summary, in this section we have summarized some major research results of the selection-based aggregation problem as reported in [94, 95, 98]. In all cases, the proposed specialized index structures have much better query performance than the existing non-specialized indices. An interesting future direction is how to extend the findings to the spatio-temporal aggregation problems. A simple case is to aggregate moving points. Each object is a weighted moving point whose location is a function of time. An aggregation query asks to compute the total weight of points which will move into a given area during a given time interval. A more complicated

case is to aggregate moving objects with non-zero extents.

## References

- [1] Steve Rozen, Dennis Shasha. A framework for automating physical database design. In *VLDB'1991*, 1991, pp.401–411.
- [2] Surajit Chaudhuri, Vivek R Narasayya. An efficient cost-driven index selection tool for microsoft SQL server. In *VLDB'1997*, 1997, pp.146–155.
- [3] Surajit Chaudhuri, Vivek R. Narasayya. Index merging. In *ICDE'1999*, 1999, pp.296–303.
- [4] Gary Valentin, Michael Zuliani, Daniel C Zilio, Guy M Lohman, Alan Skelley. DB2 advisor: An optimizer smart enough to recommend its own indexes. In *ICDE'2000*, 2000, pp.101–11.
- [5] Sanjay Agrawal, Surajit Chaudhuri, Vivek R. Narasayya. Automated selection of materialized views and indexes in SQL databases. In *VLDB'2000*, 2000, pp.496–505.
- [6] Jun Rao, Chun Zhang, Nimrod Megiddo, Guy M Lohman. Automating physical database design in a parallel database. In *SIGMOD'2002*, 2002, pp.558–569.
- [7] Philippe Bonnet, Dennis Elliott Shasha. Database Tuning: Principles, Experiments, and Troubleshooting Techniques. Morgan Kaufman, 2002.
- [8] Sanjay Agrawal, Surajit Chaudhuri, Abhinandan Das, Vivek Narasayya. Automating layout of relational databases. In *ICDE'2003*, 2003.
- [9] Carey M J *et al.* Towards heterogeneous multimedia information systems: The Garlic approach. In *Proc. RIDE-DOM'95*, 1995, pp.124–131.
- [10] Levy A, Rajaraman A, Ordille J J. Querying heterogeneous information sources using source descriptions. In *Proc. VLDB*, 1996, pp.251–262.
- [11] Chawathe S S, *et al.* The TSIMMIS project: Integration of heterogeneous information sources. *IPSSJ*, 1994, pp.7–18.
- [12] Wiederhold G. Mediators in the architecture of future information systems. *IEEE Computer*, 1992, 25(3): 38–49.
- [13] Afrati F, Li C, Ullman J D. Generating efficient plans using views. In *SIGMOD*, 2001, pp.319–330.
- [14] Halevy A. Answering queries using views: A survey. *Very Large Database Journal*, 2001.
- [15] Rajaraman A, Sagiv Y, Ullman J D. Answering queries using templates with binding patterns. In *PODS*, 1995, pp.105–112.
- [16] Yerneni R, Li C, Garcia-Molina H, Ullman J D. Computing capabilities of mediators. In *SIGMOD*, 1999, pp.443–454.
- [17] The Piazza Project. University of Washington.
- [18] The Raccoon Project on Distributed Data Integration and Sharing. University of California, Irvine.
- [19] Raman V, Hellerstein J M. Potter's wheel: An interactive data cleaning system. *The VLDB Journal*, 2001, pp.381–390.
- [20] The Flamingo Project on Data Cleansing, University of California. Irvine.
- [21] Babcock B, Babu S, Datar S, Motwani R, Widom J. Models and issues in data stream systems. In *Proc. 21st ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'02)*, Madison, WI, June 2002.

