

# Very Large Data Bases



**TORONTO**  **2004**

**30<sup>th</sup> INTERNATIONAL CONFERENCE**

**August 29<sup>th</sup> - September 3<sup>rd</sup>, 2004**

**Royal York Hotel, Toronto, Canada**

**[www.vldb04.org](http://www.vldb04.org)**



# 30<sup>th</sup> INTERNATIONAL CONFERENCE ON **VERY LARGE DATABASES**

August 29th – September 3rd, 2004  
Royal York Hotel, Toronto, Canada  
[www.vldb04.org](http://www.vldb04.org)

## TABLE OF CONTENTS

Table of Contents .....	1
Welcome Message From the General Chair .....	3
Welcome Message From The Technical Program Chairs.....	4
Toronto for Database Researchers.....	5
Social Events.....	11
Conference Venue.....	13
Conference Maps .....	15
Program At A Glance.....	17
Conference Program.....	20
Special Events.....	45
Tutorials .....	49
Panels.....	57
Demonstrations.....	61
Abstracts .....	63
Conference Co-Located Workshops.....	103
Conference Officers.....	107
My Notes .....	111

**TORONTO**  **2004**

---

Copyright © 2004

Brochure and cover design: Yannis Velegarakis (University of Toronto), Co-Editor: Ariel Fuxman (University of Toronto), Images: Periklis Andritsos, Apostolos Dimitromanolakis, Tasos Kementsietsidis, Yannis Velegarakis (University of Toronto), Conference Logo: IBM Toronto Lab Media Design Studio

## WELCOME MESSAGE FROM THE GENERAL CHAIR

The international conference series on Very Large Data Bases (VLDB) was launched in 1975 in Framingham MA, about 20 miles from Boston. The conference was a huge success for its time, attracting almost 100 papers and more than 150 participants.

We have come a long way since then! VLDB conferences attract today hundreds of submissions and participants. Thanks to the efforts of programme committees, authors, and the VLDB Endowment over the years, VLDB conferences constitute today a prestigious scientific forum for the presentation and exchange of research results and practical experiences among members of the international Databases community.

VLDB conference regulars are familiar with Canada. This is the third time the conference is visiting, after Montreal (1980) and Vancouver (1992.) In all three occasions, the Canadian Databases community served as backbone for the organizing and the programme committees. At the same time, as with other years, the committees that put together this conference were international, with participation from all regions of the globe. We are grateful to all members of these committees for their time and efforts. Special thanks to Tamer Özsu, the general programme chair of this year's conference; also Kyu-Young Whang, who served as the VLDB Endowment liaison. Their spirit of cooperation throughout was invaluable.

Of course, the technical programme is not the only attraction of this year's conference. The conference hotel is located in the core downtown area of Toronto, within walking distance of museums, parks, shopping areas and tourist attractions such as the CN tower, the Harbourfront or Lake Ontario. Toronto has been called the most multicultural city in the world and while here, you will have the opportunity to visit neighborhoods that have the distinctive ethnic flavour of parts of the Far East and Europe. More than that, Toronto is a well-run, safe city that visitors enjoy visiting and revisiting.

Welcome to VLDB'04 and Toronto. We hope that you enjoy both the technical programme and the city!!

John Mylopoulos  
General Chair, VLDB'04

## WELCOME MESSAGE FROM THE TECHNICAL PROGRAM CHAIRS

Welcome to the 30<sup>th</sup> International Conference on Very Large Databases (VLDB'04). VLDB Conferences are among the premier database meetings for dissemination of research results and for the exchange of latest ideas in the development and practice of database technology. The program includes two keynote talks, a 10-year award presentation, 81 research papers, 26 industrial papers (9 of which are invited), 5 tutorials, 2 panels and 34 demonstrations. It is a very rich program indeed.

This year we witnessed a significant jump of submissions. There were 504 research and industrial paper submissions, accounting for about 10% increase over last year and about 7% increase over 2002. Consequently, the competition was fierce with an acceptance rate of 16.1% for research papers and about 40% for industrial papers.

The technical program is the result of efforts by a large group of people. Three Program Committees were formed along themes (core database, infrastructure for information systems, and industrial and applications) consisting of 137 colleagues, each of whom reviewed about 13 papers. Raymond Ng and Matthias Jarke handled the tutorials, Jarek Gryz and Fred Lochovsky selected the panels, Bettina Kemme and David Toman assembled the demonstrations program. We thank them all for helping us put together an exciting program.

M. Tamer Özsu  
Donald Kossmann  
Renée J. Miller  
José A. Blakeley  
Berni Schiefer

## Toronto for Database Researchers

Periklis Andritsos

University of Toronto  
periklis@cs.toronto.edu

### Abstract

In the past, several guides for VLDB conferences have appeared in the literature. Having such a city guide is a vital and important thing intended to make the life of the attendees easier during the conference days.

In this paper, we present interesting issues related to the city of Toronto. The guide is intended for database researchers and presents an overview of the attractions of the city of Toronto, eating places and getting around guidelines.

Preliminary experimental results obtained by visiting some of the places listed and by obtaining recommendations from the community, suggest that the material presented herein is a promising direction for a memorable stay in the city of Toronto.

### 1. A Toronto Primer

Welcome to Toronto, the city whose name means “**the meeting place**” in one of the native dialects. In 1996 *Fortune* chose Toronto as the number one city outside the US for work and raising a family. Since then, the city has developed with restaurants and cafes offering all kinds of ethnic food, reflective of the multicultural background of the citizens. Statistics show that the city of Toronto represents more than 80 different ethnic groups speaking more than 100 languages [1].

Toronto is a safe and vibrant city. You can walk safely to many locations of the downtown area and the places mentioned in this guide.

---

*Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment*

Transportation is very well organized and if you come by car it is more practical to park it and enjoy a walk or a ride with the subway, buses or streetcars. At the end of this paper, there is a map of downtown Toronto.

Finally, the climate of Toronto for the end of summer is rather mild. Toronto is on the same latitude as Cannes on the sunny Riviera and Milan. Lake Ontario serves to moderate Toronto's weather to the point that its climate is one of the mildest in Canada. Generally speaking, end of summer temperatures range from 15°C (60°F) to 25°C (80°F).

### 2. Useful Tips

- **TTC** [2]  
(DW) The Toronto Transit Commission (TTC) operates a network of subways, buses and streetcars that provide convenient transportation throughout Toronto. There are two major subway lines. The **Yonge/University/Spadina** is a U-shaped loop that runs under Yonge Street and University Avenue in the downtown area. There are subway stations where major cross streets intersect Yonge St. and University Ave. The subway stations closest to downtown hotels are in Union Station, at King and University, at King and Yonge, at Queen and University and at Queen and Yonge. The **Bloor/Danforth** runs east/west underneath Bloor Street. There are free interchanges between these two subway lines at the Yonge/Bloor station and at the St. George station and the Spadina station. Street car or bus lines run east and west on major streets.

There are two TTC streetcar lines that start in Union Station. The Harborfront 509 line runs through the Harbor Front area along Queens Quay to the CNE grounds. The Spadina 510 line runs along Queens Quay and then north along Spadina Ave to Bloor St. The Bay bus that runs north and south on Bay St is a convenient way to reach many of the places described in this guide.

The stops for surface (bus and streetcar) routes are marked with a vertical red and white TTC sign. Current TTC fares are \$2.25 cash (bus and streetcar drivers do NOT carry change) or tokens/tickets at 5 for \$9.00. There is also a day pass (\$7.50) good for unlimited rides. Tickets and passes may be purchased at subway stations or from small shops that

display the TTC Tickets sign. For TTC information call 416-393-4636.

**FROM THE AIRPORT TO ROYAL YORK HOTEL:** Take the 192 Airport Rocket bus at Ter-minals 1, 2 or 3 to Kipling Station. At Kipling, take the Bloor-Danforth subway, change at Yonge/Bloor Station and take the Yonge-University-Spadina Line (southbound trains). Exit at Union Station. The 192 Airport Rocket bus operates every 25 minutes.

- **Federal and Provincial Taxes**

The **Goods and Services Tax (GST)** is a 7% tax that is charged on most goods and services sold or provided in Canada. And as Toronto is part of Ontario, purchases made in Toronto are also subject to the 8% **Provincial Sales Tax (PST)**. In brief, add a 15% tax on top of the price tags you see.

- **Tax refund program for visitors [3]**

Foreign visitors to Canada can apply for a rebate on the GST that is paid on accommodation (up to 30 nights per visit), and goods purchased in Canada and exported within 60 days of the purchase. Keep your receipts and read carefully the instructions given in the above website. You will need to download the corresponding form as well from <http://www.ccradrc.gc.ca/E/pbg/gf/gst176/>.

- **Toronto City Guides [4]**

A list of the most comprehensive and useful online guides are:

- <http://www.toronto.com>
- <http://www.city.toronto.on.ca>
- <http://www.math.toronto.edu/toronto>
- <http://www.torontotourism.com>

- **Emergency [5]**

(DW) The telephone number for fire department, police and ambulance service is 911. *TeleHealth Ontario* operates a 24-hour free medical advice service staffed by registered nurses at 1-866-797-0000. There are several major hospitals with Emergency Rooms on University Ave. between Dundas and College St.

- **Yellow Pages**

The Yellow pages contain a comprehensive list of services. They can be accessed online at <http://www.yellowpages.ca/>. The hard copies (that can hopefully be found in the Toronto hotels) also contain useful maps.

### 3. Major Attractions

#### 3.1 Near the Conference Hotel

- **The CN Tower [6]:** (DW) The ride from the airport downtown offers a view of the world's tallest building and freestanding structure (553+ m, 1815+ ft). On a clear day the public observation tower offers views over Toronto and as far away as Niagara Falls. The revolving restaurant at the top of the tower offers a panoramic view of the city. [entry through the CN Tower walkway at the west end of Union Station.]
- **The Sky Dome [7]:** (DW) Toronto's largest stadium with a retractable roof. Home to the Toronto Blue Jays. [entry from Front St. West.]
- **The Air Canada Center [8]:** Named after Canada's main airline, ACC is the main athletic facility for indoors sports, home of the NBA Toronto Raptors team and the NHL Toronto Maple Leafs team. During the summer it is used as a concert hall and there are stores that sell a variety of sports memorabilia. [entry from York Str and Lakeshore Blvd.]
- **Toronto Islands [9]:** (DW) A very large park on an island offshore. Includes restaurants, a children's amusement area and some long pleasant walking areas. Bicycles, roller blades, canoes and paddle boats can be rented in season. [Access via the Toronto Island Ferry at Bay Str and Queens Quay.]
- **Nathan Philips Square [10]:** A large public square featuring seasonal entertainment and nice ambience. The large building in the middle of the square is Toronto's City Hall. From 10am to 2pm every Wednesday the square hosts a *Farmer's Market*, where products from all over Ontario are being sold. [located on Queen Str West between Bay and University Ave.]
- **Harbour Front [11]:** (DW) Toronto's waterfront park. Very pleasant waterfront atmosphere. [Queens Quay, between Bay St. and Spadina Ave.]
- **Entertainment District [12]:** (DW) The main theater district in Toronto. Also home to **many** clubs and restaurants. [between King St and Queen Street, west of University Avenue.]
- **St. Lawrence Market [13]:** (DW) Toronto's traditional old style market. Butchers, bakers and fruit and vegetable vendors. Several interesting foodstuffs shops. On Saturdays the market doubles in size when local farmers come in to sell their produce [Front Street at Jarvis St.



- **The Distillery** [14]: The Distillery Historic District is a Canadian heritage site and Toronto's newest arts and entertainment cultural community. Throughout the year, it hosts celebrations and special events such as the Distillery Jazz Festival [located near the corner of Parliament and Front Str East.]

### 3.2 Downtown

- **Royal Ontario Museum** [15]: (DW) The ROM is Ontario's largest traditional museum. [On University Avenue just south of Bloor Street. TTC Museum Station on the University subway line.]
- **Art Gallery of Ontario** [16]: (DW) Home to a large collection of Canadian and World Art. [Dundas Street between University Avenue and Spadina.]
- **The Eaton's Center** [17]: One of Toronto's largest indoor shopping mall of all varieties and tastes [on Yonge Street between Dundas and Queen.]
- **University of Toronto** [18]: (DW) Canada's largest University. The St. George campus is north of College St, (mostly) west of University Avenue. [TTC Queens Park Station on the University subway.]
- **Gay Town** [19]: (DW) Toronto has one of the largest gay/lesbian communities in North America. Gay Town on Church Street between Carleton St. and Bloor St. is one focal point for this community. [TTC College St or Wellesley St stations on the Yonge subway.]

### 3.3 Elsewhere

- **Ontario Science Center** [20]: (DW) A science museum featuring exhibits of interest to all ages. [770 Don Mills Road ; TTC east on Bloor Subway to Pape station then north on the Don Mills 25 bus or north on the Yonge subway to Eglinton station then east on the Eglinton East bus to Don Mills Road; (416) 696-3127.]
- **The Toronto Zoo** [21]: (DW) On of the top ten zoos in North America and the 3rd largest in terms of area (comfortable walking shoes are recommended). On Meadowvale Road north of Highway 401 in the north east corner of Toronto. [TTC Take the Bloor/Danforth/ Scarborough LRT east to Kennedy Station. Take the 86A bus from Kennedy station to the Zoo. GO TRAIN: take a GO train from Union Station to Rouge Hill station in Scarborough, TTC

buses connect this station directly to the Zoo.; (416) 392-5900.]

- **The Casa Loma** [22]: Casa Loma is the former home of Canadian financier Sir Henry Pellatt. It is a castle with decorated suites, secret passages, an 800-foot tunnel, towers, stables, and beautiful 5-acre estate gardens. [TTC: go to Spadina station and take the Davenport 127 bus to Davenport & Spadina. Get off the bus and climb the Baldwin steps (110 steps ), or take the bus one stop further to Davenport and Walmer and walk up the hill on the west side of the castle.]
- **Paramount Canada Wonderland** [23]: (DW) A major amusement park featuring rides and shows. On Highway 400 about 15 Km north of Toronto. [Express GO busses run from Yorkdale and York Mills subway stations. An alternative is to take the GO train from Union Station north to Maple Station and then take the Vaughn Transit #4 bus to Canada Wonderland.; (905) 832-8131]

**Proposition 1:** The City of Toronto provides a **CityPass** [24], which provides entrance to 6 famous Toronto attractions for one-low-price. Includes tickets to the Art Gallery of Ontario, Casa Loma, CN Tower, Ontario Science Centre, Royal Ontario Museum, and Toronto Zoo. The cost is \$46.00 CAD (almost half the price) and tickets can be purchased online as well as in any of the 6 attractions.

**Proposition 2:** The Toronto theatre community has introduced T.O.TIX [25], which offers performing arts lovers the opportunity to purchase half-price tickets to a wide variety of theatre, dance, comedy, opera and music events **on the day of performance**.

**Proposition 3:** This year, the 29th International Toronto Film Festival, runs from September 9 to 18. Check under <http://www.e.bell.ca/filmfest/2003/> for more information on location, hours and tickets.

### 4. Eating in Toronto

(DW) Toronto has over 2,000 restaurants. The Toronto Health Department has instituted a mandatory restaurant inspection program. Restaurants that have passed this inspection display a large green PASS sign.

#### 4.1 Near the Conference Hotel

- **Mövenpick Marché** [26]: A unique open cafeteria restaurant in BCE place. An experience! Open for breakfast, lunch and dinner. Warning: unlike all other Toronto restaurants, Marché automatically adds a service charge to your bill so tipping is not appropriate. [East end of the street level in BCE place; (416) 366-8986.]
- **Marché Lino**: A smaller take out version of Marché in the PATH underground in BCE place.
- **C'est What** [27]: A bistro offering "ethno-lectic" menu serving multi-cultural food. It also offers 45 craft brewed beers on tap, fine whiskies from the world over, and a select all Ontario wine list [67 Front Street East; (416)-867-9499.]
- **Young Thailand** [28]: A small chain of very good Thai restaurants. Two locations nearby: 81 Church St (between King and Queen; (416) 368-1368) and 165 John St (north of Queen St.; (416) 593-9291).
- **Avalon** [29]: An excellent (expensive) high end restaurant. One of the best restaurants in the city. Open for lunch only on Thursdays. [270 Adelaide St. West (at John St.); (416) 979-9918.]
- **Barberian's Steak House** [30]: Carnivore's delight. One of Toronto's top steak restaurants. Expensive. [Elm Street between Yonge St and Bay St. a couple of block north of Dundas; (416) 597-0335.]
- **Bombay Palace** [31]: A good Indian restaurant with a buffet lunch. [71 Jarvis St. north of Front St.; (416) 368-8048.]
- **Sangam Restaurant**: An Indian restaurant with a good selection of vegetarian dishes. [788 Bay St at College St.; TTC Bay bus north to College St. ; (416) 591-7090.]
- **St. Lawrence Market**: There are several bakeries/cafes in the St. Lawrence market that sell sandwiches and other light snacks. Try a peameal bacon (real Candian bacon) sandwich or from Canada's north try an Arctic Char sandwich. [Front Street at Jarvis St.]
- **Lai Wah Heen** [32]: Very good Chinese restaurant with inventive (and expensive) dim sum. In the Chestnut Hotel [108 Chestnut St. one block east of University Ave, one block south of Dundas. (416) 977-9899.]
- **Gallery Grill** [33]: An excellent lunch-only restaurant in the Hart House student union on the University of Toronto Campus; (416)

978-2445 ; TTC north on University line to Queens Park or Museum stations. NOTE: It will be open starting September 1st.

- **Dundas Street China Town**: The area on Dundas Street between Bay and University has many other interesting Chinese restaurants.
- **Front Street**: There is a good selection of inexpensive to moderate priced restaurants on Front Street, west of University Avenue. These are mostly mass market restaurants catering to the crowds attending events at the Sky Dome.
- **Entertainment District**: There are many good restaurants and pubs in the Entertainment district.
- **PATH food courts**: For something fast and simple, there are a number of fast food lunch places (as well as some serious restaurants) in the PATH complex underneath the downtown area.
- **Harbour Front**: There are a number of restaurants with a pleasant waterfront ambience in the Harborfront area. The restaurants tend to have better ambience than food. Still a pleasant place to have a sandwich and a beer and watch the waterfront activity. The Kitchen Table convenience store [10 Queens Quay West; (416) 777-9875] sells a nice selection of sandwiches and other portable food for taking with you while you walk along the waterfront.

#### 4.2 Dining in Toronto

Other alternatives for those who want to explore Toronto's cuisine a bit more

- **Lee Garden** [34]: Chinese (Cantonese) restaurant, very popular. [331 Spadina Ave, between College and Dundas; (416) 593-9524.]
- **Nataraj** [35]: A good Indian restaurant in the Bloor Street West area. [394 Bloor St West, west of Spadina ; (416) 928-2925.]
- **La Bodega** [36]: A good French restaurant. Moderate to expensive. [30 Baldwin St; east of Spadina Ave, south of College St; (416) 977-1287.]
- **Avli** [37]: An upscale Greek restaurant in Toronto's Greektown. Inexpensive to moderate. 401 Danforth Ave. TTC east on the Bloor line to Chester station; (416) 461-9577.]
- **Mezes** : Very authentic Greek restaurant. Moderate. [456 Danforth Ave; TTC east on

the Bloor line to Chester station; (416) 778-5150.]

- **Le Trou Normand** [38]: A French provincial restaurant specializing in the cuisine of Normandy. [90 Yorkville Ave; TTC north on the Bay bus to Yorkville ; (416) 967-5956.]
- **Pangaea Restaurant** [39]: High end nouvelle (fusion) cuisine. Expensive. [1221 Bay Street, north of Bloor; TTC north on the Bay Street bus to Bloor St. ; (416) 920-2323.]
- **Boba** [40]: One of the finest restaurants in Toronto. Mediterranean/Asian fusion cuisine. Expensive. [90 Avenue Road, 3 blocks north of Bloor; TTC north to Museum station on the University line; (416) 961-2622.]
- **Sushi Inn** [41]: Very nice Japanese restaurant in the heart of the Yorkville area. Moderate [120 Cumberland Street, north of Bloor; TTC Bay Station and exit on Cumberland or Yorkville; (416) 923-9992.]

## 5. Conclusions

This document is for personal use only. Although there is not enough experimentation, it provides the basis for a pleasant and smooth stay in the city of Toronto, the place where the 30th International Conference on Very Large Databases is being held for the year 2004.

Even if you use this document as the basis for your getting around in Toronto, we encourage you to check the links provided. The author has done his best to assure the consistency and correctness of its contents; however, we would appreciate your input should there be missing or inaccurate information in it.

Finally, we are proposing to pursue several directions for further research. Firstly, we plan to explore variability of our evaluations with respect to factors such as ethno-cultural background (Canadian/Greek/Japanese/ etc.), age (student/junior professor/senior professor etc.), and position (academia/industry). We are also working on a framework for conducting more accurate evaluations, both with respect to outcomes and process. Therefore, we invite you to explore our city and propose your own recommendations therein.

## Acknowledgements

The contents presented in this document would not have been available to the attendees of the VLDB conference without the contributions of Mike Godfrey (MG) and Dave Wortman (DW), who compiled similar guides for ICSE 2001 and RE 2001, respectively. The present guide includes some updated and additional information.