- [22] Datar M, Gionis A, Indyk P, Motwani R. Maintaining stream statistics over sliding windows (extended abstract). [citeseer.nj.nec.com/491746.html](http://citeseer.nj.nec.com/491746.html).
- [23] Dobra A, Garofalakis M, Gehrke J, Rastogi R. Processing complex aggregate queries over data streams. In *Proc. 2002 ACM-SIGMOD Int. Conf. Management of Data (SIGMOD'02)*, Madison, Wisconsin, June 2002.
- [24] Gehrke J, Korn F, Srivastava D. On computing correlated aggregates over continuous data streams. In *Proc. 2001 ACM-SIGMOD Int. Conf. Management of Data (SIGMOD'01)*, Santa Barbara, CA, May 2001, pp.13–24.
- [25] Chen Y, Dong G, Han J, Wah B W, Wang J. Multi-dimensional regression analysis of time-series data streams. In *Proc. 2002 Int. Conf. Very Large Data Bases (VLDB'02)*, Hong Kong, China, Aug. 2002.
- [26] Guha S, Mishra N, Motwani R, O'Callaghan L. Clustering data streams. In *Proc. IEEE Symposium on Foundations of Computer Science (FOCS'00)*, Redondo Beach, CA, 2000, pp.359–366.
- [27] Domingos P, Hulten G. Mining high-speed data streams. In *Proc. 2000 ACM SIGKDD Int. Conf. Knowledge Discovery in Databases (KDD'00)*, Boston, MA, Aug. 2000, pp.71–80.
- [28] Hulten G, Spencer L, Domingos P. Mining time-changing data streams. In *Proc. 2001 ACM SIGKDD Int. Conf. Knowledge Discovery in Databases (KDD'01)*, San Francisco, CA, Aug. 2001.
- [29] Garofalakis M, Gehrke J, Rastogi R. Querying and mining data streams: You only get one look. In *Proc. 2002 Int. Conf. Very Large Data Bases (VLDB'02)*, Hong Kong, China, Aug. 2002.
- [30] Dong G, Li J. Efficient mining of emerging patterns: Discovering trends and differences. In *Proc. 1999 Int. Conf. Knowledge Discovery and Data Mining (KDD'99)*, San Diego, CA, Aug. 1999, pp.43–52.
- [31] Ganti V, Gehrke J, Ramakrishnan R. A framework for measuring changes in data characteristics. In *Proceedings of the Eighteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, May 31 – June 2, 1999, Philadelphia, Pennsylvania, ACM Press, 1999, pp.126–137.
- [32] Dong G, Han J, Lakshmanan L V S et al. Online mining of changes from data streams: Research problems and preliminary results. In *Proc. the 2003 ACM SIGMOD Workshop on Management and Processing of Data Streams*, San Diego, CA, June 2003.
- [33] Pei J, Ariwala S R, Jiang D. Online mining changes of clusters in data streams. *Submitted for publication*.
- [34] Forlizzi L, Guting R, Nardelli R, Schneider M. A data model and data structures for moving objects databases. In *SIGMOD*, 2000.
- [35] Gutting R, Bohlen M, Erwig M et al. A foundation for representing and querying moving objects. *TODS*, 2000, 25(1): 1–42.
- [36] Sistla A, Wolfson O, Chamberlain S, Dao S. Modeling and querying moving objects. *ICDE*, 1997.
- [37] Choi Y, Chung C. Selectivity estimation for spatio-temporal queries to moving objects. In *SIGMOD*, 2002.
- [38] Hadjieleftheriou M, Kollios G, Tsotras V. Performance evaluation of spatio-temporal selectivity techniques. *SS-DBM*, 2003.
- [39] Tao Y, Sun J, Papadias D. Selectivity estimation for predictive spatio-temporal queries. *ICDE*, 2003.