## References

- [1] [http://www.city.toronto.on.ca/toronto\\_facts/index.htm](http://www.city.toronto.on.ca/toronto_facts/index.htm)
- [2] <http://www.ttc.ca/>
- [3] <http://www.ccra-adrc.gc.ca/E/pub/tg/rc4031/rc4031-e.html>
- [4] <http://www.city.toronto.on.ca/links.htm#guides>
- [5] <http://www.city.toronto.on.ca/emerg/>
- [6] <http://www.cntower.ca/>
- [7] <http://www.skydome.com/>
- [8] <http://www.theaircanadacentre.com/>
- [9] <http://www.city.toronto.on.ca/parks/island/index.htm>
- [10] <http://www.toronto.com/profile/150067/>
- [11] <http://www.harbourfrontcentre.com/>
- [12] <http://www.toronto.com/infosite/310101/>
- [13] <http://www.stlawrencemarket.com/>
- [14] <http://www.thedistillerydistrict.com/>
- [15] <http://www.rom.on.ca/>
- [16] <http://www.ago.net/>
- [17] <http://www.torontoeatoncentre.com/>
- [18] <http://www.utoronto.ca>
- [19] <http://toronto.rezrez.com/entertainment/gaytoronto/index.htm>
- [20] <http://www.ontariosciencecentre.ca/>
- [21] <http://www.torontozoo.com/>
- [22] <http://www.casaloma.org/>
- [23] <http://www.canadas-wonderland.com/>
- [24] [http://www.cntower.ca/information/13\\_info\\_rates\\_citypass.htm/](http://www.cntower.ca/information/13_info_rates_citypass.htm/)
- [25] <http://www.totix.ca/>
- [26] <http://www.toronto.com/profile/150592/>
- [27] <http://www.cestwhat.com/>
- [28] <http://www.toronto.com/profile/146442/>
- [29] <http://www.toronto.com/profile/146609/>
- [30] <http://www.toronto.com/profile/146534/>
- [31] <http://www.toronto.com/profile/146446/>
- [32] <http://www.metropolitan.com/lwh/>
- [33] [http://www.harthouse.utoronto.ca/userfiles/HTML/nts\\_3\\_1320\\_1.html](http://www.harthouse.utoronto.ca/userfiles/HTML/nts_3_1320_1.html)
- [34] <http://www.toronto.com/profile/147185/>
- [35] <http://www.toronto.com/profile/221627>
- [36] <http://www.labodegarestaurant.com/>
- [37] <http://www.toronto.com/infosite/179730/>
- [38] <http://www.toronto.com/profile/146801/>
- [39] <http://www.pangaearestaurant.com/>
- [40] <http://www.toronto.com/profile/146875/>
- [41] <http://www.toronto.com/profile/115741/>

# Map of Downtown Toronto<sup>1</sup>



<sup>1</sup> This map is brought to you by WHERE Toronto/Tourism

## SOCIAL EVENTS

The Conference banquet will take place on Wednesday, September 1st at 6:00pm at the Royal Ontario Museum.

The Royal Ontario Museum<sup>2</sup> is the largest museum in Canada. It collects and exhibits the cultural and natural history of Canada and the world. The collections and research are the basis of the ROM's international reputation; the collections are diverse in their subject matter and number more than five million objects. The ROM's research, exhibitions and educational activities increase understanding of cultural and natural diversity, their relationships, significance, preservation, and conservation.



The ROM was created by an act of the provincial government on April 16, 1912 and opened to the public on March 19, 1914. The opening ceremony was attended by His Royal Highness the Duke of Connaught. Since that time, the ROM has experienced much change and growth. There have been two major expansions to the original four-storey building. On October 12, 1933, Toronto newspapers reported that the newly opened wing was a “masterpiece of architecture”. In 1978, the ROM began a \$55 million renovation of the main ROM building, to facilitate extended research and collection activities. The Queen officially opened the Terrace Galleries in a 1984 ceremony.

Located on one of the most fashionable corners in Toronto and next to the University of Toronto, the ROM is a popular destination. From galleries of art, archaeology and science, showcasing the world's culture and natural history, to exciting public programs and events, the ROM offers a truly engaging museum experience.


The Royal Ontario Museum is located at 100 Queen's Park. The best way to reach it from the conference hotel is by subway. You need to take the Yonge-University subway line and step out at the “Museum” Station.

---

<sup>2</sup> Copyright © 2003 Royal Ontario Museum

Middleware is Everywhere.

Can you see it?



**DB2**

**Key**

1. Verifies insurance on the spot.
2. Files digital claim in an instant.
3. Approves estimate at the site.
4. Orders new bumper at the scene.
5. Receives settlement in a snap.

**MIDDLEWARE IS IBM SOFTWARE.** Software like IBM DB2 Content Manager. A comprehensive and open solution that easily manages and leverages information from almost anywhere. Even content like video and scanned images are easily and securely accessed. It's how responsiveness increases, productivity soars and knowledge becomes power.

@business on demand™

Go to [ibm.com/db2/middleware](http://ibm.com/db2/middleware)

IBM, the IBM logo, the e-business logo and e-business on demand are trade-marks or registered trade-marks of International Business Machines Corporation in the United States, other countries, or both, and are used under licence by IBM Canada Ltd. Other company, product and services names may be trade-marks or service marks of others. 2004 IBM Corporation. All rights reserved.

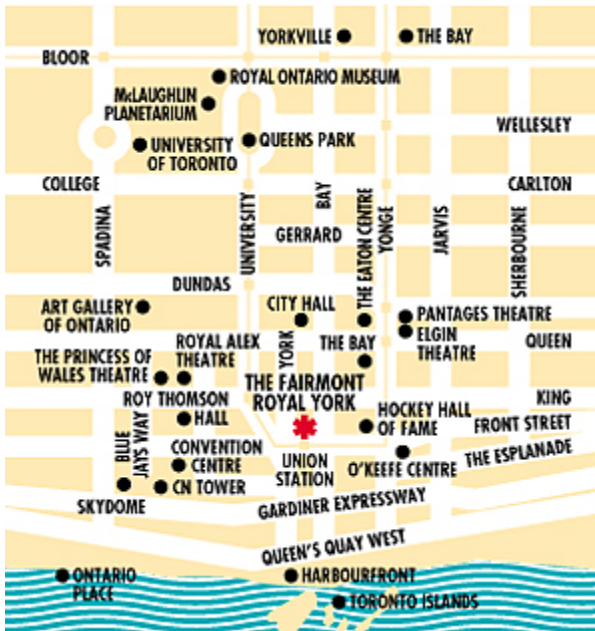
## CONFERENCE VENUE

The conference will take place at the Fairmont Royal York hotel. Just steps away from our famous doors in the heart of Canada's largest metropolis is an exciting mix of activities and attractions that will leave you exhilarated. From the theater, entertainment and financial districts, to shopping, sightseeing, and world-class sports facilities. The Fairmont Royal York truly is "at the center of it all."



The main Toronto airport is the "Pearson International Airport". To get from the airport to the hotel, you can use a taxi that can take you to the city in 50 minutes for approximately \$40. The "Airport Express" bus does the same route for approx \$16 in 1 hour and 15 minutes. The "Pacific Western" bus and they travel from the airport to all downtown hotels and vice versa. They charge \$ 14.95. Alternatively, you can use the regular Toronto transportation (TTC) for \$2.25. In particular, take the 192 Airport Rocket bus at Terminals 1, 2 or 3 to Kipling Station. At Kipling, take the Bloor-Danforth subway, change at Yonge/Bloor Station and take the Yonge-University-Spadina Line (southbound trains). Exit at Union Station. The 192 Airport Rocket bus operates every 25 minutes. You will need approximately 30 minutes to Islington subway station and from there 45 minutes ride to get downtown.

To drive there:



From Highway 401 (Eastbound or Westbound):

- Exit on the Don Valley Parkway (DVP) Southbound
- Exit on the Gardiner Expressway Westbound
- Exit on Yonge Street North
- Left on Front Street (Westbound)
- Pass one set of lights (Bay Street)
- Hotel is on right-hand side, across from Union Station

## From New York:

- Cross at the Peace Bridge at Fort Erie, take the Queen Elizabeth Way (QEW) East
- QEW becomes the Gardiner Expressway
- Exit at Bay Street North
- Take Bay Street North to Front Street
- Turn left on Front Street
- Hotel is on the north side

## From Michigan:

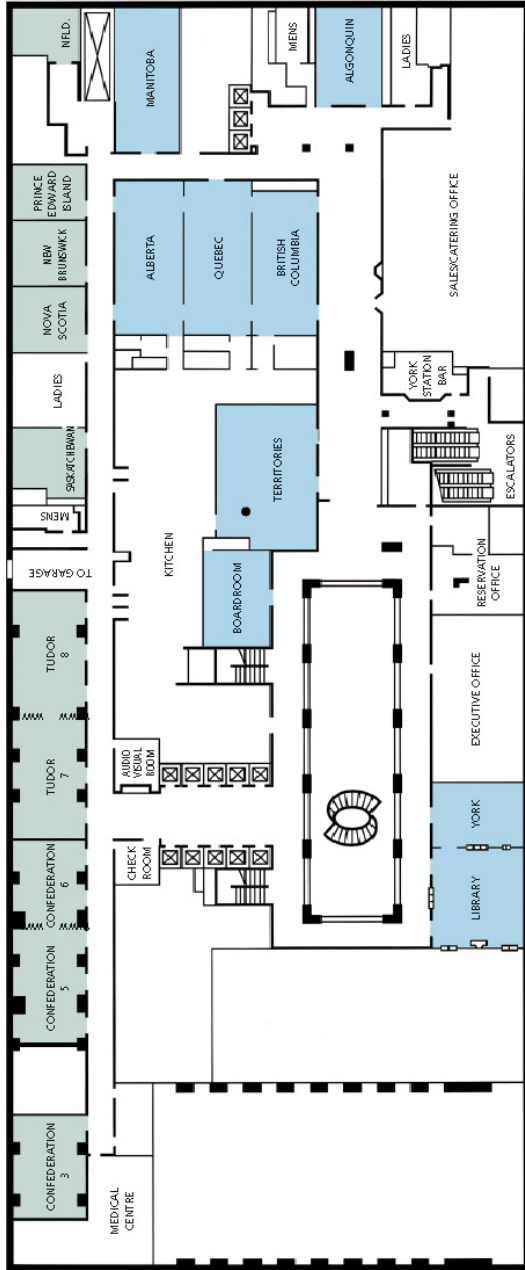
- Cross at Windsor and take Hwy 401 East
- Take Hwy 427 South to Queen Elizabeth Way (QEW) East
- QEW becomes Gardiner Expressway
- Exit at Bay Street North
- Left on Front Street
- Hotel is on the north side of Front Street

## From Quebec:

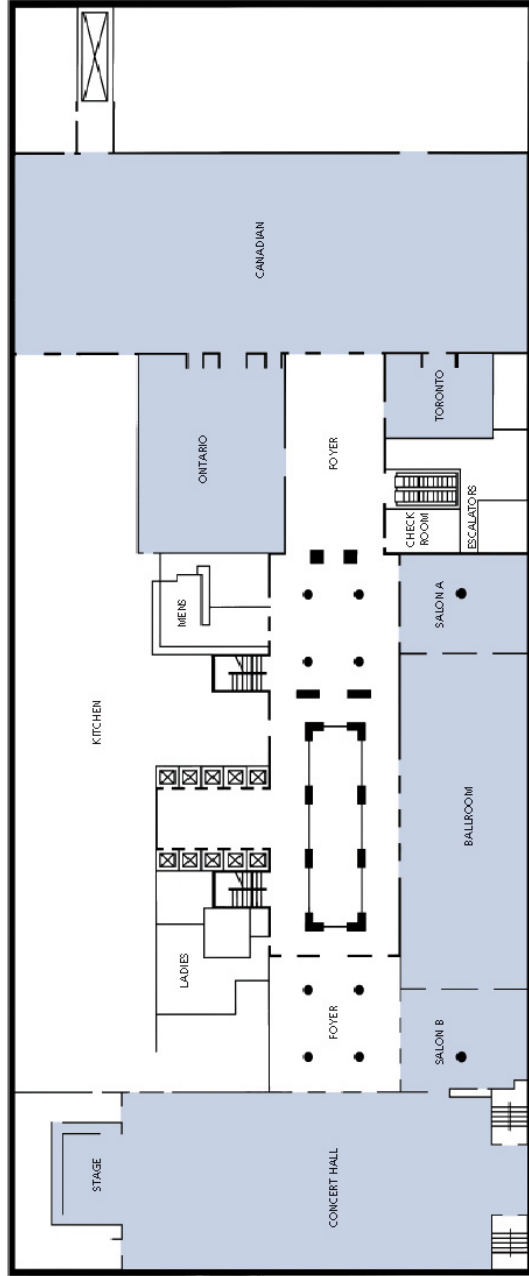
- Take Hwy 401 West to Don Valley Parkway (DVP) South to Gardiner Expressway West
- Exit north on Yonge Street left on Front Street (west)
- Pass the 1st set of lights (Bay Street)
- Hotel is on the right-hand side, across from Union station



MAIN MEZZANINE FLOOR



CONVENTION FLOOR



**MONDAY AT A GLANCE**

	Upper Canada
18:00-20:00	Welcome Reception (Upper Canada)

**TUESDAY AT A GLANCE**

	Ballroom	Ontario	Territories	Confederation 5 & 6	Tudor 7 & 8
07:30-08:30	Breakfast (Concert Hall)				
08:30-09:00	Opening Ceremony (Ballroom)				
09:00-10:30	Keynote Address 1 (Ballroom) : Databases in a Wireless World David Yach (VP Software, Research in Motion)				
10:30-11:00	Coffee Break (Ballroom Foyer)				
11:00-12:30	Research Session 1 Compression & Indexing	Research Session 2 XML Views and Schemas	Research Session 3 Controlling Access	Tutorial 1 Database Architectures for New Hardware	NO SESSION
12:30-14:00	Lunch Break (Concert Hall)				
14:00-15:30	Research Session 4 XML (I)	Research Session 5 Stream Mining	Industrial Session 1 Novel SQL Extensions	Tutorial 1 (cont) Database Architectures for New Hardware	Demo Groups 2 and 3
15:30-16:00	Coffee Break (Ballroom Foyer)				
16:00-17:30	Research Session 6 XML (II)	Industrial Session 2 New DBMS architectures and Performance	Panel 1 Biological Data Management: Research, Practice and Opportunities	Tutorial 2 Security of Shared Data in Large Systems: State of the Art and Research Directions	Demo Groups 2 and 3

**WEDNESDAY AT A GLANCE**

	Ballroom	Ontario	Territories	Confederation 5 & 6	Tudor 7 & 8
07:45-08:45	Breakfast (Concert Hall)				
09:00-10:30	Keynote Address 2 (Ballroom) : Structures, Semantics and Statistics Alon Halevy (University of Washington)				Demo Groups 1 and 2
10:30-11:00	Coffee Break (Ballroom Foyer)				
11:00-12:30	Research Session 7 XML and Relations	Research Session 8 Stream Mining (II)	Industrial Session 3 Semantic Query Approaches	Tutorial 2 (cont) Security of Shared Data in Large Systems: State of the Art and Research Directions	Demo Groups 1 and 2
12:30-14:00	Lunch Break (Concert Hall)				
14:00-16:00	Research Session 9 Stream Query Processing	Research Session 10 Managing Web Information Sources	Research Session 11 Distributed Search and Query Processing	Tutorial 3 Self-Managing Technology in Database Management Systems	Demos Groups 4 and 5
16:00-16:30	Coffee Break (Ballroom Foyer)				
16:30-17:30	Research Session 12 Stream Data Management Systems	Research Session 13 Auditing	Research Session 14 Data Warehousing	Tutorial 3 (cont) Self-Managing Technology in Database Management Systems	Demos Groups 4 and 5
18:00-21:00	Conference Banquet (Royal Ontario Museum)				

**THURSDAY AT A GLANCE**

	Ballroom	Ontario	Territories	Confederation 5 & 6	Tudor 7 & 8
07:45-08:45	Breakfast (Concert Hall)				
09:00-10:30	<b>10 Year Best Paper Award (Ballroom) : Whither Data Mining?</b> Rakesh Agrawal, Ramakrishnan Srikant (IBM Almaden Research Center)				Demo Groups 3 and 4
10:30-11:00	Coffee Break (Ballroom Foyer)				
11:00-12:30	Research Session 15 Link Analysis	Research Session 16 Sensors, Grid, Publish/Subscribe Systems	Industrial Session 4 Automatic Tuning in Commercial DBMSs	Panel 2 Where is Business Intelligence Taking Today's Database Systems?	Demo Groups 3 and 4
12:30-14:00	Lunch Break (Concert Hall)				
14:00-15:30	Research Session 17 Top-K Ranking	Research Session 18 DBMS Architecture and Performance	Industrial Session 5 XML Implementations, Automatic Physical Design, and Indexing	Tutorial 4 Architectures and Algorithms for Internet-Scale (P2P) Data Management	Demo Groups 1 and 6
15:30-16:00	Coffee Break (Ballroom Foyer)				
16:00-17:30	Research Session 19 Privacy	Research Session 20 Nearest Neighbor Search	Industrial Session 6 Data Management with RFIDs and Ease of Use	Tutorial 4 (cont) Architectures and Algorithms for Internet-Scale (P2P) Data Management	Demo Groups 1 and 6

**FRIDAY AT A GLANCE**

	Ballroom	Ontario	Territories	Confederation 5 & 6	Tudor 7 & 8
07:45-08:45	Breakfast (Concert Hall)				
09:00-10:30	Research Session 21 Similarity Search and Applications	Research Session 22 Query Processing	Industrial Session 7 Data Management Challenges in Life Sciences and Email Systems	Tutorial 5 The Continued Saga of DB-IR Integration	Demo Groups 5 and 6
10:30-11:00	Coffee Break (Ballroom Foyer)				
11:00-12:30	Research Session 23 Novel Models	Research Session 24 Query Processing and Optimization	Industrial Session 8 Issues in Data Warehousing	Tutorial 5 (cont) The Continued Saga of DB-IR Integration	Demos Groups 5 and 6

\* Check also the Conference Co-Located Workshops section of the booklet (page 103) for a description of the many interesting workshops

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page's vertical space below the 'NOTES' header.

## CONFERENCE PROGRAM

**SUNDAY, MONDAY 8:00 – 17:00**

**Workshops Registration**

**Location: British Columbia Foyer**

**MONDAY 8:00 – 17:00**

**Conference Registration**

**Location: British Columbia Foyer**

**MONDAY 18:00 – 20:00**

**Welcome Reception**

**Location: Imperial Room**

**TUESDAY - FRIDAY 8:00 – 17:00**

**Conference Registration**

**Location: Ballroom Foyer**

**TUESDAY 8:30 – 9:00**

**Opening Ceremony**

**Location: Ballroom**

**TUESDAY 9:00 – 10:30**

**Keynote Address 1**

**Location: Ballroom**

**Session Chair: M. Tamer Özsu**

**Databases in a Wireless World**

David Yach, VP Software, Research in Motion

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.

**TUESDAY 11:00 – 12:30**

**Research Session 1: Compression & Indexing**

**Location:** Ballroom

**Session Chair:** Jayant Haritsa

**Compressing Large Boolean Matrices using Reordering Techniques**

David S. Johnson, Shankar Krishnan (AT&T Labs-Research), Jatin Chhugani, Subodh Kumar (John Hopkins University), Suresh Venkatasubramanian (AT&T Labs-Research)

**On the Performance of Bitmap Indices for High Cardinality Attributes**

Kesheng Wu, Ekow Otoo, Arie Shoshani (Lawrence Berkeley National Laboratory)

**Practical Suffix Tree Construction**

Sandeep Tata, Richard A. Hankins, Jignesh M. Patel (University of Michigan)

**Research Session 2: XML Views and Schemas**

**Location:** Ontario

**Session Chair:** Alberto Mendelzon

**Answering XPath Queries over Networks by Sending Minimal Views**

Keishi Tajima, Yoshiki Fukui (JAIST)

**A Framework for Using Materialized XPath Views in XML Query Processing**

Andrey Balmin, Fatma Özcan, Kevin Beyer, Roberta Cochrane, Hamid Pirahesh (IBM Almaden)

**Schema-Free XQuery**

Yunyao Li, Cong Yu, H. V. Jagadish (University of Michigan)

**Research Session 3: Controlling Access**

**Location:** Territories

**Session Chair:** Mariano Consens

**Client-Based Access Control Management for XML Documents**

Luc Bouganim (INRIA Rocquencourt), François Dang Ngoc, Philippe Pucheral (PRiSM Laboratory)

**Secure XML Publishing without Information Leakage in the Presence of Data Inference**

Xiaochun Yang (Northeastern University), Chen Li (University of California, Irvine)

**Limiting Disclosure in Hippocratic Databases**

Kristen LeFevre (IBM Almaden and University of Wisconsin-Madison), Rakesh Agrawal (IBM Almaden), Vuk Ercegovic, Raghu Ramakrishnan (University of Wisconsin-Madison), Yirong Xu (IBM Almaden), David DeWitt (University of Wisconsin-Madison)

**Tutorial 1: Database Architectures for New Hardware**

**Location:** Confederation 5 & 6

**Presenter:**

Anastassia Ailamaki (Carnegie Mellon University)



**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page's vertical space below the 'NOTES' header.

**TUESDAY 14:00 – 15:30**

**Research Session 4: XML (I)**

**Location: Ballroom**

**Session Chair:** *David Toman*

**On Testing Satisfiability of Tree Pattern Queries**

Laks Lakshmanan, Ganesh Ramesh, Hui Wang, Zheng Zhao (University of British Columbia)

**Containment of Nested XML Queries**

Xin Dong, Alon Halevy, Igor Tatarinov (University of Washington)

**Efficient XML-to-SQL Query Translation: Where to Add the Intelligence? (short presentation)**

Rajasekar Krishnamurthy (IBM Almaden), Raghav Kaushik (Microsoft Research), Jeffrey Naughton (University of Wisconsin-Madison)

**Taming XPath Queries by Minimizing Wildcard Steps (short presentation)**

Chee-Yong Chan (National University of Singapore), Wenfei Fan (University of Edinburgh & Bell Laboratories), Yiming Zeng (National University of Singapore)

**The NEXT Framework for Logical Xquery Optimization (short presentation)**

Alin Deutsch, Yannis Papakonstantinou, Yu Xu (University of California, San Diego)

**Research Session 5: Stream Mining**

**Location: Ontario**

**Session Chair:** *Christopher Jermaine*

**Detecting Change in Data Streams**

Daniel Kifer, Shai Ben-David, Johannes Gehrke (Cornell University)

**Stochastic Consistency, and Scalable Pull-Based Caching for Erratic Data Stream Sources**

Shanzhong Zhu, Chinya V. Ravishankar (University of California, Riverside)

**False Positive or False Negative: Mining Frequent Itemsets from High Speed Transactional Data Streams**

Jeffrey Xu Yu (The Chinese University of Hong Kong), Zhihong Chong (Fudan University), Hongjun Lu (The Hong Kong University of Science and Technology), Aoying Zhou (Fudan University)

**Industrial Session 1: Novel SQL Extensions**

**Location: Territories**

**Session Chair:** *Sophie Cluet*

**Returning Modified Rows - SELECT Statements with Side Effects**

Andreas Behm, Serge Rielau, Richard Swagerman (IBM Toronto Lab)

**PIVOT and UNPIVOT: Optimization and Execution Strategies in an RDBMS**

Conor Cunningham, César Galindo-Legaria, Goetz Graefe (Microsoft Corp.)

**A Multi-Purpose Implementation of Mandatory Access Control in Relational Database Management Systems**

Walid Rjaibi, Paul Bird (IBM Toronto Lab)

**Tutorial 1: Database Architectures for New Hardware (cont)**

**Location: Confederation 5 & 6**

**Presenter:**

Anastassia Ailamaki (Carnegie Mellon University)

**Demo Session 2 & 3**

**Location: Tudor 7 & 8**

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.

**TUESDAY 16:00 – 17:30****Research Session 6: XML (II)****Location:** Ballroom**Session Chair:** Alin Deutsch**Indexing Temporal XML Documents**

Alberto Mendelzon, Flavio Rizzolo (University of Toronto), Alejandro Vaisman (University of Buenos Aires)

**Schema-based Scheduling of Event Processors and Buffer Minimization for Queries on Structured Data Streams**

Christoph Koch, Stefanie Scherzinger (Technische Universität Wien), Nicole Schweikardt (Humboldt Universität zu Berlin), Bernhard Stegmaier (Technische Univ. München)

**Bloom Histogram: Path Selectivity Estimation for XML Data with Updates**

Wei Wang (University of NSW), Haifeng Jiang, Hongjun Lu (Hong Kong University of Science and Technology), Jeffrey Xu Yu (The Chinese University of Hong Kong)

**Industrial Session 2: New DBMS Architectures and Performance****Location:** Ontario**Session Chair:** Umesh Dayal**Hardware Acceleration in Commercial Databases: A Case Study of Spatial Operations**

Nagender Bandi (University of California, Santa Barbara), Chengyu Sun (California State University, Los Angeles), Divyakant Agrawal, Amr El Abbadi (University of California, Santa Barbara)

**P\*TIME: Highly Scalable OLTP DBMS for Managing Update-Intensive Stream Workload**

Sang K. Cha (Transact In Memory, Inc.), Changbin Song (Seoul National University)

**Generating Thousand Benchmark Queries in Seconds**

Meikel Poess (Oracle Corp.), John M. Stephens, Jr. (Gradient Systems)

**Panel 1: Biological Data Management: Research, Practice and Opportunities****Location:** Territories**Moderator:**

Thodoros Topaloglou (MDS Proteomics)

**Panelists:**

Susan B. Davidson (University of Pennsylvania), H. V. Jagadish (University of Michigan, Ann Arbor), Victor M. Markowitz (JGI and Lawrence Berkeley National Laboratory), Evan W. Steeg (Independent Consultant and Author, and former Founder/CEO of Molecular Mining Corp.), Mike Tyers (Samuel Lunenfeld Research Institute and U. of Toronto)

**Tutorial 2: Security of Shared Data in Large Systems: State of the Art and Research Directions****Location:** Confederation 5 & 6**Presenter:**

Annie Rosenthal (MITRE) and Marianne Winslett (University of Illinois)

**Demo Session 2 & 3 (cont)****Location:** Tudor 7 & 8

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.

**WEDNESDAY 9:00 – 10:30****Keynote Address 2****Location:** Ballroom**Session Chair:** Donald Kossmann**Structures, Semantics and Statistics**

Alon Halevy (U. of Washington)

**Demo Session 1 & 2****Location:** Tudor 7 & 8**WEDNESDAY 11:00 – 12:30****Research Session 7: XML and Relations****Location:** Ballroom**Session Chair:** Christoph Koch**XQuery on SQL Hosts**

Torsten Grust, Sherif Sakr, Jens Teubner (University of Konstanz)

**ROX: Relational Over XML**

Alan Halverson (University of Wisconsin-Madison), Vanja Josifovski, Guy Lohman, Hamid Pirahesh (IBM Almaden), Mathias Mörschel

**From XML View Updates to Relational View Updates: Old Solutions to a New Problem**

Vanessa Braganholo (Universidade Federal do Rio Grande do Sul), Susan Davidson (University of Pennsylvania, and INRIA-FUTURS), Carlos Heuser (Universidade Federal do Rio Grande do Sul)

**Research Session 8: Stream Mining (II)****Location:** Ontario**Session Chair:** Johannes Gehrke**XWAVE: Optimal and Approximate Extended Wavelets for Streaming Data**

Sudipto Guha (University of Pennsylvania), Chulyun Kim, Kyuseok Shim (Seoul National University)

**REHIST: Relative Error Histogram Construction Algorithms**

Sudipto Guha (University of Pennsylvania), Kyuseok Shim, Jungchul Woo (Seoul National University)

**Distributed Set-Expression Cardinality Estimation**

Abhinandan Das (Cornell University), Sumit Ganguly (IIT Kanpur), Minos Garofalakis, Rajeev Rastogi (Bell Labs, Lucent Technologies)

**Industrial Session 3: Semantic Query Approaches****Location:** Territories**Session Chair:** Erhard Rahm**Supporting Ontology-Based Semantic Matching in RDBMS**

Souripriya Das, Eugene Chong, George Eadon, Jagannathan Srinivasan (Oracle Corp.)

**BioPatentMiner: An Information Retrieval System for BioMedical Patents**

Sougata Mukherjee, Bhuvan Bamba (IBM India Research Lab)

**Flexible String Matching Against Large Databases in Practice**

Nick Koudas, Amit Marathe, Divesh Srivastava (AT&amp;T Labs-Research)

**Tutorial 2: Security of Shared Data in Large Systems: State of the Art and Research Directions (cont.)****Location:** Confederation 5 & 6**Presenter:**

Arnie Rosenthal (MITRE) and Marianne Winslett (University of Illinois)

**Demo Session 1 & 2 (cont)****Location:** Tudor 7 & 8

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.

**WEDNESDAY 14:00 – 16:00****Research Session 9: Stream Query Processing****Location:** Ballroom**Session Chair:** *Nick Koudas***Memory-Limited Execution of Windowed Stream Joins**

Utkarsh Srivastava, Jennifer Widom (Stanford University)

**Resource Sharing in Continuous Sliding-Window Aggregates**

Arvind Arasu, Jennifer Widom (Stanford University)

**Remembrance of Streams Past: Overload-Sensitive Management of Archived Streams**

Sirish Chandrasekaran, Michael J. Franklin (University of California, Berkeley)

**WIC: A General-Purpose Algorithm for Monitoring Web Information Sources**

Sandeep Pandey, Kedar Dhamdhere, Christopher Olston (Carnegie Mellon University)

**Research Session 10: Managing Web Information Sources****Location:** Ontario**Session Chair:** *Davood Rafiei***Similarity Search for Web Services**

Xin Dong, Alon Halevy, Jayant Madhavan, Ema Nemes, Jun Zhang (University of Washington)

**AWESOME: A Data Warehouse-based System for Adaptive Website Recommendations**

Andreas Thor, Erhard Rahm (University of Leipzig)

**Accurate and Efficient Crawling for Relevant Websites**

Martin Ester (Simon Fraser University), Hans-Peter Kriegel, Matthias Schubert (University of Munich)

**Instance-based Schema Matching for Web Databases by Domain-specific Query Probing**

Jiyang Wang (Hong Kong University of Science and Technology), Ji-Rong Wen (Microsoft Research Asia), Fred Lochovsky (Hong Kong University of Science and Technology), Wei-Ying Ma (Microsoft Research Asia)

**Research Session 11: Distributed Search and Query Processing****Location:** Territories**Session Chair:** *Karl Aberer***Computing PageRank in a Distributed Internet Search System**

Yuan Wang, David DeWitt (University of Wisconsin-Madison)

**Enhancing P2P File-Sharing with an Internet-Scale Query Processor**

Boon Thau Loo (University of California, Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley), Ryan Huebsch (University of California, Berkeley), Scott Shenker (University of California, Berkeley and International Computer Science Institute), Ion Stoica (University of California, Berkeley)

**Online Balancing of Range-Partitioned Data with Applications to Peer-to-Peer Systems**

Prasanna Ganesan, Mayank Bawa, Hector Garcia-Molina (Stanford University)

**Network-Aware Query Processing for Stream-Based Applications (short presentation)**

Yanif Ahmad, Uğur Çetintemel (Brown University)

**Data Sharing Through Query Translation in Autonomous Sources (short presentation)**

Anastasios Kementsietsidis, Marcelo Arenas (University of Toronto)

**Tutorial 3: Self-Managing Technology in Database Management Systems****Location:** Confederation 5 & 6**Presenters:**

Surajit Chaudhuri (Microsoft Research), Benoit Dageville (Oracle Corp.), Guy Lohman (IBM Almaden)

**Demo Session 4 & 5****Location:** Tudor 7 & 8



**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.

**WEDNESDAY 16:30 – 17:30****Research Session 12: Stream Data Management Systems****Location:** Ballroom**Session Chair:** *Pat Martin***Linear Road: A Stream Data Management Benchmark**

Arvind Arasu (Stanford University), Mitch Cherniack, Eduardo Galvez (Brandeis University), David Maier (OHSU/OGI), Anurag Maskey, Esther Ryvkina (Brandeis University), Michael Stonebraker, Richard Tibbetts (MIT)

**Query Languages and Data Models for Database Sequences and Data Streams**

Yan-Nei Law (UCLA), Haixun Wang (IBM T. J. Watson Research), Carlo Zaniolo (UCLA)

**Research Session 13: Auditing****Location:** Ontario**Session Chair:** *Rajeev Rastogi***Tamper Detection in Audit Logs**

Richard Snodgrass, Shilong Yao, Christian Collberg (University of Arizona)

**Auditing Compliance with a Hippocratic Database**

Rakesh Agrawal, Roberto Bayardo, Christos Faloutsos, Jerry Kiernan, Ralf Rantzau, Ramakrishnan Srikant (IBM Almaden)

**Research Session 14: Data Warehousing****Location:** Territories**Session Chair:** *Martin Kersten***High-Dimensional OLAP: A Minimal Cubing Approach**

Xiaolei Li, Jiawei Han, Hector Gonzalez (University of Illinois at Urbana-Champaign)

**The Polynomial Complexity of Fully Materialized Coalesced Cubes**

Yannis Sismanis, Nick Roussopoulos (University of Maryland)

**Tutorial 3: Self-Managing Technology in Database Management Systems (cont)****Location:** Confederation 5 & 6**Presenters:**

Surajit Chaudhuri (Microsoft Research), Benoit Dageville (Oracle Corp.), Guy Lohman (IBM Almaden)

**Demo Session 4 & 5 (cont)****Location:** Tudor 7 & 8**THURSDAY 9:00 – 10:30****Ten-Year Best Paper Award****Location:** Ballroom**Session Chair:** *Renée J. Miller***Whither Data Mining?**

Rakesh Agrawal, Ramakrishnan Srikant (IBM Almaden Research Center)

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.

**THURSDAY 11:00 – 12:30****Research Session 15: Link Analysis****Location:** Ballroom**Session Chair:** *Divesh Srivastava***Relational Link-based Ranking**

Floris Geerts (University of Edinburgh), Heikki Mannila, Evimaria Terzi (University of Helsinki)

**ObjectRank: Authority-Based Keyword Search in Databases**

Andrey Balmin (IBM Almaden), Vagelis Hristidis (Florida International University), Yannis Papakonstantinou (University of California, San Diego)

**Combating Web Spam with TrustRank**

Zoltan Gyöngyi, Hector Garcia-Molina (Stanford University), Jan Pedersen (Yahoo! Inc.)

**Research Session 16: Sensors, Grid, Pub/Sub****Location:** Ontario**Session Chair:** *Alexandros Labrinidis***Model-Driven Data Acquisition in Sensor Networks (Best Paper Award winner)**

Amol Deshpande (University of California, Berkeley), Carlos Guestrin (Intel Research Berkeley), Samuel Madden (MIT and Intel Research Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley), Wei Hong (Intel Research Berkeley)

**GridDB: A Data-Centric Overlay for Scientific Grids**

David Liu, Michael J. Franklin (University of California, Berkeley)

**Towards an Internet-Scale XML Dissemination Service**

Yanlei Diao, Shariq Rizvi, Michael J. Franklin (University of California, Berkeley)

**Industrial Session 4: Automatic Tuning in Commercial DBMSs****Location:** Territories**Session Chair:** *Berni Schiefer***DB2 Design Advisor: Integrated Automatic Physical Database Design**

Daniel Zilio (IBM Toronto Lab), Jun Rao (IBM Almaden), Sam Lightstone (IBM Toronto Lab), Guy Lohman (IBM Almaden), Adam Storm, Christian, Garcia-Arellano (IBM Toronto Lab), Scott Fadden (IBM Portland)

**Automatic SQL Tuning in Oracle 10g**

Benoit Dageville, Dinesh Das, Karl Dias, Khaled Yagoub, Mohamed Zait, Mohamed Ziauddin (Oracle Corp.)

**Database Tuning Advisor for Microsoft SQL Server 2005**

Sanjay Agrawal, Surajit Chaudhuri, Lubor Kollar, Arun Marathe, Vivek Narasayya, Manoj Syamala (Microsoft Corp.)

**Panel 2:****Location:** Confederation 5 & 6**Moderator:**

William O'Connell (IBM Toronto Lab)

**Panelists:**

Andy Witkowski (Oracle), Ramesh Bhashyam (NCR-Teradata), Surajit Chaudhuri (Microsoft Research), Nigel Campbell

**Demo Session 3 & 4 (cont)****Location:** Tudor 7 & 8

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page's vertical space below the 'NOTES' header.

**THURSDAY 14:00 – 15:30****Research Session 17: Top-K Ranking****Location:** Ballroom**Session Chair:** Themis Palpanas**Efficiency-Quality Tradeoffs for Vector Score Aggregation**

Pavan Kumar C Singitham, Mahathi Mahabhashyam (Stanford University), Prabhakar Raghavan (Verity Inc.)

**Merging the Results of Approximate Match Operations**

Sudipto Guha (University of Pennsylvania), Nick Koudas, Amit Marathe, Divesh Srivastava (AT&amp;T Labs-Research)

**Top-k Query Evaluation with Probabilistic Guarantees**

Martin Theobald, Gerhard Weikum, Ralf Schenkel (Max-Planck Institute of Computer Science)

**Research Session 18: DBMS Architecture and Performance****Location:** Ontario**Session Chair:** Ken Salem**STEPS Towards Cache-resident Transaction Processing**

Stavros Harizopoulos, Anastassia Ailamaki (Carnegie Mellon University)

**Write-Optimized B-Trees**

Goetz Graefe (Microsoft)

**Cache-Conscious Radix-Decluster Projections (short presentation)**

Stefan Manegold, Peter Boncz, Martin Kersten, Niels Nes (CWI)

**Clotho: Decoupling Memory Page Layout from Storage Organization (short presentation)**

Minglong Shao, Jiri Schindler, Steven Schlosser, Anastassia Ailamaki, Gregory Ganger (Carnegie Mellon University)

**Industrial Session 5: XML Implementations, Automatic Physical Design and Indexing****Location:** Territories**Session Chair:** Viswanathan Krishnamurthy**Query Rewrite for XML in Oracle XML DB**

Muralidhar Krishnaprasad, Zhen Liu, Anand Manikutty, James W. Warner, Vikas Arora, Susan Kotsovolos (Oracle Corp.)

**Indexing XML Data Stored in a Relational Database**

Shankar Pal, Istvan Cseri, Gideon Schaller, Oliver Seeliger, Leo Giakoumakis, Vasili Zolotov (Microsoft Corp.)

**Automated Statistics Collection in DB2 UDB (short presentation)**

Ashraf Aboulnaga, Peter Haas (IBM Almaden), Mokhtar Kandil, Sam Lightstone (IBM Toronto Lab), Guy Lohman, Volker Markl (IBM Almaden), Ivan Popivanov (IBM Toronto Lab), Vijayshankar Raman (IBM Almaden)

**High Performance Index Build Algorithms for Intranet Search Engines (short presentation)**

Marcus Fontoura, Eugene Shekita, Jason Zien, Sridhar Rajagopalan, Andreas Neumann (IBM Almaden)

**Automating the Design of Multi-dimensional Clustering Tables in Relational Databases (short presentation)**

Sam Lightstone (IBM Toronto Lab), Bishwaranjan Bhattacharjee (IBM T.J. Watson Research Center)

**Tutorial 4: Architectures and Algorithms for Internet-Scale (P2P) Data Management****Location:** Confederation 5 & 6**Presenter:**

Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley)

**Demo Session 1 & 6****Location:** Tudor 7 & 8

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page below the 'NOTES' header.

## THURSDAY 16:00 – 17:30

### Research Session 19: Privacy

#### Location: Ballroom

**Session Chair:** *S. Sudarshan*

#### Viewion paper: Enabling Privacy for the Paranoids

Gagan Aggarwal, Mayank Bawa, Prasanna Ganesan, Hector Garcia-Molina, Krishnaram Kenthapadi, Nina Mishra, Rajeev Motwani, Utkarsh Srivastava, Dilys Thomas, Jennifer Widom, Ying Xu (Stanford University)

#### A Privacy-Preserving Index for Range Queries

Bijit Hore, Sharad Mehrotra, Gene Tsudik (University of California, Irvine)

#### Resilient Rights Protection for Sensor Streams

Radu Sion, Mikhail Atallah, Sunil Prabhakar (Purdue University)

### Research Session 20: Nearest Neighbor Search

#### Location: Ontario

**Session Chair:** *Hans-Peter Kriegel*

#### Reverse kNN Search in Arbitrary Dimensionality

Yufei Tao (City University of Hong Kong), Dimitris Papadias, Xiang Lian (Hong Kong University of Science and Technology)

#### GORDER: An Efficient Method for KNN Join Processing

Chenyi Xia (National University of Singapore), Hongjun Lu (Hong Kong University of Science and Technology), Beng Chin Ooi, Jin Hu (National University of Singapore)

#### Query and Update Efficient B+-Tree Based Indexing of Moving Objects

Christian S. Jensen (Aalborg University), Dan Lin, Beng Chin Ooi (National University of Singapore)

### Industrial Session 6: Data Management with RFIDs and Ease of Use

#### Location: Territories

**Session Chair:** *José A. Blakeley*

#### Integrating Automatic Data Acquisition with Business Processes - Experiences with SAP's Auto-ID Infrastructure (invited)

Christof Bornhövd, Tao Lin, Stephan Haller, Joachim Schaper (SAP Research)

#### Managing RFID Data (invited)

Sudarshan S. Chawathe (University of Maryland), Venkat Krishnamurthy, Sridhar Ramachandran (OAT Systems), and Sanjay Sarma (MIT)

#### Production Database Systems: Making Them Easy is Hard Work (invited)

David Campbell (Microsoft Corporation)

### Tutorial 4: Architectures and Algorithms for Internet-Scale (P2P) Data Management

#### Location: Confederation 5 & 6

**Presenter:**

Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley)

### Demo Session 1 & 6

#### Location: Tudor 7 & 8



**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page's vertical space below the 'NOTES' header.

FRIDAY 9:00 – 10:30

**Research Session 21: Similarity Search and Applications****Location: Ballroom****Session Chair:** *Christian Jensen***Indexing Large Human-Motion Databases**

Eamonn Keogh, Themistoklis Palpanas, Victor B. Zordan, Dimitrios Gunopulos (University of California, Riverside), Marc Cardle (University of Cambridge)

**On The Marriage of Lp-norms and Edit Distance (short presentation)**

Lei Chen (University of Waterloo), Raymond Ng (University of British Columbia)

**Approximate NN queries on Streams with Guaranteed Error/performance Bounds (short presentation)**

Nick Koudas (AT&amp;T Labs-Research), Beng Chin Ooi, Kian-Lee Tan, Rui Zhang (National University of Singapore)

**Object Fusion in Geographic Information Systems (short presentation)**

Catriel Beeri, Yaron Kanza, Elyahu Safra, Yehoshua Sagiv (The Hebrew University)

**Maintenance of Spatial Semijoin Queries on Moving Points (short presentation)**

Glenn Iwerks, Hanan Samet (University of Maryland at College Park), Kenneth Smith (MITRE Corp.)

**Voronoi-Based K Nearest Neighbor Search for Spatial Network Databases (short presentation)**

Mohammad Kolahdouzan, Cyrus Shahabi (University of Southern California)

**A Framework for Projected Clustering of High Dimensional Data Streams (short presentation)**

Charu Aggarwal (T.J. Watson Res. Ctr.), Jiawei Han, Jianyong Wang (UIUC), Philip Yu (T.J. Watson Res. Ctr.)