- [40] Pfoser D, Jensen C, Theodoridis Y. Novel approaches to the indexing of moving object trajectories. *VLDB*, 2000.
- [41] Tao Y, Papadias D. Time-parameterized queries in spatio-temporal databases. *SIGMOD*, 2002.
- [42] Benetis R, Jensen C, Karciauskas G, Saltenis S. Nearest neighbor and reverse nearest neighbor queries for moving objects. In *IDEAS*, 2002.
- [43] Song Z, Roussopoulos N. K-nearest neighbor search for moving query point. *SSTD*, 2001.
- [44] Tao Y, Papadias D. Spatial queries in dynamic environments. *To appear in TODS*, 2003.
- [45] Tao Y, Papadias D, Shen Q. Continuous nearest neighbor search. *VLDB*, 2002.
- [46] Zhang J, Zhu M, Papadias D, Tao Y, Lee D. Location-based spatial queries. *SIGMOD*, 2003.
- [47] Beckmann N, Kriegel H, Schneider R, Seeger B. The  $r^*$ -tree: An efficient and robust access method for points and rectangles. In *SIGMOD*, 1990.
- [48] Guttman A. R-trees: A dynamic index structure for spatial searching. In *SIGMOD*, 1984.
- [49] Becker B, Gschwind S, Ghler T, Seeger B, Widmayer P. An asymptotically optimal multiversion B-trees. *VLDB Journal*, 1996, 5(4): 264–275.
- [50] Salzberg B, Tsotras V. A comparison of access methods for temporal data. *ACM Computing Survey*, 1999, 31(2): 158–221.
- [51] Nascimento M, Silva J. Towards historical r-trees. *ACM Symposium on Applied Computing*, 1998.
- [52] Tao Y, Papadias D, Zhang J. Cost models for overlapping and multi-version structures. *TODS*, 2002, 27(3): 299–342.
- [53] Kumar A, Tsotras V, Faloutsos C. Designing access methods for bitemporal databases. *TKDE*, 1998, 10(1): 1–20.
- [54] Hadjieleftheriou M, Kollios G, Tsotras V, Gunopulos D. Efficient indexing of spatiotemporal objects. In *EDBT*, 2002.
- [55] Kollios G, Gunopulos D, Tsotras V, Delis A, Hadjieleftheriou M. Indexing animated objects using spatiotemporal access methods. *TKDE*, 2001.
- [56] Tao Y, Papadias D. The mv3r-tree: A spatio-temporal access method for timestamp and interval queries. In *VLDB*, 2001.
- [57] Vazirgiannis M, Theodoridis Y, Sellis T. Spatio-temporal composition and indexing for large multimedia applications. *Multimedia Systems*, 1998, 6(4): 284–298.
- [58] Tayeb J, Ulusoy O, Wolfson O. A quadtree-based dynamic attribute indexing method. *The Computer Journal*, 1998, 41(3): 185–200.
- [59] Samet H. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley Publishing Company, 1990.
- [60] Saltenis S, Jensen C, Leutenegger S, Lopez M. Indexing the positions of continuously moving objects. In *SIGMOD*, 2000.
- [61] Saltenis S, Jensen C. Indexing of moving objects for location-based services. *ICDE*, 2002.
- [62] Tao Y, Papadias D, Sun J. The tpr\*-tree: An optimized spatio-temporal access method for predictive queries. *VLDB*, 2003.
- [63] Agarwal P, Arge L, Erickson J. Indexing moving points. *PODS*, 2000.
- [64] Kollios G, Gunopulos D, Tsotras V. On indexing mobile objects. *PODS*, 1999.
- [65] Procopiuc C, Agarwal P, Har-Peled S. Star-tree: An efficient self-adjusting index for moving points. *ALENEX*, 2000.
- [66] Tao Y, Mamoulis N, Papadias D. Validity information retrieval for spatio-temporal queries. *SSTD*, 2003.