**Research Session 22: Query Processing****Location: Ontario****Session Chair:** *Joseph M. Hellerstein***Efficient Query Evaluation on Probabilistic Databases**

Nilesh Dalvi, Dan Suciu (University of Washington)

**Efficient Indexing Methods for Probabilistic Threshold Queries over Uncertain Data**

Reynold Cheng, Yuni Xia, Sunil Prabhakar, Rahul Shah, Jeffrey S. Vitter (Purdue University)

**Probabilistic Ranking of Database Query Results**

Surajit Chaudhuri, Gautam Das (Microsoft Research), Vagelis Hristidis (Florida International University), Gerhard Weikum (MPI Informatik)

**Industrial Session 7: Data Management Challenges in Life Sciences and Email Systems****Location: Territories****Session Chair:** *Guy Lohman***Managing Data from High-Throughput Genomic Processing: A Case Study (invited)**

Toby Bloom, Ted Sharpe (Broad Institute of MIT and Harvard)

**Database Challenges in the Integration of Biomedical Data Sets (invited)**

Rakesh Nagarajan (Washington University, School of Medicine), Mushtaq Ahmed (Persistent Systems Pvt. Ltd.), Aditya Phatak (Persistent Systems Pvt. Ltd.)

**The Bloomba Personal Content Database (invited)**

Raymie Stata (Stata Labs and University of California, Santa Cruz), Patrick Hunt (Stata Labs), Thiruvalluvan M. G. (iSoftTech Ltd.)

**Tutorial 5: The Continued Saga of DB-IR Integration****Location: Confederation 5 & 6****Presenter:**

Ricardo Baeza-Yates (University of Chile), Mariano Consens (University of Toronto)

**Demo Session 5 & 6****Location: Tudor 7 & 8**

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes during the conference. The box occupies most of the page below the 'NOTES' header.

## FRIDAY 11:00 – 12:30

**Research Session 23: Novel Models****Location: Ballroom****Session Chair:** *Hongjun Lu***An Annotation Management System for Relational Databases**

Deepavali Bhagwat, Laura Chiticariu, Wang-Chiew Tan, Gaurav Vijayvargiya (University of California, Santa Cruz)

**Symmetric Relations and Cardinality-Bounded Multisets in Database Systems**

Kenneth Ross, Julia Stoyanovich (Columbia University)

**Algebraic Manipulation of Scientific Datasets**

Bill Howe, David Maier (Oregon Health &amp; Science University)

**Research Session 24: Query Processing and Optimization****Location: Ontario****Session Chair:** *Leo Bertossi***Multi-objective Query Processing for Database Systems**

Wolf-Tilo Balke (University of California, Berkeley), Ulrich Gütntzer (University of Tübingen)

**Lifting the Burden of History from Adaptive Query Processing**

Amol Deshpande (University of California, Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research, Berkeley)

**A Combined Framework for Grouping and Order Optimization (short presentation)**

Thomas Neumann, Guido Moerkotte (University of Mannheim)

**The Case for Precision Sharing (short presentation)**

Sailesh Krishnamurthy, Michael J. Franklin (University of California, Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley), Garrett Jacobson (University of California, Berkeley)

**Industrial Session 8: Issues in Data Warehousing****Location: Territories****Session Chair:** *Glenn Pauley***Trends in data Warehousing: A Practitioner's View (invited)**

William O'Connell (IBM Toronto Lab)

**Technology Challenges in a Data Warehouse (invited)**

Ramesh Bhashyam (Teradata Corporation)

**Sybase IQ Multiplex – Designed For Analytics (invited)**

Roger MacNicol, Blaine French (Sybase)

**Tutorial 5: The Continued Saga of DB-IR Integration (cont.)****Location: Confederation 5 & 6****Presenter:**

Ricardo Baeza-Yates (University of Chile), Mariano Consens (University of Toronto)

**Demo Session 5 & 6****Location: Tudor 7 & 8**

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page area below the 'NOTES' header.



Anything they can crunch, we can crunch better. The Altix® Family of servers and superclusters combine the industry leading performance of the Intel® Itanium® 2 processor with a production-class Linux® solution. The result is a new class of reliable, cost-effective servers that meet the most demanding requirements. Last year, the Altix® 3000 series took the open standards community by storm delivering scalability across hundreds of processors and terabytes of memory. Now Altix® 350 brings best-in-class price and best-in-class performance to the mid-range. Whether you need a solution with infinite scalability, or dream of incredible power at an affordable price, there's an Altix system for you. [sgi.com/altix](http://sgi.com/altix)

SILICON GRAPHICS | The Source of Innovation and Discovery™

© 2004 Silicon Graphics, SGI, Altix and the SGI logo are registered trademarks and The Source of Innovation and Discovery is a trademark of Silicon Graphics, Inc. in the United States and/or other countries worldwide. Intel, Itanium, and Intel Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a registered trademark of Linus Torvalds, used with permission by Silicon Graphics, Inc.

## SPECIAL EVENTS

### Keynote Speech 1: Databases in a Wireless World

**David Yach, VP Software, Research in Motion**

**Ballroom, Tuesday 9:00-10:30**

The traditional view of distributed databases is based on a number of database servers with regular communication. Today information is stored not only in these central databases, but on a myriad of computers and computer-based devices in addition to the central storage. These range from desktop and laptop computers to PDA's and wireless devices such as cellular phones and BlackBerry's. The combination of large centralized databases with a large number and variety of associated edge databases effectively form a large distributed database, but one where many of the traditional rules and assumptions for distributed databases are no longer true. This keynote will discuss some of the new and challenging attributes of this new environment, particularly focusing on the challenges of wireless and occasionally connected devices. It will look at the new constraints, how these impact the traditional distributed database model, the techniques and heuristics being used to work within these constraints, and identify the potential areas where future research might help tackle these difficult issues.



**David Yach, VP Software, Research in Motion**

*David is Senior Vice President of Software at Research In Motion. David oversees and manages the development of software that has helped Research In Motion become a world leader in the mobile data communications market. With a Bachelor's degree in Math and an MBA, David has a solid portfolio of technical, operational and managerial experience from several senior-level positions. Prior to joining Research In Motion, David held the position of Vice President and Chief Architect at Sybase Inc.*

### Keynote Speech 2: Structures, Semantics and Statistics

**Alon Halevy, University of Washington**

**Ballroom, Wednesday 9:00-10:30**

Integration of data from multiple sources is one of the longest standing problems facing the database research community. In addition to being a problem in most enterprises and in large-scale science projects, research on this topic has been fueled by the promise of querying the WWW. I will begin by highlighting some of the significant recent achievements in the field of data integration, and will then focus on what I consider to be the main challenges going forward, namely, large-scale reconciliation of semantic heterogeneity, and on-the-fly information integration. To address these challenges, I argue for an approach based on computing statistics over corpora of database structures. At a fundamental level, the key challenge in data integration is to reconcile the semantics of disparate data sets, each expressed with different database structures. Computing statistics over a large corpus of schemas and mappings offers a powerful methodology for producing semantic mappings, the expressions that specify such reconciliation. In essence, the statistics offer hints about the semantics of the symbols in the structures, thereby enabling to detect when two symbols, from disparate schemas, should be matched to each other. The same methodology can be applied to several other data management tasks that involve search in a space of complex structures. I will illustrate several examples where this approach has been successful.



**Alon Halevy, University of Washington**

*Alon Halevy is an Associate Professor in the Computer Science and Engineering Department at the University of Washington. His interests are in data integration, management of XML data, peer-data management systems, query optimization, and the intersection of Database and AI technologies. His research developed several systems, such as the Information Manifold and Tukwila data integration systems, the Strudel web-site management system, and Piazza Peer-data Management System. In 1999, Dr. Halevy co-founded Nimble Technology, a data integration company, and in 2004 he founded Transformic Inc., a company focused on tools for bridging semantic heterogeneity. Dr. Halevy was a Sloan Fellow, and received the Presidential Early Career Award for Scientists and Engineers (PECASE) in 2000. He serves on the editorial board of the VLDB Journal and the advisory board of the Journal of Artificial Intelligence Research, and served as the program chair of SIGMOD 2003. He received his B.Sc from the Hebrew University in Jerusalem in 1988, and his Ph.D in Computer Science from Stanford University in 1993.*

### **10 Year Best Paper Award: Whither Data Mining?**

**Rakesh Agrawal, Ramakrishnan Srikant (IBM Almaden Research Center)**

**Ballroom, Thursday 9:00-10:30**

The last decade has witnessed tremendous advances in data mining. We take a retrospective look at these developments, focusing on association rules discovery, and discuss the challenges and opportunities ahead.



**Rakesh Agrawal, IBM Almaden Research Center**

*Rakesh Agrawal is an IBM Fellow, who leads the Intelligent Information Systems Research at the IBM Almaden Research Center. His current research interests include privacy and security technologies for data systems, web technologies, data mining and OLAP. He has pioneered fundamental concepts in data privacy, including Hippocratic Database, Sovereign Information Sharing, and Privacy-Preserving Data Mining. He earlier developed key data mining concepts and technologies. IBM's commercial data mining product, Intelligent Miner, grew out of this work. Rakesh Agrawal has published more than 100 research papers and he has been granted more than 50 patents. He is the recipient of the ACM-SIGKDD First Innovation Award, ACM-SIGMOD Edgar F. Codd Innovations Award, as well as the ACM-SIGMOD Test of Time Award. He is also a Fellow of IEEE and a Fellow of ACM. Scientific American named him to the list of 50 top scientists and technologists in 2003.*





**Ramakrishnan Srikant, IBM Almaden Research Center**

*Dr. Ramakrishnan Srikant is a Research Staff Member in the Intelligent Information Systems Research department at IBM Almaden Research Center, San Jose, CA. His research interests include privacy technologies for data systems, data mining, and web technologies. Dr. Srikant received the 2002 ACM Grace Murray Hopper Award "for his seminal work on mining association rules". Dr. Srikant has published more than 30 research papers that have been extensively cited. He has been granted 12 patents, and has another 9 pending patent applications. Dr. Srikant was a key architect for IBM's commercial data mining product, Intelligent Miner. He was named an IBM Research Division Master Inventor in 1999, and has also received 2 Outstanding Technical Achievement Awards for his contributions to Intelligent Miner. Dr. Srikant was the Program co-Chair of SIGKDD 2001. He has also served as the Tutorials Chair of KDD 2003, the Industrial Track co-Chair of PAKDD 2003, and co-Chair of the 1999 SIGMOD DMKD Workshop. Dr. Srikant received his M.S. and Ph.D. degrees in Computer Science from the University of Wisconsin, Madison in 1996. He also has a B. Tech. degree in Computer Science & Engineering from the Indian Institute of Technology, Madras, India.*



GREAT MINDS FOR A GREAT FUTURE



<http://www.cs.toronto.edu/db>

The University of Toronto's Database Group studies a broad array of topics involving the efficient, effective use of large collections of information and knowledge. Our projects cover both core database systems issues (query processing, indexing, and data mining), and new applications spanning data sharing, integration of information retrieval with data management, networked and web-based data management. The group has a vibrant graduate program that has produced a long, distinguished set of alumni active in database research and related areas. Our alumni have served on the faculties of major research universities including Harvard, the University of Washington, the University of Maryland, and Carnegie Mellon, and have led and participated in database research efforts in influential industry research labs that include Microsoft, AT&T Labs, GTE Labs, and IBM. We are actively recruiting applications for postdocs to work with our group.

John Mylopoulos  
Ken Sevcik  
Alberto Mendelzon  
Renée J. Miller  
Leonid Libkin  
Nick Koudas  
Mariano Consens

## TUTORIALS

### Tutorial 1. Database Architectures for New Hardware

Anastassia Ailamaki

Confederation 5 & 6, Tuesday 11:00-12:30 and 14:00-15:30

Thirty years ago, DBMS stored data on disks and cached recently used data in main memory buffer pools, while designers worried about improving I/O performance and maximizing main memory utilization. Today, however, databases live in multi-level memory hierarchies that include disks, main memories, and several levels of processor caches. Four (often correlated) factors have shifted the performance bottleneck of data-intensive commercial workloads from I/O to the processor and memory subsystem. First, storage systems are becoming faster and more intelligent (now disks come complete with their own processors and caches). Second, modern database storage managers aggressively improve locality through clustering, hide I/O latencies using prefetching, and parallelize disk accesses using data striping. Third, main memories have become much larger and often hold the application's working set. Finally, the increasing memory/processor speed gap has pronounced the importance of processor caches to database performance.

This tutorial will first survey techniques proposed in the computer architecture and database literature on understanding and evaluating database application performance on modern hardware. We will present approaches and methodologies used to produce time breakdowns when executing database workloads on modern processors. Then, we will survey techniques proposed to alleviate the problem, with major emphasis on data placement and prefetching techniques and their evaluation. Finally, we will discuss open problems and future directions: Is it only the memory subsystem database software architects should worry about? How important are other decisions processors make to database workload behavior? Given the emerging multi-threaded, multi-processor computers with modular, deep cache hierarchies, how feasible is it to create database systems that will adapt to their environment and will automatically take full advantage of the underlying hierarchy?



**Anastassia Ailamaki, Carnegie Mellon University**

*Anastassia Ailamaki received a B.Sc. degree in Computer Engineering from the Polytechnic School of the University of Patra, Greece, M.Sc. degrees from the Technical University of Crete, Greece and from the University of Rochester, NY, and a Ph.D. degree in Computer Science from the University of Wisconsin-Madison. In 2001, she joined the Computer Science Department at Carnegie Mellon University as an Assistant Professor. Her research interests are in the broad area of database systems and applications, with emphasis on database system behavior on modern processor hardware and disks. Her projects at Carnegie Mellon (including Staged Database Systems, Cache-Resident Data Bases, and the Fates Storage Manager), aim at building systems to strengthen the interaction between the database software and the underlying hardware and I/O devices. Her other research interests include automated database design for scientific databases, storage device modeling, and internet querying. She has received three best-paper awards (VLDB 2001, Performance 2002, and ICDE 2004), an NSF CAREER award (2002), and IBM Faculty Partnership awards in 2001, 2002, and 2003. She is a member of IEEE and ACM.*

## Tutorial 2: Security of Shared Data in Large Systems: State of the Art and Research Directions

**Arnon Rosenthal, Marianne Winslett**

**Confederation 5 & 6, Tuesday 16:00-17:30 and Wednesday 11:00-12:30**

Security is increasingly recognized as a key impediment to sharing data in enterprise systems, virtual enterprises, and the semantic web. Yet the topic has not been a focus for mainstream database research, industrial progress in data security has been slow, and (too) much security enforcement is in application code, or else is coarse grained and insensitive to data contents. Today the database research community is in an excellent position to make significant improvements in the way people think about security policies, due to the community's experience with declarative and logic-based specifications, automated compilation and physical design, and both semantic and efficiency issues for federated systems. These strengths provide a foundation for improving both theory and practice. This tutorial aims to enlighten the VLDB community about the state of the art in data security, especially for enterprise or larger systems, and to engage the community's interest in improving the current state of affairs by showing how database researchers can help improve the state of the art in data security.



**Arnie Rosenthal, MITRE**

*Arnie Rosenthal is a Principal Scientist at MITRE. He has broad interests in problems that arise when data is shared between communities, including a long-standing interest in the security issues that arise in data warehouses, federated databases, and enterprise information systems. He has also had a first-hand look at many security problems that arise in large government and military organizations.*



**Marianne Winslett, University of Illinois**

*Marianne Winslett has been a professor at the University of Illinois since 1987. She started working on database security issues in the early 1990s, focusing on semantic issues in MLS databases. Her interests soon shifted to issues of trust management for data on the web. Trust negotiation is her main current research focus.*

## Tutorial 3: Self-Managing Technology in Database Management Systems

**Surajit Chaudhuri, Benoit Dageville, Guy Lohman**

**Confederation 5 & 6, Wednesday 14:00-16:00 and 16:30-17:30**

The rising cost of labor relative to hardware and software means that the total cost of ownership of information technology is increasingly dominated by people costs. In response, all major providers of information technology are attempting to make their products more self-managing. In this tutorial, we will introduce the core concepts of self-managing technology and discuss their implications for database management systems. We will review the state of the art of self-managing technology for the IBM DB2, Microsoft SQL Server, and Oracle products, and describe the wealth of research topics remaining to be solved.



**Surajit Chaudhuri, Microsoft Research**

*Surajit Chaudhuri leads the Data Management and Exploration Group at Microsoft Research <http://research.microsoft.com/dmx>. In 1996, Surajit started the AutoAdmin project at Microsoft Research to address the challenges in building self-tuning database systems. This project developed novel automated index tuning technology that shipped with SQL Server 7.0 in 1998. The AutoAdmin project has since then continued to develop the self-tuning and self-manageability technology further in collaboration with the Microsoft SQL Server product team. Surajit's other project is Data Exploration, which looks at the problem of querying and discovery of information in a flexible manner information that spans text as well as relational data. Surajit is also interested in the problems of data cleaning and integration. Surajit did his Ph.D. from Stanford University in 1991 and worked at Hewlett-Packard Laboratories, Palo Alto from 1991-1995. He has published widely in major database conferences. <http://research.microsoft.com/users/surajit>*



**Benoit Dageville, Oracle Corp.**

*Benoit Dageville is a consulting member in the Oracle Database Server Technologies division at Oracle Corporation in Redwood Shores, California. His main areas of expertise include parallel query processing, ETL (Extract Transform and Load), large Data Warehouse benchmarks (e.g. TPC-H), SQL execution and optimization. Since 1999, he is one of the lead architects of the self-managing database initiative at Oracle. Major features resulting from this effort include the Automatic SQL Memory Management (Oracle9i), Automatic Workload Repository, Automatic Database Diagnostic Monitor and Automatic SQL Tuning (Oracle10g). Dr Benoit Dageville graduated from the University of Paris VI, France, in 1995 with a Ph.D. degree in computer science, specializing in Parallel Database Management Systems under the supervision of Dr. Patrick Valduriez. His research and industrial work resulted in several refereed papers in international conferences and journals.*



**Guy M. Lohman, IBM Almaden**

*Guy M. Lohman is Manager of Advanced Optimization in the Advanced Database Solutions Department at IBM Research Division's Almaden Research Center in San Jose, California, and has 22 years of experience in relational query optimization at IBM. He is the architect of the Optimizer of the DB2 Universal Data Base (UDB) for Linux, Unix, and Windows, and was responsible for its development from 1992 to 1997. Prior to that, he was responsible for the optimizers of the Starburst extensible object-relational and R\* distributed prototype DBMSs. More recently, Dr. Lohman co-invented and designed the DB2 Index Advisor (predecessor to today's Design Advisor), and in 2000 co-founded the DB2 Autonomic Computing project (formerly known as SMART -- Self-Managing And Resource Tuning), now part of IBM's company-wide Autonomic Computing initiative. In 2002, Dr. Lohman was elected to the IBM Academy of Technology. Dr. Lohman received his Ph.D. in Operations Research in 1976 from Cornell University. He is the author of over 40 papers in the refereed literature and the holder of 13 U.S. patents. His current research interests involve query optimization and self-managing database systems.*

## Tutorial 4: Architectures and Algorithms for Internet-Scale (P2P) Data Management

Joseph M. Hellerstein

Confederation 5 & 6, Thursday 14:00-15:30 and 16:00-17:30

The database community prides itself on scalable data management solutions. In recent years, a new set of scalability challenges have arisen in the context of peer-to-peer (p2p) systems on the Internet, in which the scaling metric is the number of participating computers, rather than the number of bytes stored. The best-known application of p2p technology to date has been file sharing, but there are compelling new application agendas in Internet monitoring, content distribution, distributed storage, multi-user games and next-generation Internet routing. The energy behind p2p technology has led to a renaissance in the distributed algorithms and distributed systems communities, much of which directly addresses issues in massively distributed data management. Moreover, many of these ideas have applicability beyond the "pure p2p" context of Internet end-user systems, with ramifications for any large-scale distributed system in which the scale of the system makes traditional administrative models untenable. Internet-scale systems present numerous unique technical challenges, including steady-state "churn" (nodes joining and leaving), the need to evolve and scale without reconfiguration, an absence of ongoing system administration, and adversarial participants in the processing. In this tutorial, we will focus on key data management building blocks including network indirection architectures, persistence models, network embeddings of computations, resource management, and security/trust challenges. We will also discuss motivations for the use of these technologies. We will ground the presentation in experiences from a set of deployed systems, and present open challenges that have arisen in this context.



**Joseph M. Hellerstein, University of California, Berkeley**

*Joseph M. Hellerstein is a Professor of Computer Science at the University of California, Berkeley, and is the Director of Intel Research, Berkeley. He is an Alfred P. Sloan Research Fellow, and a recipient of multiple awards, including ACM-SIGMOD's "Test of Time" award for his first published paper, NSF CAREER, NASA New Investigator, Okawa Foundation Fellowship, and IBM's Best Paper in Computer Science. In 1999, MIT's Technology Review named him one of the top 100 young technology innovators worldwide in their inaugural "TR100" list. Hellerstein's research focuses on data management and movement, including database systems, sensor networks, peer-to-peer and distributed systems. Prior to his position at Intel Research, Hellerstein was a co-founder of Cohera Corporation (now part of PeopleSoft), where he served as Chief Scientist from 1998-2001. Key ideas from his research have been incorporated into commercial and open-source database systems including IBM's DB2 and Informix, PeopleSoft's Catalog Management, and the open-source PostgreSQL system. Hellerstein currently serves on the technical advisory boards of a number of software companies, and has served as a member of the advisory boards of ACM SIGMOD and Ars Digita University. Hellerstein received his Ph.D. at the University of Wisconsin, Madison, a masters degree from University of California, Berkeley, and a bachelor's degree from Harvard College. He spent a pre-doctoral internship at IBM Almaden Research Center, and a post-doctoral internship at the Hebrew University in Jerusalem.*

## Tutorial 5: The Continued Saga of DB-IR Integration

Ricardo Baeza-Yates, Mariano Consens

Confederation 5 & 6, Friday 9:00-10:30 and 11:00-12:30

The world of data has been developed from two main points of view: the structured relational data model and the unstructured text model. The two distinct cultures of database and information retrieval now have a natural meeting place in the Web with its semi-structured XML model. As web-style searching becomes the ubiquitous tool, the need for integrating these two viewpoints becomes even more important. In this tutorial we explore the differences, the problems and the techniques for DB-IR integration for a range of applications. The tutorial will provide an overview of the different approaches put forward by the IR and DB communities and survey the DB-IR integration efforts. Both earlier proposals as well as recent ones (in the context of XML in particular) will be discussed. A variety of application scenarios for DB-IR integration will be covered. The objective of this tutorial is to provide an overview of the issues and approaches developed for the integration of database and information retrieval systems. The target audience of this tutorial includes researchers in database systems, as well as developers of Web and database/information retrieval applications.



**Ricardo Baeza-Yates, University of Chile**

*Ricardo Baeza-Yates is a professor at the Computer Science Department of the University of Chile, where he was the chair between 1993-95. He is also the director of the Center for Web Research, a project funded by the Millenium Scientific Initiative of the Ministry of Planning. He received the bachelor degree in Computer Science (CS) in 1983 from the University of Chile. Later, in 1985, he received also the M.Sc. in computer science and the professional title in electrical engineering (EE). One year later, he obtained M.Eng. in EE from the same university. He received his Ph.D. in CS from the University of Waterloo, Canada, in 1989, doing a six months post-doctoral position the same year. His research interests include information retrieval, algorithms, and information visualization. He is co-author of the book *Modern Information Retrieval*, published in 1999 by Addison-Wesley, as well as co-author of the 2nd edition of the *Handbook of Algorithms and Data Structures*, Addison-Wesley, 1991; and co-editor of *Information Retrieval: Algorithms and Data Structures*, Prentice-Hall, 1992, between other publications in journals published by ACM, IEEE, SIAM, etc. Ricardo received, in 1993, the Organization of American States award for young researchers in exact sciences. In 1994 he received the award to the best engineering research in the last 4 years from the Institute of Engineers of Chile, and was invited by the U.S.A. Presidential Office for a one month scientific tour in that country. In 1996, he won a scholarship from the Education and Science Ministry of Spain to have a sabbatical year at the Polytechnic University of Catalunya. In 1997 with two Brazilian colleagues obtained the COMPAQ prize to the best Brazilian research article in computer science. As a Ph.D. student he won several scholarships including the Ontario Graduate scholarship, the Institute for Computer Research scholarship for graduate studies, Information Technology and Research Centre graduate scholarship, the University of Waterloo graduate student award, and the Department of Computer Science Fellowship. Ricardo's professional activities include Presidency of the Chilean Computer Science Society (SCCC) between 1992-1995 and 1997-1998. From 1998-2000 he was in charge of the IEEE-CS chapter in Chile and has been involved in the South American ACM Programming Contest since 1998. He is currently the president of CLEI, a Latin American association of CS departments; and coordinates the Iberoamerican cooperation program in Electronics and Informatics. He was recently elected to the IEEE CS Board of Governors for the period 2002-04. In 2002 he was appointed to the Chilean Academy of Sciences, being the first person from computer science to achieve this position in Chile.*



**Mariano Consens, University of Toronto**

*Mariano Consens is a faculty member in Information Engineering at the MIE Department, University of Toronto, which he joined in 2003. Before that, he was research faculty at the School of Computer Science, University of Waterloo, from 1994 to 1999. He received his PhD and MSc degrees in Computer Science from the University of Toronto. He also holds a Computer Systems Engineer degree from the Universidad de la Republica, Uruguay. Mariano's research interests are in the areas of Data Management Systems and the Web, with a current focus on XML searching, autonomic systems and pervasive computing. He has over 20 publications and two patents, including journal publications selected from best conference papers. In addition to his academic positions, he has been active in the software industry as a founder and Director of Freedom Intelligence (a query engine provider), as the CTO of Classwave Wireless (a Bluetooth software infrastructure company) and as a technology advisor for Xign (an electronic payment systems supplier), OpenText (an early web search engine turned knowledge management software vendor) and others.*



**NOTES**

# ROADMAP

## to REAL-TIME NIRVANA

your guide to the  
UNWIRED ENTERPRISE



### Data Goes To The Edge

Sybase mobile middleware and mobile databases enable real-time data everywhere.

### Data Gets A Lot Smarter

Sybase Integration Orchestrator links and correlates isolated events into meaningful business metrics.

### Data Goes Into Action

Sybase Adaptive Server<sup>®</sup> Enterprise lets your business processes respond to exception events in real time.



Implementing a  
real-time initiative?  
Start with our  
new White Paper.

**SYBASE**<sup>®</sup>  
The Enterprise. Unwired.

Free White Paper  
[www.sybase.com/rts](http://www.sybase.com/rts)

©2004 Sybase, Inc. Sybase and the Sybase logo are registered trademarks of Sybase, Inc. EXPLANATIONS™ by Xplane.com™. All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

## PANELS

### Panel 1: Biological Data Management: Research, Practice and Opportunities Territories, Tuesday 16:00-17:30

#### Moderator:

Thodoros Topaloglou

#### Panelists:

Susan B. Davidson, H. V. Jagadish, Victor M. Markowitz, Evan W. Steeg, Mike Tyers.

Biological research and drug development are routinely producing terabytes of data that need to be organized, queried and reduced to useful scientific knowledge. Although data management technology can provide solutions to problems, in practice the data needs of biomedical research are not well served. The goal of this panel is to expose the barriers blocking the effective application of advanced data management technology to biological data.

The current state of biological data management ranges from “malpractice” of database principles, to the reinvention of well-known data management methods, to the contribution of valuable new ideas to the state of the art in database research. On the other side, database research is often incompatible with the production requirements of the biomedical data operations: integration and interoperability of biological data sources, support for more meaningful data types, domain-aware querying interfaces, practical workflow management, and methods for evaluating data quality including data provenance are still considered open problems in biological data management, just as a decade ago. Our panelists bring wealth of experience in translating database research into biological data management tools and in communicating data requirements back to the database community. They will identify and debate the key challenges and opportunities to which the database community should contribute.



**Thodoros Topaloglou, MDS Proteomics**

*Thodoros Topaloglou has worked with a wide range of biological data types including sequence, micro-array, proteomics and clinical data for which he developed querying and integration tools, and production infrastructures. He is presently a Senior Vice President of Scientific Computing at MDS Proteomics. In the past decade Thodoros held senior technical positions in several biotechnology and research organizations. From 2001 to 2002, he was Director of Data Management in Target Validation and Drug Discovery at Incyte Genomics. He was the Director of Gene Expression Data Management and played a lead role in the development of Gene Logic's Gene Express product between 1997 and 2001. He also held staff scientist positions at Lawrence Berkeley Laboratory and Affymetrix. Thodoros holds a PhD in Computer Science from the University of Toronto.*



**Susan B. Davidson, University of Pennsylvania**

*Susan B. Davidson received the B.A. degree in Mathematics from Cornell University, Ithaca, NY, in 1978, and the M.A. and Ph.D. degrees in Electrical Engineering and Computer Science from Princeton University, Princeton NJ, in 1980 and 1982. Dr. Davidson is currently a Professor in the Department of Computer and Information Science at the University of Pennsylvania (UPenn), where she has been since 1982. She is an ACM Fellow, a Fulbright scholar, and recently stepped down as founding co-Director of the Center for Bioinformatics at UPenn (PCBI). Preceding the formation of the PCBI, Dr. Davidson was involved with planning and administering a 1994 NSF funded research training program in computational biology. She also helped establish degree programs in*

*bioinformatics and computational biology run through the departments of Biology and Computer and Information Science, at UPenn. Dr. Davidson's research interests include database systems, database modeling, database integration, distributed systems, bioinformatics and real-time systems. Within bioinformatics she is best known for her work with the Kleisli data integration system, and more recently with XML as a data exchange and integration strategy.*



**H. V. Jagadish, University of Michigan, Ann Arbor**

*H. V. Jagadish is Professor of Electrical Engineering and Computer Science and Professor of Bioinformatics at University of Michigan, Ann Arbor since 1999. Prior to joining the University of Michigan, he worked at AT&T where he headed the database department. He has recently co-organized a workshop on data management issues for molecular and cell biology and guest edited a special issue, devoted to data management for biology, of the OMICS journal on integrative biology. He obtained his Ph. D. from Stanford in 1985.*



**Victor M. Markowitz, JGI and Lawrence Berkeley National Laboratory**

*Victor M. Markowitz, D.Sc., is Chief Informatics Officer at JGI and Head of the recently established Biological Data Management and Technology Center at Lawrence Berkeley National Laboratory. Until 2003 he was Chief Information Officer and Senior Vice President, Data Management Systems, at Gene Logic Inc., where he was responsible for the development and deployment of Gene Logic's data management products and software tools. Prior to joining Gene Logic, Dr. Markowitz was at Lawrence Berkeley National Laboratory where he lead the development of the Object Protocol Model data management and integration tools that have been used for developing public and commercial genome databases. Dr. Markowitz received his M.Sc. and D.Sc. degrees in computer science from Technion, the Israel Institute of Technology. Dr. Markowitz has authored articles and book chapters on various aspects of data management. He has served on review panels and program committees for database and bioinformatics programs and conferences.*



**Evan W. Steeg, Independen Consultant and Author**

*Evan W. Steeg has developed algorithms and software for pattern discovery and computational molecular biology for twenty years in academia, government and industry. From DNA sequence analysis to protein structure prediction and x-ray crystallography, to microarray experiments and clinical outcome prediction, Dr. Steeg has faced many of the challenges of pulling valuable knowledge from biological databases. Formerly a scientific co-founder and President & CEO of Molecular Mining Corporation, he now consults and is writing a book on life sciences data mining for John Wiley & Sons. He received a B.A. in Mathematics from Cornell University, and M.Sc. and Ph.D. degrees in Computer Science from the University of Toronto.*



**Mike Tyers, Samuel Lunenfeld Research Institute and University of Toronto**

*Mike Tyers is a Senior Investigator at the Samuel Lunenfeld Research Institute of Mount Sinai Hospital and Professor in the Department of Medical Genetics and Microbiology at the University of Toronto. He received his Ph.D in Biochemistry from McMaster University in 1989 and was a postdoctoral fellow at Cold Spring Harbor Laboratory from 1989-1992. Dr. Tyers currently holds a Canada Research in Proteomics, Functional Genomics and Bioinformatics. His laboratory uses biochemical, genetic and genome-wide approaches to elucidate mechanisms that control the cell division cycle. Current research interests include ubiquitin-dependent degradation pathways that control the G1/S transition, genome stability and gene expression, phosphorylation-dependent protein interactions, mechanisms that couple growth and division, high throughput analysis, archival and visualization of protein and genetic interactions, and mathematical modeling of signal transduction events.*

## **Panel 2: Where is Business Intelligence taking today's database systems?**

**Confederation 5 & 6, Thursday 11:00-12:30**

### **Moderator:**

William O'Connell

### **Panelists:**

Andy Witkowski, Ramesh Bhashyam, Surajit Chaudhuri, Nigel Campbell

The invention of technology made Business Intelligence (BI) possible over relational engines, but now the experiences of putting them into production has unearthed a new set of problems in need of further invention. Over a period of few past years, academia has provided very performant and storage efficient technologies for fundamental BI objects. Database industry either incorporated into their sql engines some of these algorithms (like data mining algorithms, OLAP engines), or tried to integrate better stand alone BI engines like OLAP, or provided their own unique solutions for BI. This panel will discuss a few of these emerging issues and trends. The intent is not to overview individual products and or solutions, nor to provide a background on BI solutions. But, it is to point out select trends and issues, as well as old issues that are still very real. The panelists will also describe why these issues are important for the research community. The initial questions posed to the panelists will be twofold. First, is where should the BI community go on extending the SQL Language bindings (such as for data mining, reporting, and data analysis), as well as associated DBMS implications to support new and existing BI extensions. And, Second, what does XML mean to BI, as well as associated DBMS implications.



**William O'Connell, IBM Toronto Lab**

*William O'Connell is the lead BI architect for DB2 database engineering. Bill is Senior Technical Staff Member within the IBM Toronto Lab and is responsible for DB2 BI and Warehousing technology for the database server. He is also driving DB2 enablement within the end-to-end BI Data Warehousing solution stack. Bill has been working in database development for 18 years, and received his Ph.D at Illinois Institute of Technology. Bill has also worked at NCR-Teradata, as well as AT&T Bell Labs in database development.*



**Andy Witkowski, Oracle**

*Andy Witkowski is an Architect for the Oracle Data Warehousing and Business Intelligence Groups. He has actively worked on extending RDBMS with language, algorithms and data structures for BI for the past six years. His interest include materialized views, Business Intelligence calculations in SQL like SQL Spreadsheet, Analytic and Statistical functions, compressed cubes, and query optimization. Andy has been working in database technology for 15 years, and received his Ph.D. at University of Southern California.*



**Ramesh Bhashyam, NCR-Teradata**

*Ramesh Bhashyam is the CTO for Teradata Database Engineering at NCR-Teradata. He is responsible for all the database server aspects of Teradata DBMS engineering. Ramesh has a Masters in Computer Science and a Bachelor's in Electrical Engineering.*



**Surajit Chaudhuri, Microsoft Research**

*Surajit Chaudhuri leads the Data Management and Exploration Group at Microsoft Research <http://research.microsoft.com/dmx>. In 1996, Surajit started the AutoAdmin project on self-tuning database systems at Microsoft Research and developed novel automated physical design tuning technology for SQL Server 7.0, SQL Server 2000 and SQL Server 2005. Technology on data mining from the group was incorporated in SQL Server 2000. More recently, Surajit has initiated work in the area of data cleaning and integration. Surajit's most recent project is Data Exploration which looks at the problem of flexible querying of information that spans text as well as relational data. Surajit did his Ph.D. from Stanford University in 1991 and worked at Hewlett-Packard Laboratories, Palo Alto from 1991-1995. He has published widely in major database conferences.*

## DEMONSTRATIONS

### Group 1: Data Mining

**Location: Tudor 7 & 8**

**Wednesday 9:00-10:15 and 11:00-12:30**

**Thursday 2:00-3:30 and 4:00-5:30**

**GPX: Interactive Mining of Gene Expression Data**

Daxin Jiang, Jian Pei, Aidong Zhang (State University of New York at Buffalo)

**Computing Frequent Itemsets Inside Oracle 10G**

Wei Li, and Ari Mozes (Oracle Corporation)

**StreamMiner: A Classifier Ensemble-based Engine to Mine Concept-drifting Data Streams**

Wei Fan (IBM T.J. Watson Research)

**Semantic Mining and Analysis of Gene Expression Data**

Xin Xu, Gao Cong, Beng Chin Ooi, Kian-Lee Tan, Anthony K.H.Tung (National University of Singapore)

**HOS-Miner: A System for Detecting Outlying Subspaces of High-dimensional Data**

Ji Zhang, Meng Lou (University of Toronto), Tok Wang Ling (National University of Singapore), Hai Wang (University of Toronto)

**VizTree: a Tool for Visually Mining and Monitoring Massive Time Series Databases**

Jessica Lin, Eamonn Keogh, Stefano Lonardi (University of California, Riverside), Jeffrey P. Lankford, Daonna M. Nystrom (The Aerospace Corporation)

### Group 2: Distributed Databases

**Location: Tudor 7 & 8,**

**Tuesday 2:00-3:30 and 4:00-5:30**

**Wednesday 9:00-10:15 and 11:00-12:30**

**An Electronic Patient Record “on Steroids”: Distributed, Peer-to-Peer, Secure and Privacy-conscious**

Serge Abiteboul (INRIA), Bogdan Alexe (Bell Labs, Lucent Technologies), Omar Benjelloun, Bogdan Cautis (INRIA, France), Irini Fundulaki (Bell Labs, Lucent Technologies), Tova Milo (INRIA and Tel-Aviv University), Arnaud Sahuguet (Tel-Aviv University)

**Queries and Updates in the coDB Peer to Peer Database System**

Enrico Franconi (Free University of Bozen-Bolzano), Gabriel Kuper (University of Trento), Andrei Lopatenko (Free University of Bozen-Bolzano and University of Manchester), Ilya Zaihrayeu (University of Trento)

**A-ToPSS: A Publish/Subscribe System Supporting Imperfect Information Processing**

Haifeng Liu, Hans-Arno Jacobsen (University of Toronto)

**Efficient Constraint Processing for Highly Personalized Location Based Services**

Zhengdao Xu, Hans-Arno Jacobsen (University of Toronto)

**LH\*RS: A Highly Available Distributed Data Storage**

Witold Litwin, Rim Moussa (Université Paris Dauphine), Thomas Schwarz (Santa Clara University)

### Group 3: XML

**Location: Tudor 7 & 8**

**Tuesday 2:00-3:30 and 4:00-5:30**

**Thursday 9:00-10:15 and 11:00-12:30**

**Semantic Query Optimization in an Automata-Algebra Combined XQuery Engine over XML Streams**

Hong Su, Elke A. Rundensteiner, Murali Mani (Worcester Polytechnic Institute)

**ShreX: Managing XML Documents in Relational Databases**

Fang Du (OGI/OHSU), Sihem Amer-Yahia (AT&T Labs), Juliana Freire (OGI/OHSU)

**A Uniform System for Publishing and Maintaining XML Data**

Byron Choi (University of Pennsylvania), Wenfei Fan (University of Edinburgh and Bell Laboratories), Xibei Jia (University of Edinburgh), Arek Kasprzyk (European Bioinformatics Institute)

**An Injection of Tree Awareness: Adding Staircase Join to PostgreSQL**

Sabine Mayer, Torsten Grust (University of Konstanz), Maurice van Keulen (University of Twente, The Netherlands), Jens Teubner (University of Konstanz)

**FluXQuery: An Optimizing XQuery Processor for Streaming XML Data**

Christoph Koch, Stefanie Scherzinger (Technische Univ. Wien), Nicole Schweikardt (Humboldt University Berlin), Bernhard Stegmaier (Technische Univ. München)

**Group 4: Web and Web-Services****Location: Tudor 7 & 8****Wednesday 2:00-4:00 and 4:30-5:30****Thursday 9:00-10:15 and 11:00-12:30****COMPASS: A Concept-based Web Search Engine for HTML, XML, and Deep Web Data**

Jens Graupmann, Michael Biber, Christian Zimmer, Patrick Zimmer, Matthias Bender, Martin Theobald, Gerhard Weikum (Max-Planck Institute for Computer Science, Saarbruecken)

**Discovering and Ranking Semantic Associations over a Large RDF Metabase**

Chris Halaschek, Boanerges Aleman-Meza, I. Budak Arpinar, and Amit P. Sheth (University of Georgia)

**An Automatic Data Grabber for Large Web Sites**

Valter Crescenzi (University of Roma), Giansalvatore Mecca (University of Potenza), Paolo Merialdo, Paolo Missier (University of Roma)

**WS-CatalogNet: An Infrastructure for Creating, Peering, and Querying e-Catalog Communities**

Karim Baina, Boualem Benatallah, Hye-young Paik (University of New South Wales), Farouk Toumani, Christophe Rey (ISIMA), Agnieszka Rutkowska, Bryan Harianto (University of New South Wales)

**Trust-Serv: A Lightweight Trust Negotiation Service**

Halvard Skogsrud, Boualem Benatallah (University of New South Wales), Fabio Casati (Hewlett-Packard Labs), Manh Q. Dinh (University of New South Wales)

**Group 5: Query Optimization and Database Integrity****Location: Tudor 7 & 8****Wednesday 2:00-4:00 and 4:30-5:30****Friday 10:15-11:45 and 12:00-1:30****Green Query Optimization: Taming Query Optimization Overheads through Plan Recycling**

Parag Sarda, Jayant R. Haritsa (Indian Institute of Science)

**Progressive Optimization in Action**

Vijayshankar Raman, Volker Markl, David Simmen, Guy Lohman, Hamid Pirahesh (IBM Almaden Research Center)

**CORDS: Automatic Generation of Correlation Statistics in DB2**

Ihab F. Ilyas (University of Waterloo), Volker Markl, Peter J. Haas, Paul G. Brown, Ashraf Aboulnaga (IBM Almaden Research Center)

**CHICAGO: A Test and Evaluation Environment for Coarse-Grained Optimization**

Tobias Kraft, Holger Schwarz (University of Stuttgart)

**SVT: Schema Validation Tool for Microsoft SQL-Server**

Ernest Teniente, Carles Farré, Toni Urpí, Carlos Beltrán, David Gañán (University of Catalunya)

**Group 6: Streaming Data and Spatial and Multimedia Databases****Tudor 7 & 8,****Thursday 2:00-3:30 and 4:00-5:30****Friday 10:15-11:45 and 12:00-1:30****CAPE: Continuous Query Engine with Heterogeneous-Grained Adaptivity**

Elke A. Rundensteiner, Luping Ding, Timothy Sutherland, Yali Zhu, Brad Pielech, Nishant Mehta (Worcester Institute of Technology)

**HiFi: A Unified Architecture for High Fan-in Systems**

Owen Cooper, Anil Edakkunni, Michael J. Franklin (University of California, Berkeley), Wei Hong (Intel Research Berkeley), Shawn R. Jeffery, Sailesh Krishnamurthy, Fredrick Reiss, Shariq Rizvi, Eugene Wu (University of California, Berkeley)

**An Integration Framework for Sensor Networks and Data Stream Management Systems**

Daniel J. Abadi, Wolfgang Lindner, Samuel Madden (MIT), Jörg Schuler (Tufts University)

**QStream: Deterministic Querying of Data Streams**

Sven Schmidt, Henrike Berthold, Wolfgang Lehner (Dresden University of Technology)

**AIDA: an Adaptive Immersive Data Analyzer**

Mehdi Sharifzadeh, Cyrus Shahabi, Bahareh Navai, Farid Parvini, Albert A. Rizzo (University of Southern California)

**BiVideo Video Database Management System**

Özgür Ulusoy, Uğur Güdükbay, Mehmet Emin Dönderler, Ediz Şaykol, Cemil Alper (Bilkent University)

**PLACE: A Query Processor for Handling Real-time Spatio-temporal Data Streams**

Mohamed F. Mokbel, Xiaopeng Xiong, Walid G. Aref, Susanne E. Hambrusch, Sunil Prabhakar, Moustafa A. Hammad (Purdue University)



## ABSTRACTS

### Research Session 1: Compression And Indexing

#### Compressing Large Boolean Matrices using Reordering Techniques

David S. Johnson, Shankar Krishnan (AT&T Labs-Research), Jatin Chhugani, Subodh Kumar (John Hopkins University), Suresh Venkatasubramanian (AT&T Labs-Research)

Large boolean matrices are a basic representational unit in a variety of applications, with some notable examples being interactive visualization systems, mining large graph structures, and association rule mining. Designing space and time efficient scalable storage and query mechanisms for such large matrices is a challenging problem. We present a lossless compression strategy to store and access such large matrices efficiently on disk. Our approach is based on viewing the columns of the matrix as points in a very high dimensional Hamming space, and then formulating an appropriate optimization problem that reduces to solving an instance of the Traveling Salesman Problem on this space. Finding good solutions to large TSP's in high dimensional Hamming spaces is itself a challenging and little-explored problem -- we cannot readily exploit geometry to avoid the need to examine all  $N^2$  inter-city distances and instances can be too large for standard TSP codes to run in main memory. Our multi-faceted approach adapts classical TSP heuristics by means of instance-partitioning and sampling, and may be of independent interest. For instances derived from interactive visualization and telephone call data we obtain significant improvement in access time over standard techniques, and for the visualization application we also make significant improvements in compression.