- [67] Chung C, Min J, Shim K. APEX: An adaptive path index for XML data. In *ACM SIGMOD*, June 2002.
- [68] Cooper B F, Sample N, Franklin M, Hjaltason G, Shadmon M. A fast index for semistructured data. In *VLDB*, September 2001, pp.341–350.
- [69] Goldman R, Widom J. DataGuides: Enable query formulation and optimization in semistructured databases. In *VLDB*, August 1997, pp.436–445.
- [70] Kaushik R, Bohannon P, Naughton J, Korth H. Covering indexes for branching path queries. In *ACM SIGMOD*, June 2002.
- [71] Li Q, Moon B. Indexing and querying XML data for regular path expressions. In *VLDB*, September 2001, pp.361–370.
- [72] Milo T, Suciu D. Index structures for path expression. In *Proc. 7th International Conference on Database Theory (ICDT)*, January 1999, pp.277–295.
- [73] Haixun Wang, Sanghyun Park, Wei Fan, Philip S Yu. ViST: A dynamic index method for querying XML data by tree structures. In *SIGMOD*, 2003.
- [74] Gusfield D. Algorithms on Strings, Trees, and Sequences. Cambridge University Press, 1997.
- [75] Yang J, Wang W. CLUSEQ: Efficient and efficient sequence clustering. In *Proc. 19th IEEE Int. Conf. Data Engineering (ICDE)*, 2003, pp.101–112.
- [76] Han J, Dong G, Yin Y. Efficient mining partial periodic patterns in time series database. In *Proc. Int. Conf. Data Engineering*, 1999, pp.106–115.
- [77] Yang J, Wang W, Yu P. Mining asynchronous periodic patterns in time series data. In *Proc. of the 6th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2000, pp.275–279.
- [78] Yang J, Wang W, Yu P. Info-miner: Mining surprising periodic patterns. In *Proc. the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2001, pp.395–400.
- [79] Ashish Gupta, Inderpal Singh Mumick (eds.). Materialized Views: Techniques, Implementations and Applications. MIT Press, June 1999.
- [80] Chaudhuri S, Dayal U. An overview of data warehousing and OLAP technology. *ACM SIGMOD Record*, 1997, 26(1): 65–74.
- [81] Gibbons P B, Matias Y. Synopsis data structures for massive data sets. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science: Special Issue on External Memory Algorithms and Visualization*, 1999, A: 39–70.
- [82] Dar S, Franklin M J, Jónsson B, Srivastava D, Tan M. Semantic data caching and replacement. In *Proc. the 1996 Int. Conf. Very Large Data Bases*, Bombay, India, September 1996, pp.330–341.
- [83] Candan K S, Li W S, Luo Q, Hsiung W-P, Agrawal D. Enabling dynamic content caching for database-driven web sites. In *Proc. the 2001 ACM SIGMOD Intl. Conf. Management of Data*, Santa Barbara, California, USA, May 2001.
- [84] Amiri K, Park S, Tewari R, Padmanabhan S. DBProxy: A dynamic data cache for web applications. In *Proc. the 2003 Int. Conf. Data Engineering*, Bangalore, India, March 2003, pp.821–831.
- [85] Luo Q, Naughton J F, Krishnamurthy R, Cao P, Li Y. Active query caching for database web servers. In *Proc. the 2000 Int. Workshop on the Web and Databases*, Dallas, Texas, USA, May 2000, pp.92–104.
- [86] Akinde M O, Jensen O G, Böhlen M H. Minimizing detail data in data warehouses. In *Proc. the 1998 Int. Conf. Extending Database Technology*, 1998, pp.293–307.
- [87] Quass D, Gupta G, Mumick I S, Widom J. Making views self-maintainable for data warehousing. In *Proc. the 1996 Int. Conf. Parallel and Distributed Information Systems*, December 1996, pp.158–169.
- [88] Yang J, Widom J. Temporal view self-maintenance in a warehousing environment. In *Proc. the 2000 Int. Conf. Extending Database Technology*, Konstanz, Germany, March 2000, pp.395–412.
- [89] Yi K, Yu H, Yang J, Xia G, Chen Y. Efficient maintenance of top-k views. In *Proc. the 2003 Int. Conf. Data Engineering*, Bangalore, India, March 2003, pp.189–200.
- [90] Yang J, Widom J. Incremental computation and maintenance of temporal aggregates. In *Proc. the 2001 Int. Conf. Data Engineering*, Heidelberg, Germany, April 2001.
- [91] Olston C, Jiang J, Widom J. Adaptive filters for continuous queries over distributed data streams. In *Proc. the 2003 ACM SIGMOD Int. Conf. Management of Data*, San Diego, California, USA, June 2003.
- [92] Olston C, Widom J. Best-effort cache synchronization with source cooperation. In *Proc. of the 2002 ACM SIGMOD Intl. Conf. on Management of Data*, Madison, Wisconsin, USA, June 2002.
- [93] Yang J, Widom J. Incremental computation and maintenance of temporal aggregates. In *Proc. Int. Conf. Data Engineering (ICDE)*, 2001.
- [94] Zhang D, Gunopulos D, Tsotras V J, Seeger B. Temporal aggregation over data streams using multiple granularities. In *Proc. Int. Conf. Extending Database Technology (EDBT)*, 2002.
- [95] Zhang D, Markowetz A, Tsotras V J, Gunopulos D, Seeger B. Efficient computation of temporal aggregates with range predicates. In *ACM Int. Symp. Principles of Database Systems (PODS)*, 2001.
- [96] Lazaridis I, Mehrotra S. Progressive approximate aggregate queries with a multi-resolution tree structure. In *Proceedings of ACM/SIGMOD Annual Conference on Management of Data (SIGMOD)*, 2001.
- [97] Papadias D, Kalnis P, Zhang J, Tao Y. Efficient olap operations in spatial data warehouses. In *Proc. Symp. Spatial and Temporal Databases (SSTD)*, 2001.
- [98] Zhang D, Tsotras V J, Gunopulos D. Efficient aggregation over objects with extent. In *ACM Int. Symp. Principles of Database Systems (PODS)*, 2002.