#### On the Performance of Bitmap Indices for High Cardinality Attributes

Kesheng Wu, Ekow Otoo, Arie Shoshani (Lawrence Berkeley National Laboratory)

It is well established that bitmap indices are efficient for read-only attributes with low attribute cardinalities. For an attribute with a high cardinality, the size of the bitmap index can be very large. To overcome this size problem, specialized compression schemes are used. Even though there are empirical evidences that some of these compression schemes work well, there has not been any systematic analysis of their effectiveness. In this paper, we systematically analyze the two most efficient bitmap compression techniques, the Byte-aligned Bitmap Code(BBC) and the Word-Aligned Hybrid (WAH) code. Our analyses show that both compression schemes can be optimal. We propose a novel strategy to select the appropriate algorithms so that this optimality is achieved in practice. In addition, our analyses and tests show that the compressed indices are relatively small compared with commonly used indices such as B-trees. Given these facts, we conclude that bitmap index is efficient on attributes of low cardinalities as well as on those of high cardinalities.

#### Practical Suffix Tree Construction

Sandeep Tata, Richard A. Hankins, Jignesh M. Patel (University of Michigan)

Large string datasets are common in a number of emerging text and biological database applications. Common queries over such datasets include both exact and approximate string matches. These queries can be evaluated very efficiently by using a suffix tree index on the string dataset. Although suffix tree scan be constructed quickly in memory for small input datasets, constructing persistent trees for large datasets has been challenging. In this paper, we explore suffix tree construction algorithms over a wide spectrum of data sources and sizes. First, we show that on modern processors, a cache-efficient algorithm with  $O(n^2)$  complexity outperforms the popular  $O(n)$  Ukkonen algorithm, even for in-memory construction. For larger datasets, the disk I/O requirement quickly becomes the bottleneck in each algorithm's performance. To address this problem, we present a buffer management strategy for the  $O(n^2)$  algorithm, creating a new disk-based construction algorithm that scales to sizes much larger than have been previously described in the literature. Our approach far outperforms the best known disk-based construction algorithms.

## Research Session 2: XML Views And Schemas

### Answering XPath Queries over Networks by Sending Minimal Views

Keishi Tajima, Yoshiki Fukui (JAIST)

When a client submits a set of XPath queries to a XML database on a network, the set of answer sets sent back by the database may include redundancy in two ways: some elements may appear in more than one answer set, and some elements in some answer sets may be subelements of other elements in other (or the same) answer sets. Even when a client submits a single query, the answer can be self-redundant because some elements may be subelements of other elements in that answer. Therefore, sending those answers as they are is not optimal with respect to communication costs. In this paper, we propose a method of minimizing communication costs in XPath processing over networks. Given a single or a set of queries, we compute a minimal-size view set that can answer all the original queries. The database sends this view set to the client, and the client produces answers from it. We show algorithms for computing such a minimal view set for given queries. This view set is optimal; it only includes elements that appear in some of the final answers, and each element appears only once.

### A Framework for Using Materialized XPath Views in XML Query Processing

Andrey Balmin, Fatma Özcan, Kevin Beyer, Roberta Cochrane, Hamid Pirahesh (IBM Almaden)

XML languages, such as XQuery, XSLT and SQL/XML, employ XPath as the search and extraction language. XPath expressions often define complicated navigation, resulting in expensive query processing, especially when executed over large collections of documents. In this paper, we propose a framework for exploiting materialized XPath views to expedite processing of XML queries. We explore a class of materialized XPath views, which may contain XML fragments, typed data values, full paths, node references or any combination thereof. We develop an XPath matching algorithm to determine when such views can be used to answer a user query containing XPath expressions. We use the match information to identify the portion of an XPath expression in the user query which is not covered by the XPath view. Finally, we construct, possibly multiple, compensation expressions which need to be applied to the view to produce the query result. Experimental evaluation, using our prototype implementation, shows that the matching algorithm is very efficient and usually accounts for a small fraction of the total query compilation time.

### Schema-Free XQuery

Yunyao Li, Cong Yu, H. V. Jagadish (University of Michigan)

The widespread adoption of XML holds out the promise that document structure can be exploited to specify precise database queries. However, the user may have only a limited knowledge of the XML structure, and hence may be unable to produce a correct XQuery, especially in the context of a heterogeneous information collection. The default is to use keyword-based search and we are all too familiar with how difficult it is to obtain precise answers by these means. We seek to address these problems by introducing the notion of Meaningful Lowest Common Ancestor Structure (MLCAS) for finding related nodes within an XML document. By automatically computing MLCAS and expanding ambiguous tag names, we add new functionality to XQuery and enable users to take full advantage of XQuery in querying XML data precisely and efficiently without requiring (perfect) knowledge of the document structure. Such a Schema-Free XQuery is potentially of value not just to casual users with partial knowledge of schema, but also to experts working in a data integration or data evolution context. In such a context, a schema-free query, once written, can be applied universally to multiple data sources that supply similar content under different schemas, and applied "forever" as these schemas evolve. Our experimental evaluation found that it was possible to express a wide variety of queries in a schema-free manner and have them return correct results over a broad diversity of schemas. Furthermore, the evaluation of a schema-free query is not expensive using a novel stack-based algorithm we develop for computing MLCAS: from 1 to 4 times the execution time of an equivalent schema-aware query.

## Research Session 3: Controlling Access

### Client-Based Access Control Management for XML Documents

Luc Bouganim (INRIA Rocquencourt), François Dang Ngoc, Philippe Pucheral (PRISM Laboratory)

The erosion of trust put in traditional database servers and in Database Service Providers, the growing interest for different forms of data dissemination and the concern for protecting children from suspicious Internet content are different factors that lead to move the access control from servers to clients. Several encryption schemes can be used to serve this purpose but all suffer from a static way of sharing data. With the emergence of hardware and software security elements on client devices, more dynamic client-based access control schemes can be devised. This paper proposes an efficient client-based evaluator of access control rules for regulating access to XML documents. This evaluator takes benefit from a dedicated index to quickly converge towards the authorized parts of a – potentially streaming – document. Additional security mechanisms guarantee that prohibited data can never be disclosed during the processing and that the input document is protected from any form of tampering. Experiments on synthetic and real datasets demonstrate the effectiveness of the approach.

### Secure XML Publishing without Information Leakage in the Presence of Data Inference

Xiaochun Yang (Northeastern University), Chen Li (University of California, Irvine)

Recent applications are seeing an increasing need that publishing XML documents should meet precise security requirements. In this paper, we consider data-publishing applications where the publisher specifies what information is sensitive and should be protected. We show that if a partial document is published carelessly, users can use common knowledge (e.g., “all patients in the same ward have the same disease”) to infer more data, which can cause leakage of sensitive information. The goal is to protect such information in the presence of data inference with common knowledge. We consider common knowledge represented as semantic XML constraints. We formulate the process how users can infer data using three types of common XML constraints. Interestingly, no matter what sequences users follow to infer data, there is a unique, maximal document that contains all possible inferred documents. We develop algorithms for finding a partial document of a given XML document without causing information leakage, while allowing publishing as much data as possible. Our experiments on real data sets show that effect of inference on data security, and how the proposed techniques can prevent such leakage from happening.

### Limiting Disclosure in Hippocratic Databases

Kristen LeFevre (IBM Almaden and University of Wisconsin-Madison), Rakesh Agrawal (IBM Almaden), Vuk Ercegovic, Raghu Ramakrishnan (University of Wisconsin-Madison), Yirong Xu (IBM Almaden), David DeWitt (University of Wisconsin-Madison)

We present a practical and efficient approach to incorporating privacy policy enforcement into an existing application and database environment, and we explore some of the semantics tradeoffs introduced by enforcing these rules at cell-level granularity. Through a comprehensive set of performance experiments, we show that the cost of privacy enforcement is small, and scalable to large databases.

## Research Session 4: XML (I)

### On Testing Satisfiability of Tree Pattern Queries

Laks Lakshmanan, Ganesh Ramesh, Hui Wang, Zheng Zhao (University of British Columbia)

XPath and XQuery (which includes XPath as a sublanguage) are the major query languages for XML. An important issue arising in efficient evaluation of queries expressed in these languages is satisfiability, i.e., whether there exists a database, consistent with the schema if one is available, on which the query has a non-empty answer. Our experience shows satisfiability check can effect substantial savings in query evaluation. We systematically study satisfiability of tree pattern queries (which capture a useful fragment of XPath) together with additional constraints, with or without a

schema. We identify cases in which this problem can be solved in polynomial time and develop novel efficient algorithms for this purpose. We also show that in several cases, the problem is NP-complete. We ran a comprehensive set of experiments to verify the utility of satisfiability check as a preprocessing step in query processing. Our results show that this check takes a negligible fraction of the time needed for processing the query while often yielding substantial savings.

#### **Containment of Nested XML Queries**

Xin Dong, Alon Halevy, Igor Tatarinov (University of Washington)

Query containment is the most fundamental relationship between a pair of database queries: a query  $Q$  is said to be contained in a query  $Q'$  if the answer for  $Q$  is always a subset of the answer for  $Q'$ , independent of the current state of the database. Query containment is an important problem in a wide variety of data management applications, including verification of integrity constraints, reasoning about contents of data sources in data integration, semantic caching, verification of knowledge bases, determining queries independent of updates, and most recently, in query reformulation for peer data management systems. Query containment has been studied extensively in the relational context and for XPath queries, but not for XML queries with nesting. We consider the theoretical aspects of the problem of query containment for XML queries with nesting. We begin by considering conjunctive XML queries (c-XQueries), and show that containment is in polynomial time if we restrict the fan-out (number of sibling sub-blocks) to be 1. We prove that for arbitrary fan-out, containment is coNP-hard already for queries with nesting depth 2, even if the query does not include variables in the return clauses. We then show that for queries with fixed nesting depth, containment is coNP-complete. Next, we establish the computational complexity of query containment for several practical extensions of c-XQueries, including queries with union and arithmetic comparisons, and queries where the XPath expressions may include descendant edges and negation. Finally, we describe a few heuristics for speeding up query containment checking in practice by exploiting properties of the queries and the underlying schema.

#### **Efficient XML-to-SQL Query Translation: Where to Add the Intelligence ?**

Rajasekar Krishnamurthy (IBM Almaden), Raghav Kaushik (Microsoft Research), Jeffrey Naughton (University of Wisconsin-Madison)

We consider the efficiency of queries generated by XML to SQL translation. We first show that published XML-to-SQL query translation algorithms are suboptimal in that they often translate simple path expressions into complex SQL queries even when much simpler equivalent SQL queries exist. There are two logical ways to deal with this problem. One could generate suboptimal SQL queries using a fairly naive translation algorithm, and then attempt to optimize the resulting SQL; or one could use a more intelligent translation algorithm with the hopes of generating efficient SQL directly. We show that optimizing the SQL after it is generated is problematic, becoming intractable even in simple scenarios; by contrast, designing a translation algorithm that exploits information readily available at translation time is a promising alternative. To support this claim, we present a translation algorithm that exploits translation time information to generate efficient SQL for path expression queries over tree schemas.

#### **Taming XPath Queries by Minimizing Wildcard Steps**

Chee-Yong Chan (National University of Singapore), Wenfei Fan (University of Edinburgh & Bell Laboratories), Yiming Zeng (National University of Singapore)

This paper presents a novel and complementary technique to optimize an XPath query by minimizing its wildcard steps. Our approach is based on using a general composite axis called the layer axis, to rewrite a sequence of XPath steps (all of which are wildcard steps except for possibly the last) into a single layer-axis step. We describe an efficient implementation of the layer axis and present a novel and efficient rewriting algorithm to minimize both non-branching as well as branching wildcard steps in XPath queries. We also demonstrate the usefulness of wildcard-step elimination by proposing an optimized evaluation strategy for wildcard-free Path queries that enables selective loading of only the relevant input XML data for query evaluation. Our experimental results not only validate the scalability and efficiency of our optimized evaluation strategy, but also demonstrate the effectiveness of our rewriting algorithm for minimizing wildcard

steps in XPath queries. To the best of our knowledge, this is the first effort that addresses this new optimization problem.

### **The NEXT Framework for Logical Xquery Optimization**

Alin Deutsch, Yannis Papakonstantinou, Yu Xu (University of California, San Diego)

Classical logical optimization techniques rely on a logical semantics of the query language. The adaptation of these techniques to XQuery is precluded by its definition as a functional language with operational semantics. We introduce Nested XML Tableaux which enable a logical foundation for XQuery semantics and provide the logical plan optimization framework of our XQuery processor. As a proof of concept, we develop and evaluate a minimization algorithm for removing redundant navigation within and across nested subqueries. The rich XQuery features create key challenges that fundamentally extend the prior work on the problems of minimizing conjunctive and tree pattern queries.

## **Research Session 5: Stream Mining**

### **Detecting Change in Data Streams**

Daniel Kifer, Shai Ben-David, Johannes Gehrke (Cornell University)

Detecting changes in a data stream is an important area of research with many applications. In this paper, we present a novel method for the detection and estimation of change. In addition to providing statistical guarantees on the reliability of detected changes, our method also provides meaningful descriptions and quantification of these changes. Our approach assumes that the points in the stream are independently generated, but otherwise makes no assumptions on the nature of the generating distribution. Thus our techniques work for both continuous and discrete data. In an experimental study we demonstrate the power of our techniques.

### **Stochastic Consistency, and Scalable Pull-Based Caching for Erratic Data Stream Sources**

Shanzhong Zhu, China V. Ravishankar (University of California, Riverside)

We introduce the notion of stochastic consistency, and propose a novel approach to achieving it for caches of highly erratic data. Erratic data sources, such as stock prices, sensor data, and network traffic statistics, are common and important in practice. However their erratic patterns of change can make caching hard. Stochastic consistency guarantees that errors in cached values of erratic data remain within a user-specified bound, with a user-specified probability. We use a Brownian motion model to capture the behavior of data changes, and use its underlying theory to predict when caches should initiate pulls to refresh cached copies to maintain stochastic consistency. Our approach allows servers to remain totally stateless, thus achieving excellent scalability and reliability. We also discuss a new real-time scheduling approach for servicing pull requests at the server. Our scheduler delivers prompt response whenever possible, and minimizes the aggregate cache-source deviation due to delays during server overload. We conduct extensive experiments to validate our model on real-life datasets, and show that our scheme outperforms current schemes.

### **False Positive or False Negative: Mining Frequent Item sets from High Speed Transactional Data Streams**

Jeffrey Xu Yu (The Chinese University of Hong Kong), Zhihong Chong (Fudan University), Hongjun Lu (The Hong Kong University of Science and Technology), Aoying Zhou (Fudan University)

The problem of finding frequent items has been recently studied over high speed data streams. However, mining frequent item sets from transactional data streams has not been well addressed yet in terms of its bounds of memory consumption. The main difficulty is due to the nature of the exponential explosion of item sets. Given a domain of  $I$  unique items, the possible number of item sets can be up to  $2^I - 1$ . When the length of data streams approaches to a very large number  $N$ , the possibility of an item set to be frequent becomes larger and difficult to track with limited memory. However, the real killer of effective frequent item set mining is that most of existing algorithms are false-positive oriented. That is, they control memory consumption in the counting processes by an error parameter  $\epsilon$ , and allow items with support below the specified minimum support  $s$  but above  $s - \epsilon$  counted as frequent ones. Such false-positive items increase the number of false-positive frequent

item sets exponentially, which may make the problem computationally intractable with bounded memory consumption. In this paper, we developed algorithms that can effectively mine frequent item(set)s from high speed transactional data streams with a bound of memory consumption. While our algorithms are false-negative oriented, that is, certain frequent item sets may not appear in the results, the number of false-negative item sets can be controlled by a predefined parameter so that desired recall rate of frequent item sets can be guaranteed. We developed algorithms based on Chernoff bound. Our extensive experimental studies show that the proposed algorithms have high accuracy, require less memory, and consume less CPU time. They significantly outperform the existing false-positive algorithms.

## Research Session 6: XML (II)

### Indexing Temporal XML Documents

Alberto Mendelzon, Flavio Rizzolo (University of Toronto), Alejandro Vaisman (University of Buenos Aires)

Different models have been proposed recently for representing temporal data, tracking historical information, and recovering the state of the document as of any given time, in XML documents. We address the problem of indexing temporal XML documents. In particular we show that by indexing "continuous paths", i.e. paths that are valid continuously during a certain interval in a temporal XML graph, we can dramatically increase query performance. We describe in detail the indexing scheme, denoted TempIndex, and compare its performance against both a system based on a non-temporal path index, and one based on DOM.

### Schema-based Scheduling of Event Processors and Buffer Minimization for Queries on Structured Data Streams

Christoph Koch, Stefanie Scherzinger (Technische Universität Wien), Nicole Schweikardt (Humboldt Universität zu Berlin), Bernhard Stegmaier (Technische Univ. München)

We introduce an extension of the XQuery language, FluX, that supports event-based query processing and the conscious handling of main memory buffers. Purely event-based queries of this language can be executed on streaming XML data in a very direct way. We then develop an algorithm that allows to efficiently rewrite XQueries into the event-based FluX language. This algorithm uses order constraints from a DTD to schedule event handlers and to thus minimize the amount of buffering required for evaluating a query. We discuss the various technical aspects of query optimization and query evaluation within our framework. This is complemented with an experimental evaluation of our approach.

### Bloom Histogram: Path Selectivity Estimation for XML Data with Updates

Wei Wang (University of NSW), Haifeng Jiang, Hongjun Lu (Hong Kong University of Science and Technology), Jeffrey Xu Yu (The Chinese University of Hong Kong)

Cost-based XML query optimization calls for accurate estimation of the selectivity of path expressions. Some other interactive and internet applications can also benefit from such estimations. While there are a number of estimation techniques proposed in the literature, almost none of them has any guarantee on the estimation accuracy within a given space limit. In addition, most of them assume that the XML data are more or less static, i.e., with few updates. In this paper, we present a framework for XML path selectivity estimation in a dynamic context. Specifically, we propose a novel data structure, bloom histogram, to approximate XML path frequency distribution within a small space budget and to estimate the path selectivity accurately with the bloom histogram. We obtain the upper bound of its estimation error and discuss the trade-offs between the accuracy and the space limit. To support updates of bloom histograms efficiently when underlying XML data change, a dynamic summary layer is used to keep exact or more detailed XML path information. We demonstrate through our extensive experiments that the new solution can achieve significantly higher accuracy with an even smaller space than the previous methods in both static and dynamic environments.

## Research Session 7: XML And Relations

### XQuery on SQL Hosts

Torsten Grust, Sherif Sakr, Jens Teubner (University of Konstanz)

Relational database systems may be turned into efficient XML and XPath processors if the system is provided with a suitable relational tree encoding. This paper extends this relational XML processing stack and shows that an RDBMS can also serve as a highly efficient XQuery runtime environment. Our approach is purely relational: XQuery expressions are compiled into SQL code which operates on the tree encoding. The core of the compilation procedure trades XQuery's notions of variable scopes and nested iteration (FLWOR blocks) for equi-joins. The resulting relational XQuery processor closely adheres to the language semantics, e.g., it obeys node identity as well as document and sequence order, and can support XQuery's full axis feature. The system exhibits quite promising performance figures in experiments. Somewhat unexpectedly, we will also see that the XQuery compiler can make good use of SQL's OLAP functionality.

### ROX: Relational Over XML

Alan Halverson (University of Wisconsin-Madison), Vanja Josifovski, Guy Lohman, Hamid Pirahesh (IBM Almaden), Mathias Mörschel

An increasing percentage of the data needed by business applications is being generated in XML format. Storing the XML in its native format will facilitate new applications that exchange business objects in XML format and query portions of XML documents using XQuery. This paper explores the feasibility of accessing natively-stored XML data through traditional SQL interfaces, called Relational Over XML (ROX), in order to avoid the costly conversion of legacy applications to XQuery. It describes the forces that are driving the industry to evolve toward the ROX scenario as well as some of the issues raised by ROX. The impact of denormalization of data in XML documents is discussed both from a semantic and performance perspective. We also weigh the implications of ROX for manageability and query optimization. We experimentally compared the performance of a prototype of the ROX scenario to today's SQL engines, and found that good performance can be achieved through a combination of utilizing XML's hierarchical storage to store relations "pre-joined" as well as creating indices over the remaining join columns. We have developed an experimental framework using DB2 8.1 for Linux, Unix and Windows, and have gathered initial performance results that validate this approach.

### From XML View Updates to Relational View Updates: old solutions to a new problem

Vanessa Braganholo (Universidade Federal do Rio Grande do Sul), Susan Davidson (University of Pennsylvania, and INRIA-FUTURS), Carlos Heuser (Universidade Federal do Rio Grande do Sul)

This paper addresses the question of updating relational databases through XML views. Using "query trees" to capture the notions of selection, projection, nesting, grouping, and heterogeneous sets found throughout most XML query languages, we show how XML views expressed using query trees can be mapped to a set of corresponding relational views. We then show how updates on the XML view are mapped to updates on the corresponding relational views. Existing work on updating relational views can then be leveraged to determine whether or not the relational views are updatable with respect to the relational updates, and if so, to translate the updates to the underlying relational database.

## Research Session 8: Stream Mining (II)

### XWAVE: Optimal and Approximate Extended Wavelets for Streaming Data

Sudipto Guha (University of Pennsylvania), Chulyun Kim, Kyuseok Shim (Seoul National University)

Wavelet synopses have been found to be of interest in query optimization and approximate query answering. Recently, extended wavelets were proposed by Deligiannakis and Roussopoulos for datasets containing multiple measures. Extended wavelets optimize the storage utilization by attempting to store the same wavelet coefficient across different measures. This reduces the bookkeeping overhead and more coefficients can be stored. An optimal algorithm for minimizing the error in representation and an approximation algorithm for the complementary problem was

provided. However, both their algorithms take linear space. Synopsis structures are often used in environments where space is at a premium and the data arrives as a continuous stream which is too expensive to store. In this paper, we give algorithms for extended wavelets which are space sensitive, i.e., use space which is dependent on the size of the synopsis (and at most on the logarithm of the total data) and operates in a streaming fashion. We present better optimal algorithms based on dynamic programming and a near optimal approximate greedy algorithm. We also demonstrate the performance benefits of our algorithms compared to previous ones through experiments on real-life and synthetic datasets.

#### **REHIST: Relative Error Histogram Construction Algorithms**

Sudipto Guha (University of Pennsylvania), Kyuseok Shim, Jungchul Woo (Seoul National University)

Histograms and Wavelet synopses provide useful tools in query optimization and approximate query answering. Traditional histogram construction algorithms, such as V-Optimal, optimize absolute error measures for which the error in estimating a true value of 10 by 20 has the same effect of estimating a true value of 1000 by 1010. However, several researchers have recently pointed out the drawbacks of such schemes and proposed wavelet based schemes to minimize relative error measures. None of these schemes provide satisfactory guarantees -- and we provide evidence that the difficulty may lie in the choice of wavelets as the representation scheme. In this paper, we consider histogram construction for the known relative error measures. We develop optimal as well as fast approximation algorithms. We provide a comprehensive theoretical analysis and demonstrate the effectiveness of these algorithms in providing significantly more accurate answers through synthetic and real life data sets.

#### **Distributed Set-Expression Cardinality Estimation**

Abhinandan Das (Cornell University), Sumit Ganguly (IIT Kanpur), Minos Garofalakis, Rajeev Rastogi (Bell Labs, Lucent Technologies)

We consider the problem of estimating set-expression cardinality in a distributed streaming environment where rapid update streams originating at remote sites are continually transmitted to a central processing system. At the core of our algorithmic solutions for answering set-expression cardinality queries are two novel techniques for lowering data communication costs without sacrificing answer precision. Our first technique exploits global knowledge of the distribution of certain frequently occurring stream elements to significantly reduce the transmission of element state information to the central site. Our second technical contribution involves a novel way of capturing the semantics of the input set expression in a boolean logic formula, and using models (of the formula) to determine whether an element state change at a remote site can affect the set expression result. Results of our experimental study with real-life as well as synthetic data sets indicate that our distributed set-expression cardinality estimation algorithms achieve substantial reductions in message traffic compared to naive approaches that provide the same accuracy guarantees.

### **Research Session 9: Stream Query Processing**

#### **Memory-Limited Execution of Windowed Stream Joins**

Utkarsh Srivastava, Jennifer Widom (Stanford University)

We address the problem of computing approximate answers to continuous sliding-window joins over data streams when the available memory may be insufficient to keep the entire join state. One approximation scenario is to provide a maximum subset of the result, with the objective of losing as few result tuples as possible. An alternative scenario is to provide a random sample of the join result, e.g., if the output of the join is being aggregated. We show formally that neither approximation can be addressed effectively for a sliding-window join of arbitrary input streams. Previous work has only addressed the maximum-subset problem, and has implicitly used a frequency-based model of stream arrival. We address the sampling problem for this model. More importantly, we point out a broad class of applications for which an age-based model of stream arrival is more appropriate, and we address both approximation scenarios under this new model. Finally, for the case of multiple joins being executed with an overall memory constraint, we provide an algorithm for memory allocation



across the joins that optimizes a combined measure of approximation in all scenarios considered. All of our algorithms are implemented and experimental results demonstrate their effectiveness.

### **Resource Sharing in Continuous Sliding-Window Aggregates**

Arvind Arasu, Jennifer Widom (Stanford University)

We consider the problem of resource sharing when processing large numbers of continuous queries. We specifically address sliding-window aggregates over data streams, an important class of continuous operators for which sharing has not been addressed. We present a suite of sharing techniques that cover a wide range of possible scenarios: different classes of aggregation functions (algebraic, distributive, holistic), different window types (time-based, tuple-based, suffix, historical), and different input models (single stream, multiple substreams). We provide precise theoretical performance guarantees for our techniques, and show their practical effectiveness through a thorough experimental study.

### **Remembrance of Streams Past: Overload-Sensitive Management of Archived Streams**

Sirish Chandrasekaran, Michael J. Franklin (University of California, Berkeley)

This paper studies Data Stream Management Systems that combine real-time data streams with historical data, and hence access incoming streams and archived data simultaneously. A significant problem for these systems is the I/O cost of fetching historical data which inhibits processing of the live data streams. Our solution is to reduce the I/O cost for accessing the archive by retrieving only a reduced (summarized or sampled) version of the historical data. This paper does not propose new summarization or sampling techniques, but rather a framework in which multiple resolutions of summarization/sampling can be generated efficiently. The query engine can select the appropriate level of summarization to use depending on the resources currently available. The central research problem studied is whether to generate the multiple representations of archived data eagerly upon data-arrival, lazily at query-time, or in a hybrid fashion. Concrete techniques for each approach are presented, which are tied to a specific data reduction technique (random sampling). The tradeoffs among the three approaches are studied both analytically and experimentally.

### **WIC: A General-Purpose Algorithm for Monitoring Web Information Sources**

Sandeep Pandey, Kedar Dhamdhere, Christopher Olston (Carnegie Mellon University)

The Web is becoming a universal information dissemination medium, due to a number of factors including its support for content dynamicity. A growing number of Web information providers post near real-time updates in domains such as auctions, stock markets, bulletin boards, news, weather, roadway conditions, sports scores, etc. External parties often wish to capture this information for a wide variety of purposes ranging from online data mining to automated synthesis of information from multiple sources. There has been a great deal of work on the design of systems that can process streams of data from Web sources, but little attention has been paid to how to produce these data streams, given that Web pages generally require "pull-based" access. In this paper we introduce a new general-purpose algorithm for monitoring Web information sources, effectively converting pull-based sources into push-based ones. Our algorithm can be used in conjunction with continuous query systems that assume information is fed into the query engine in a push-based fashion. Ideally, a Web monitoring algorithm for this purpose should achieve two objectives: (1) timeliness and (2) completeness of information captured. However, we demonstrate both analytically and empirically using real-world data that these objectives are fundamentally at odds. When resources available for Web monitoring are limited, and the number of sources to monitor is large, it may be necessary to sacrifice some timeliness to achieve better completeness, or vice versa. To take this fact into account, our algorithm is highly parameterized and targets an application-specified balance between timeliness and completeness. In this paper we formalize the problem of optimizing for a flexible combination of timeliness and completeness, and prove that our parameterized algorithm is a 2-approximation in all cases, and in certain cases is optimal.

## Research Session 10: Managing Web Information Sources

### Similarity Search for Web Services

Xin Dong, Alon Halevy, Jayant Madhavan, Ema Nemes, Jun Zhang (University of Washington)

Web services are loosely coupled software components, published, located, and invoked across the web. The growing number of web services available within an organization and on the Web raises a new and challenging search problem: locating desired web services. Traditional keyword search is insufficient in this context: the specific types of queries users require are not captured, the very small text fragments in web services are unsuitable for keyword search, and the underlying structure and semantics of the web services are not exploited. We describe the algorithms underlying the Woogoo search engine for web services. Woogoo supports similarity search for web services, such as finding similar web-service operations and finding operations that compose with a given one. We describe novel techniques to support these types of searches, and an experimental study on a collection of over 1500 web-service operations that shows the high recall and precision of our algorithms.

### AWESOME - A Data Warehouse-based System for Adaptive Website Recommendations

Andreas Thor, Erhard Rahm (University of Leipzig)

Recommendations are crucial for the success of large websites. While there are many ways to determine recommendations, the relative quality of these recommenders depends on many factors and is largely unknown. We propose a new classification of recommenders and comparatively evaluate their relative quality for a sample website. The evaluation is performed with AWESOME (Adaptive WebSite reCOMmEndations), a new data warehouse-based recommendation system capturing and evaluating user feedback on presented recommendations. Moreover, we show how AWESOME performs an automatic and adaptive closed-loop website optimization by dynamically selecting the most promising recommenders based on continuously measured recommendation feedback. We propose and evaluate several alternatives for dynamic recommender selection including a powerful machine learning approach.

### Accurate and Efficient Crawling for Relevant Websites

Martin Ester (Simon Fraser University), Hans-Peter Kriegel, Matthias Schubert (University of Munich)

Focused web crawlers have recently emerged as an alternative to the well-established web search engines. While the well-known focused crawlers retrieve relevant web pages, there are various applications which target whole websites instead of single web pages. For example, companies are represented by websites, not by individual web pages. To answer queries targeted at websites, web directories are an established solution. In this paper, we introduce a novel focused website crawler to employ the paradigm of focused crawling for the search of relevant websites. The proposed crawler is based on a two-level architecture and corresponding crawl strategies with an explicit concept of websites. The external crawler views the web as a graph of linked websites, selects the websites to be examined next and invokes internal crawlers. Each internal crawler views the web pages of a single given website and performs focused (page) crawling within that website. Our experimental evaluation demonstrates that the proposed focused website crawler clearly outperforms previous methods of focused crawling which were adapted to retrieve websites instead of single web pages.

### Instance-based Schema Matching for Web Databases by Domain-specific Query Probing

Jiyang Wang (Hong Kong University of Science and Technology), Ji-Rong Wen (Microsoft Research Asia), Fred Lochovsky (Hong Kong University of Science and Technology), Wei-Ying Ma (Microsoft Research Asia)

In a Web database that dynamically provides information in response to user queries, two distinct schemas, interface schema (the schema users can query) and result schema (the schema users can browse), are presented to users. Each partially reflects the actual schema of the Web database. Most previous work only studied the problem of schema matching across query interfaces of Web databases. In this paper, we propose a novel schema model that distinguishes the interface and the result schema of a Web database in a specific domain. In this model, we address two significant Web database schema-matching problems: intra-site and inter-site. The first problem is crucial in

automatically extracting data from Web databases, while the second problem plays a significant role in meta-retrieving and integrating data from different Web databases. We also investigate a unified solution to the two problems based on query probing and instance-based schema matching techniques. Using the model, a cross validation technique is also proposed to improve the accuracy of the schema matching. Our experiments on real Web databases demonstrate that the two problems can be solved simultaneously with high precision and recall.

## Research Session 11: Distributed Search And Query Processing

### Computing PageRank in a Distributed Internet Search System

Yuan Wang, David DeWitt (University of Wisconsin-Madison)

Existing Internet search engines use web crawlers to download data from the Web. Page quality is measured on central servers, where user queries are also processed. This paper argues that using crawlers has a list of disadvantages. Most importantly, crawlers do not scale. Even Google, the leading search engine, indexes less than 1% of the entire Web. This paper proposes a distributed search engine framework, in which every web server answers queries over its own data. Results from multiple web servers will be merged to generate a ranked hyperlink list on the submitting server. This paper presents a series of algorithms that compute PageRank in such framework. The preliminary experiments on a real data set demonstrate that the system achieves comparable accuracy on PageRank vectors to Google's well-known PageRank algorithm and, therefore, high quality of query results.

### Enhancing P2P File-Sharing with an Internet-Scale Query Processor

Boon Thau Loo (University of California, Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley), Ryan Huebsch (University of California, Berkeley), Scott Shenker (University of California, Berkeley and International Computer Science Institute), Ion Stoica (University of California, Berkeley)

In this paper, we address the problem of designing a scalable, accurate query processor for peer-to-peer filesharing and similar distributed keyword search systems. Using a globally-distributed monitoring infrastructure, we perform an extensive study of the Gnutella file sharing network, characterizing its topology, data and query workloads. We observe that Gnutella's query processing approach performs well for popular content, but quite poorly for rare items with few replicas. We then consider an alternate approach based on Distributed Hash Tables (DHTs). We describe our implementation of PIER Search, a DHT-based system, and propose a hybrid system where Gnutella is used to locate popular items, and PIER Search for handling rare items. We develop an analytical model of the two approaches, and use it in concert with our Gnutella traces to study the trade-off between query recall and system overhead of the hybrid system. We evaluate a variety of localized schemes for identifying items that are rare and worth handling via the DHT. Lastly, we show in a live deployment on fifty nodes on two continents that it nicely complements Gnutella in its ability to handle rare items.

### Online Balancing of Range-Partitioned Data with Applications to Peer-to-Peer Systems

Prasanna Ganesan, Mayank Bawa, Hector Garcia-Molina (Stanford University)

We consider the problem of horizontally partitioning a dynamic relation across a large number of disks/nodes by the use of range partitioning. Such partitioning is often desirable in large-scale parallel databases, as well as in peer-to-peer (P2P) systems. As tuples are inserted and deleted, the partitioning may need to be adjusted, and data moved, in order to achieve storage balance across the participant disks/nodes. We propose efficient, asymptotically optimal algorithms that ensure storage balance at all times, even against an adversarial insertion and deletion of tuples. We combine the above algorithms with distributed routing structures to architect a P2P system that supports efficient range queries, while simultaneously guaranteeing storage balance.

### Network-Aware Query Processing for Stream-Based Applications

Yanif Ahmad, Uğur Çetintemel (Brown University)

This paper investigates the benefits of network awareness when processing queries in widely-distributed environments such as the Internet. We present algorithms that leverage knowledge of network characteristics (e.g., topology, bandwidth, etc.) when deciding on the network locations where the query operators are executed. Using a detailed emulation study based on realistic network models, we analyze and experimentally evaluate the proposed approaches for distributed stream processing. Our results quantify the significant benefits of the network-aware approaches and reveal the fundamental trade-off between bandwidth efficiency and result latency that arises in networked query processing.

### Data Sharing Through Query Translation in Autonomous Sources

Anastasios Kementsietsidis, Marcelo Arenas (University of Toronto)

Given a point  $q$ , a reverse  $k$  nearest neighbor (RkNN) query retrieves all the data points that have  $q$  as one of their  $k$  nearest neighbors. Existing methods for processing such queries have at least one of the following deficiencies: (i) they do not support arbitrary values of  $k$  (ii) they cannot deal efficiently with database updates, (iii) they are applicable only to 2D data (but not to higher dimensionality), and (iv) they retrieve only approximate results. Motivated by these shortcomings, we develop algorithms for exact processing of RkNN with arbitrary values of  $k$  on dynamic datasets of any dimensionality. Our methods utilize a conventional multi-dimensional index on the dataset and do not require any pre-computation. In addition to their flexibility, we experimentally verify that the proposed algorithms outperform the existing ones even in their restricted focus.

## Research Session 12: Stream Data Management Systems

### Linear Road: A Stream Data Management Benchmark

Arvind Arasu (Stanford University), Mitch Cherniack, Eduardo Galvez (Brandeis University), David Maier (OHSU/OGI), Anurag Maskey, Esther Ryvkina (Brandeis University), Michael Stonebraker, Richard Tibbetts (MIT)

This paper specifies the Linear Road Benchmark for Stream Data Management Systems (SDMS). Stream Data Management Systems process streaming data by executing continuous and historical queries while producing query results in real-time. This benchmark makes it possible to compare the performance characteristics of SDMS' relative to each other and to alternative (e.g., relational) systems. Linear Road has been endorsed as an SDMS benchmark by both Aurora (out of Brandeis University, Brown University and MIT) and STREAM (out of Stanford University). Linear Road is inspired by the increasing prevalence of "variable tolling" on highway systems throughout the world. Variable tolling uses dynamically determined factors such as congestion levels and accident proximity to calculate toll charges. Linear Road specifies a variable tolling system for a fictional urban area including such features as accident detection and alerts, traffic congestion measurements, toll calculations and historical queries. After specifying the benchmark, we describe experimental results involving two implementations: one using a commercially available Relational Database Management System and the other using Aurora. Our results show that a dedicated Stream Data Management System can outperform a Relational Database Management System by at least a factor of 5 on streaming data applications.

### Query Languages and Data Models for Database Sequences and Data Streams

Yan-Nei Law (UCLA), Haixun Wang (IBM T. J. Watson Research), Carlo Zaniolo (UCLA)

We study the fundamental limitations of relational algebra (RA) and SQL in supporting sequence and stream queries, and present effective query language and data model enrichments to deal with them. We begin by observing the well-known limitations of SQL in application domains which are important for data streams, such as sequence queries and data mining. Then we present a formal proof that, for continuous queries on data streams, SQL suffers from additional expressive power problems. We begin by focusing on the notion of non blocking (NB) queries that are the only continuous queries that can be supported on data streams. We characterize the notion of non blocking queries by showing that they are equivalent to monotonic queries. Therefore the notion of

NB-completeness for RA can be formalized as its ability to express all monotonic queries expressible in RA using only the monotonic operators of RA. We show that RA is not NB-complete, and SQL is not more powerful than RA for monotonic queries. To solve these problems, we propose extensions that allow SQL to support all the monotonic queries expressible by a Turing machine using only monotonic operators. We show that these extensions are (i) user-defined aggregates (UDAs) natively coded in SQL (rather than in an external language), and (ii) a generalization of the union operator to support the merging of multiple streams according to their timestamps. These query language extensions require matching extensions to basic relational data model to support sequences explicitly ordered by timestamps. Along with the formulation of very powerful queries, the proposed extensions entail more efficient expressions for many simple queries. In particular, we show that non blocking queries are simple to characterize according to their syntactic structure.

### Research Session 13: Auditing

#### Tamper Detection in Audit Logs

Richard Snodgrass, Shilong Yao, Christian Collberg (University of Arizona)

Audit logs are considered good practice for business systems, and are required by federal regulations for secure systems, drug approval data, medical information disclosure, financial records, and electronic voting. Given the central role of audit logs, it is critical that they are correct and inalterable. It is not sufficient to say, "our data is correct, because we store all interactions in a separate audit log." The integrity of the audit log itself must also be guaranteed. This paper proposes mechanisms within a database management system (DBMS), based on cryptographically strong one-way hash functions, that prevent an intruder, including an auditor or an employee or even an unknown bug within the DBMS itself, from silently corrupting the audit log. We propose that the DBMS store additional information in the database to enable a separate audit log validator to examine the database along with this extra information and state conclusively whether the audit log has been compromised. We show with an implementation on a high-performance storage engine that the overhead for auditing is low and that the validator can efficiently and correctly determine if the audit log has been compromised.

#### Auditing Compliance with a Hippocratic Database

Rakesh Agrawal, Roberto Bayardo, Christos Faloutsos, Jerry Kiernan, Ralf Rantzau, Ramakrishnan Srikant (IBM Almaden)

We introduce an auditing framework for determining whether a database system is adhering to its data disclosure policies. Users formulate audit expressions to specify the (sensitive) data subject to disclosure review. An audit component accepts audit expressions and returns all queries (deemed "suspicious") that accessed the specified data during their execution. The overhead of our approach on query processing is small, involving primarily the logging of each query string along with other minor annotations. Database triggers are used to capture updates in a backlog database. At the time of audit, a static analysis phase selects a subset of logged queries for further analysis. These queries are combined and transformed into an SQL audit query, which when run against the backlog database, identifies the suspicious queries efficiently and precisely. We describe the algorithms and data structures used in a DB2-based implementation of this framework. Experimental results reinforce our design choices and show the practicality of the approach.

### Research Session 14: Data Warehousing

#### High-Dimensional OLAP: A Minimal Cubing Approach

Xiaolei Li, Jiawei Han, Hector Gonzalez (University of Illinois at Urbana-Champaign)

Data cube has been playing an essential role in fast OLAP (online analytical processing) in many multi-dimensional data warehouses. However, there exist data sets in applications like bioinformatics, statistics, and text processing that are characterized by high dimensionality, e.g., over 100 dimensions, and moderate size, e.g., around  $10^6$  tuples. No feasible data cube can be constructed with such data sets. In this paper we will address the problem of developing an efficient algorithm to perform OLAP on such data sets. Experience tells us that although data analysis

tasks may involve a high dimensional space, most OLAP operations are performed only on a small number of dimensions at a time. Based on this observation, we propose a novel method that computes a thin layer of the data cube together with associated value-list indices. This layer, while being manageable in size, will be capable of supporting flexible and fast OLAP operations in the original high dimensional space. Through experiments we will show that the method has I/O costs that scale nicely with dimensionality and that are comparable to the costs of accessing a fully materialized data cube.