**Zhiyuan Chen** is currently a post-doctoral researcher at Microsoft Research, Redmond, WA. He received his B.S. (1995) and M.S. (1997) degrees from Fudan University, China. From 1997 to 2002, he was a graduate student at Cornell University under the guidance of Prof. Johannes Gehrke and Praveen Seshadri. He received his Ph.D. degree in August, 2002. His research interests include automatic database tuning and administration, database compression, selectivity estimation, and XML indexing.

**Chen Li** is an assistant professor in the School of Information and Computer Science at the University of California, Irvine. He received his Ph.D. degree in computer science from Stanford University in 2001, and his B.S. degree in computer science from Tsinghua University, China. His research interests are in the fields of

database and information systems, including data integration, data warehouses, data cleansing, multimedia databases, and XML.

**Jian Pei** received the B. Eng. and the M. Eng. degrees, both in computer science, from Shanghai Jiaotong University, China, in 1991 and 1993, respectively, and the Ph.D. degree in computing science from Simon Fraser University, Canada, in 2002. He was a Ph.D. candidate in Peking University in 1997–1999. He is currently an assistant professor of computer science and engineering, the State University of New York at Buffalo, USA. His research interests include data mining, data warehousing, online analytical processing, database systems, and bio-informatics. He has published over 40 research papers in refereed journals, conferences, and workshops. He is a member of the editorial board of the ACM SIGMOD Digital Symposium Collection (DiSC). He has served in the program committees of international conferences and workshops, including KDD'03, ICDE'02 and '04, ICDM'03, CIKM'02, and SIAM Data Mining'04. He is a reviewer for some leading academic journals, including ACM Transaction on Database Systems, IEEE Transaction on Knowledge and Data Engineering, Data Mining and Knowledge Discovery, and Knowledge and Information Systems. Dr. Pei is a member of the ACM, the ACM SIGMOD, the ACM SIGKDD and the IEEE Computer Society, and a professional member of the ASEE. He is the Chair of the East Coast Regional Chapter of ACM SIGKDD.

**Yufei Tao** became a research associate in Department of Computer Science, Hong Kong University of Science and Technology, after obtaining his Ph.D. degree from the same department in July 2002. Currently he is a visiting scientist at the CS Department of Carnegie Mellon University, USA. His research mainly focuses on the development of efficient query algorithms in spatio-temporal databases, as well as the application of related techniques to other areas such as temporal, spatial databases, and data warehouses. He has been awarded the Hong Kong Young Scientist Award 2002 (in physical and mathematical science) by the Hong Kong Institution of Science (HKIS) for his work on spatio-temporal data.

**Haixun Wang** received his Ph.D. degree in computer science from UCLA in 2000. He also holds the B.S. degree in computer science and the M.S. degree in computer science from Shanghai Jiao Tong University. In 2000, he joined IBM T. J. Watson Research

Center as a research staff member. His research interests include data mining, machine learning, database language and systems, bioinformatics, and XML.

**Wei Wang** received the M.S. degree from the State University of New York at Binghamton in 1995 and the Ph.D. degree in computer science from the University of California at Los Angeles in 1999. She is currently an assistant professor in the Department of Computer Science at the University of North Carolina at Chapel Hill. Her research interests include data mining, database systems and bioinformatics. She is the member of the editorial board of the Journal of Data Management and has published more than 50 research papers in international journals and referred conference proceedings.

**Jiong Yang** is currently a visiting assistant professor at University of Illinois at Urbana-Champaign. His current research interests include data mining, data integration, bio-informatics, mobile computing, and sensor networks. Dr. Yang received the B.S. degree from the University of California at Berkeley at 1994, the M.S. and Ph.D. degrees in computer science from the University of California, Los Angeles in 1996 and 1999, respectively. He is the author of more than forty research papers.

**Jun Yang** is an assistant professor of computer science at Duke University. Jun received his Ph.D. degree in computer science from Stanford University in the area of data warehousing. He is broadly interested in database and information management, with special emphasis on the management of derived data. He received his B.A. in computer science from University of California at Berkeley.

**Donghui Zhang's** primary research area is database systems. He has been doing researches on temporal, spatial and spatio-temporal database indexing, aggregation queries and join processing; efficiently storing and querying XML documents which evolve over time. His current research interests include data streams, indexing and querying moving objects, image database processing, biological data processing. Prof. Zhang earned his Ph.D. degree from University of California, Riverside in August 2002. His dissertation title is "Aggregation Computation over Complex Objects", in which he addressed the problem of computing aggregates over a large set of temporal and spatial objects. He proposed algorithms to compute such queries in logarithmic time, while previous approaches have linear performance.