### **The Polynomial Complexity of Fully Materialized Coalesced Cubes**

Yannis Sismanis, Nick Roussopoulos (University of Maryland)

The data cube operator encapsulates all possible groupings of a dataset and has proved to be an invaluable tool in analyzing vast amounts of data. However its apparent exponential complexity has significantly limited its applicability to low dimensional datasets. Recently the idea of the coalesced cube was introduced, and showed that high-dimensional coalesced cubes are orders of magnitudes smaller in size than the original data cubes even when they calculate and store every possible aggregation with 100% precision. In this paper we present an analytical framework for estimating the size of coalesced cubes. By using this framework on uniform coalesced cubes we show that their size and the required computation time scales polynomially with the dimensionality of the data set and, therefore, a full data cube at 100% precision is not inherently cursed by high dimensionality. Additionally, we show that such coalesced cubes scale polynomially (and close to linearly) with the number of tuples on the dataset. We were also able to develop an efficient algorithm for estimating the size of coalesced cubes before actually computing them, based only on metadata about the cubes. Finally, we complement our analytical approach with an extensive experimental evaluation using real and synthetic data sets, and demonstrate that not only uniform but also zipfian and real coalesced cubes scale polynomially.

## **Research Session 15: Link Analysis**

### **Relational Link-based Ranking**

Floris Geerts (University of Edinburgh), Heikki Mannila, Evimaria Terzi (University of Helsinki)

Link analysis methods show that the interconnections between web pages have lots of valuable information. The link analysis methods are, however, inherently oriented towards analyzing binary relations. We consider the question of generalizing link analysis methods for analyzing relational databases. To this aim, we provide a generalized ranking framework and address its practical implications. More specifically, we associate with each relational database and set of queries a unique weighted directed graph, which we call the database graph. We explore the properties of database graphs. In analogy to link analysis algorithms, which use the Web graph to rank web pages, we use the database graph to rank partial tuples. In this way we can, e.g., extend the PageRank link analysis algorithm to relational databases and give this extension a random querier interpretation. Similarly, we extend the HITS link analysis algorithm to relational databases. We conclude with some preliminary experimental results.

### **ObjectRank: Authority-Based Keyword Search in Databases**

Andrey Balmin (IBM Almaden), Vagelis Hristidis (Florida International University), Yannis Papakonstantinou (University of California, San Diego)

The ObjectRank system applies authority-based ranking to keyword search in databases modeled as labeled graphs. Conceptually, authority originates at the nodes (objects) containing the keywords and flows to objects according to their semantic connections. Each node is ranked according to its authority with respect to the particular keywords. One can adjust the weight of global importance, the weight of each keyword of the query, the importance of a result actually containing the keywords versus being referenced by nodes containing them, and the volume of authority flow via each type of semantic connection. Novel performance challenges and opportunities are addressed. First, schemas impose constraints on the graph, which are exploited for performance purposes. Second, in order to address the issue of authority ranking with respect to the given keywords (as opposed to Google's global PageRank) we precompute single keyword ObjectRanks and combine them during run time.

We conducted user surveys and a set of performance experiments on multiple real and synthetic datasets, to assess the semantic meaningfulness and performance of ObjectRank.

### **Combating Web Spam with TrustRank**

Zoltan Gyöngyi, Hector Garcia-Molina (Stanford University), Jan Pedersen (Yahoo! Inc.)

Web spam pages use various techniques to achieve higher-than-deserved rankings in a search engine's results. While human experts can identify spam, it is too expensive to manually evaluate a large number of pages. Instead, we propose techniques to semi-automatically separate reputable, good pages from spam. We first select a small set of seed pages to be evaluated by an expert. Once we manually identify the reputable seed pages, we use the link structure of the web to discover other pages that are likely to be good. In this paper we discuss possible ways to implement the seed selection and the discovery of good pages. We present results of experiments run on the World Wide Web indexed by AltaVista and evaluate the performance of our techniques. Our results show that we can effectively filter out spam from a significant fraction of the web, based on a good seed set of less than 200 sites.

## **Research Session 16: Sensors, Grid, Pub/Sub Systems**

### **Model-Driven Data Acquisition in Sensor Networks**

Amol Deshpande (University of California, Berkeley), Carlos Guestrin (Intel Research Berkeley), Samuel Madden (MIT and Intel Research Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley), Wei Hong (Intel Research Berkeley)

Declarative queries are proving to be an attractive paradigm for interacting with networks of wireless sensors. The metaphor that "the sensor net is a database" is problematic, however, because sensors do not exhaustively represent the data in the real world. In order to map the raw sensor readings onto physical reality, a "model" of that reality is required to complement the readings. In this paper, we enrich interactive sensor querying with statistical modeling techniques. We demonstrate that such models can help provide answers that are both more meaningful, and, by introducing approximations with probabilistic confidences, significantly more efficient to compute in both time and energy. Utilizing the combination of a model and live data acquisition raises the challenging optimization problem of selecting the best sensor readings to acquire, balancing the increase in the confidence of our answer against the communication and data acquisition costs in the network. We describe an exponential time algorithm for finding the optimal solution to this optimization problem, and a polynomial-time heuristic for identifying solutions that perform well in practice. We evaluate our approach on several real-world sensor-network data sets, taking into account the real measured data and communication quality, demonstrating that our model-based approach provides a high-fidelity representation of the real phenomena and leads to significant performance gains versus traditional data acquisition techniques.

### **GridDB: A Data-Centric Overlay for Scientific Grids**

David Liu, Michael J. Franklin (University of California, Berkeley)

We present GridDB, a data-centric overlay for scientific grid data analysis. In contrast to currently deployed process-centric middleware, GridDB manages data entities rather than processes. GridDB provides a suite of services important to data analysis: a declarative interface, type-checking, interactive query processing, and memorization. We discuss several elements of GridDB: workflow/data model, query language, software architecture and query processing; and a prototype implementation. We validate GridDB by showing its modeling of real-world physics and astronomy analyses, and measurements on our prototype.

### **Towards an Internet-Scale XML Dissemination Service**

Yanlei Diao, Shariq Rizvi, Michael J. Franklin (University of California, Berkeley)

Publish/subscribe systems have demonstrated the ability to scale to large numbers of users and high data rates when providing content-based data dissemination services on the Internet. However, their services are limited by the data semantics and query expressiveness that they support. On the other

hand, the recent work on selective dissemination of XML data has made significant progress in moving from XML filtering to the richer functionality of transformation for result customization, but in general has ignored the challenges of deploying such XML-based services on an Internet-scale. In this paper, we address these challenges in the context of incorporating the rich functionality of XML data dissemination in a highly scalable system. We present the architectural design of ONYX, a system based on an overlay network. We identify the salient technical challenges in supporting XML filtering and transformation in this environment and propose techniques for solving them.

## Research Session 17: Top-K Ranking

### Efficiency-Quality Tradeoffs for Vector Score Aggregation

Pavan Kumar C Singitham, Mahathi Mahabhashyam (Stanford University), Prabhakar Raghavan (Verity Inc.)

Finding the 1-nearest neighbors to a query in a vector space is an important primitive in text and image retrieval. Here we study an extension of this problem with applications to XML and image retrieval: we have multiple vector spaces, and the query places a weight on each space. Match scores from the spaces are weighted by these weights to determine the overall match between each record and the query; this is a case of score aggregation. We study approximation algorithms that use a small fraction of the computation of exhaustive search through all records, while returning nearly the best matches. We focus on the tradeoff between the computation and the quality of the results. We develop two approaches to retrieval from such multiple vector spaces. The first is inspired by resource allocation. The second, inspired by computational geometry, combines the multiple vector spaces together with all possible query weights into a single larger space. While mathematically elegant, this abstraction is intractable for implementation. We therefore devise an approximation of this combined space. Experiments show that all our approaches (to varying extents) enable retrieval quality comparable to exhaustive search, while avoiding its heavy computational cost.

### Merging the Results of Approximate Match Operations

Sudipto Guha (University of Pennsylvania), Nick Koudas, Amit Marathe, Divesh Srivastava (AT&T Labs-Research)

Data Cleaning is an important process that has been at the center of research interest in recent years. An important end goal of effective data cleaning is to identify the relational tuple or tuples that are "most related" to a given query tuple. Various techniques have been proposed in the literature for efficiently identifying approximate matches to a query string against a single attribute of a relation. In addition to constructing a ranking (i.e., ordering) of these matches, the techniques often associate, with each match, scores that quantify the extent of the match. Since multiple attributes could exist in the query tuple, issuing approximate match operations for each of them separately will effectively create a number of ranked lists of the relation tuples. Merging these lists to identify a final ranking and scoring, and returning the top-K tuples, is a challenging task. In this paper, we adapt the well-known foot rule distance (for merging ranked lists) to effectively deal with scores. We study efficient algorithms to merge rankings, and produce the top-K tuples, in a declarative way. Since techniques for approximately matching a query string against a single attribute in a relation are typically best deployed in a database, we introduce and describe two novel algorithms for this problem and we provide SQL specifications for them. Our experimental case study, using real application data along with a realization of our proposed techniques on a commercial data base system, highlights the benefits of the proposed algorithms and attests to the overall effectiveness and practicality of our approach.

### Top-k Query Evaluation with Probabilistic Guarantees

Martin Theobald, Gerhard Weikum, Ralf Schenkel (Max-Planck Institute of Computer Science)

Top-k queries based on ranking elements of multidimensional datasets are a fundamental building block for many kinds of information discovery. The best known general-purpose algorithm for evaluating top-k queries is Fagin's threshold algorithm (TA). Since the user's goal behind top-k queries is to identify one or a few relevant and novel data items, it is intriguing to use approximative variants of TA to reduce run-time costs. This paper introduces a family of approximative top-k



algorithms based on probabilistic arguments. When scanning index lists of the underlying multidimensional data space in descending order of local scores, various forms of convolution and derived bounds are employed to predict when it is safe, with high probability, to drop candidate items and to prune the index scans. The precision and the efficiency of the developed methods are experimentally evaluated based on a large Web corpus and a structured data collection.

## Research Session 18: DBMS Architecture and Performance

### STEPS Towards Cache-resident Transaction Processing

Stavros Harizopoulos, Anastassia Ailamaki (Carnegie Mellon University)

Online transaction processing (OLTP) is a multi-billion dollar industry with high-end database servers employing state-of-the-art processors to maximize performance. Unfortunately, recent studies show that CPUs are far from realizing their maximum intended throughput because of delays in the processor caches. When running OLTP, instruction-related delays in the memory subsystem account for 25 to 40% of the total execution time. In contrast to data, instruction misses cannot be overlapped with out-of-order execution, and instruction caches cannot grow as the slower access time directly affects the processor speed. The challenge is to alleviate the instruction-related delays without increasing the cache size. We propose Steps, a technique that minimizes instruction cache misses in OLTP workloads by multiplexing concurrent transactions and exploiting common code paths. One transaction paves the cache with instructions, while close followers enjoy a nearly miss-free execution. Steps yields up to 96.7% reduction in instruction cache misses for each additional concurrent transaction, and at the same time eliminates up to 64% of mispredicted branches by loading a repeating execution pattern into the CPU. This paper (a) describes the design and implementation of Steps, (b) analyzes Steps using microbenchmarks, and (c) shows Steps performance when running TPC-C on top of the Shore storage manager.

### Write-Optimized B-Trees

Goetz Graefe (Microsoft)

Large writes are beneficial both on individual disks and on disk arrays, e.g., RAID-5. The presented design enables large writes of internal B-tree nodes and leaves. It supports both in-place updates and large append-only ("log-structured") write operations within the same storage volume, within the same B-tree, and even at the same time. The essence of the design is to make page migration inexpensive, to migrate pages while writing them, and to make such migration optional rather than mandatory as in log-structured file systems. The inexpensive page migration also aids traditional defragmentation as well as consolidation of free space needed for future large writes. These advantages are achieved with a very limited modification to conventional B-trees that also simplifies other B-tree operations, e.g., key range locking and compression. Prior proposals and prototypes implemented transacted B-tree on top of log-structured file systems and added transaction support to log-structured file systems. Instead, the presented design adds techniques and performance characteristics of log-structured file systems to traditional B-trees and their standard transaction support, notably without adding a layer of indirection for locating B-tree nodes on disk. The result retains fine-granularity locking, full transactional ACID guarantees, fast search performance, etc. expected of a modern B-tree implementation, yet adds efficient transacted page relocation and large, high-bandwidth writes.

### Cache-Conscious Radix-Decluster Projections

Stefan Manegold, Peter Boncz, Martin Kersten, Niels Nes (CWI)

As CPUs become more powerful with Moore's law and memory latencies stay constant, the impact of the memory access performance bottleneck continues to grow on relational operators like join, which can exhibit random access on a memory region larger than the hardware caches. While cache-conscious variants for various relational algorithms have been described, previous work has mostly ignored (the cost of) projection columns. However, real-life joins almost always come with projections, such that proper projection column manipulation should be an integral part of any generic join algorithm. In this paper, we analyze cache-conscious hash-join algorithms including projections on two storage schemes: N-ary Storage Model (NSM) and Decomposition Storage

Model (DSM). It turns out, that the strategy of first executing the join and only afterwards dealing with the projection columns (i.e., post-projection) on DSM, in combination with a new finely tunable algorithm called Radix-Deccluster, outperforms all previously reported projection strategies. To make this result generally applicable, we also outline how DSM Radix-Deccluster can be integrated in a NSM-based RDBMS using projection indices.

**Clotho: Decoupling Memory Page Layout from Storage Organization**

Minglong Shao, Jiri Schindler, Steven Schlosser, Anastassia Ailamaki, Gregory Ganger (Carnegie Mellon University)

As database application performance depends on the utilization of the memory hierarchy, smart data placement plays a central role in increasing locality and in improving memory utilization. Existing techniques, however, do not optimize accesses to all levels of the memory hierarchy and for all the different workloads, because each storage level uses different technology (cache, memory, disks) and each application accesses data using different patterns. Clotho is a new buffer pool and storage management architecture that decouples in-memory page layout from data organization on non-volatile storage devices to enable independent data layout design at each level of the storage hierarchy. Clotho can maximize cache and memory utilization by (a) transparently using appropriate data layouts in memory and non-volatile storage, and (b) dynamically synthesizing data pages to follow application access patterns at each level as needed. Clotho creates in-memory pages individually tailored for compound and dynamically changing workloads, and enables efficient use of different storage technologies (e.g., disk arrays or MEMS-based storage devices). This paper describes the Clotho design and prototype implementation and evaluates its performance under a variety of workloads using both disk arrays and simulated MEMS-based storage devices.

**Research Session 19: Privacy**

**Vision Paper: Enabling Privacy for the Paranoids**

Gagan Aggarwal, Mayank Bawa, Prasanna Ganesan, Hector Garcia-Molina, Krishnaram Kenthapadi, Nina Mishra, Rajeev Motwani, Utkarsh Srivastava, Dilys Thomas, Jennifer Widom, Ying Xu (Stanford University)

P3P is a set of standards that allow corporations to declare their privacy policies. Hippocratic Databases have been proposed to implement such policies within a corporation's datastore. From an end-user individual's point of view, both of these rest on an uncomfortable philosophy of trusting corporations to protect his/her privacy. Recent history chronicles several episodes when such trust has been willingly or accidentally violated by corporations facing bankruptcy courts, civil subpoenas or lucrative mergers. We contend that data management solutions for information privacy must restore controls in the individual's hands. We suggest that enabling such control will require a radical re-think on modeling, release/acquisition, and management of personal data.

**A Privacy-Preserving Index for Range Queries**

Bijit Hore, Sharad Mehrotra, Gene Tsudik (University of California, Irvine)

Database outsourcing is an emerging data management paradigm which has the potential to transform the IT operations of corporations. In this paper we address privacy threats in database outsourcing scenarios where trust in the service provider is limited. Specifically, we analyze the data partitioning (bucketization) technique and algorithmically develop this technique to build privacy-preserving indices on sensitive attributes of a relational table. Such indices enable an untrusted server to evaluate obfuscated range queries with minimal information leakage. We analyze the worst-case scenario of inference attacks that can potentially lead to breach of privacy (e.g., estimating the value of a data element within a small error margin) and identify statistical measures of data privacy in the context of these attacks. We also investigate precise privacy guarantees of data partitioning which form the basic building blocks of our index. We then develop a model for the fundamental privacy-utility tradeoff and design a novel algorithm for achieving the desired balance between privacy and utility (accuracy of range query evaluation) of the index.

### Resilient Rights Protection for Sensor Streams

Radu Sion, Mikhail Atallah, Sunil Prabhakar (Purdue University)

Today's world of increasingly dynamic computing environments naturally results in more and more data being available as fast streams. Applications such as stock market analysis, environmental sensing, web clicks and intrusion detection are just a few of the examples where valuable data is streamed. Often, streaming information is offered on the basis of a non-exclusive, single-user customer license. One major concern, especially given the digital nature of the valuable stream, is the ability to easily record and potentially "re-play" parts of it in the future. If there is value associated with such future re-plays, it could constitute enough incentive for a malicious customer (Mallory) to duplicate segments of such recorded data, subsequently re-selling them for profit. Being able to protect against such infringements becomes a necessity. In this paper we introduce the issue of rights protection for discrete streaming data through watermarking. This is a novel problem with many associated challenges including: operating in a finite window, single-pass, (possibly) high-speed streaming model, surviving natural domain specific transforms and attacks (e.g. extreme sparse sampling and summarizations), while at the same time keeping data alterations within allowable bounds. We propose a solution and analyze its resilience to various types of attacks as well as some of the important expected domain-specific transforms, such as sampling and summarization. We implement a proof of concept software (wms.\*) and perform experiments on real sensor data from the NASA Infrared Telescope Facility at the University of Hawaii, to assess encoding resilience levels in practice. Our solution proves to be well suited for this new domain. For example, we can recover an over 97% confidence watermark from a highly down-sampled (e.g. less than 8%) stream or survive stream summarization (e.g. 20%) and random alteration attacks with very high confidence levels, often above 99%.

### Research Session 20: Nearest Neighbor Search

#### Reverse kNN Search in Arbitrary Dimensionality

Yufei Tao (City University of Hong Kong), Dimitris Papadias, Xiang Lian (Hong Kong University of Science and Technology)

Given a point  $q$ , a reverse  $k$  nearest neighbor (RkNN) query retrieves all the data points that have  $q$  as one of their  $k$  nearest neighbors. Existing methods for processing such queries have at least one of the following deficiencies: (i) they do not support arbitrary values of  $k$  (ii) they cannot deal efficiently with database updates, (iii) they are applicable only to 2D data (but not to higher dimensionality), and (iv) they retrieve only approximate results. Motivated by these shortcomings, we develop algorithms for exact processing of RkNN with arbitrary values of  $k$  on dynamic datasets of any dimensionality. Our methods utilize a conventional multi-dimensional index on the dataset and do not require any pre-computation. In addition to their flexibility, we experimentally verify that the proposed algorithms outperform the existing ones even in their restricted focus.

#### GORDER: An Efficient Method for KNN Join Processing

Chenyi Xia (National University of Singapore), Hongjun Lu (Hong Kong University of Science and Technology), Beng Chin Ooi, Jin Hu (National University of Singapore)

An important but very expensive primitive operation of high-dimensional databases is the  $K$ -Nearest Neighbor (KNN) similarity join. The operation combines each point of one dataset with its KNNs in the other dataset and it provides more meaningful query results than the range similarity join. Such an operation is useful for data mining and similarity search. In this paper, we propose a novel KNN-join algorithm, called the Gorder (or the  $G$ -ordering KNN) join method. Gorder is a block nested loop join method that exploits sorting, join scheduling and distance computation filtering and reduction to reduce both I/O and CPU costs. It sorts input datasets into the  $G$ -order and applied the scheduled block nested loop join on the  $G$ -ordered data. The distance computation reduction is employed to further reduce CPU cost. It is simple and yet efficient, and handles high-dimensional data efficiently. Extensive experiments on both synthetic cluster and real life datasets were conducted, and the results illustrate that Gorder is an efficient KNN-join method and outperforms existing methods by a wide margin.

### Query and Update Efficient B+-Tree Based Indexing of Moving Objects

Christian S. Jensen (Aalborg University), Dan Lin, Beng Chin Ooi (National University of Singapore)

A number of emerging applications of data management technology involve the monitoring and querying of large quantities of continuous variables, e.g., the positions of mobile service users, termed moving objects. In such applications, large quantities of state samples obtained via sensors are streamed to a database. Indexes for moving objects must support queries efficiently, but must also support frequent updates. Indexes based on minimum bounding regions (MBRs) such as the R-tree exhibit high concurrency overheads during node splitting, and each individual update is known to be quite costly. This motivates the design of a solution that enables the B+-tree to manage moving objects. We represent moving-object locations as vectors that are time stamped based on their update time. By applying a novel linearization technique to these values, it is possible to index the resulting values using a single B+-tree that partitions values according to their timestamp and otherwise preserves spatial proximity. We develop algorithms for range and  $k$  nearest neighbor queries, as well as continuous queries. The proposal can be grafted into existing database systems cost effectively. An extensive experimental study explores the performance characteristics of the proposal and also shows that it is capable of substantially outperforming the R-tree based TPR-tree for both single and concurrent access scenarios.

## Research Session 21: Similarity Search And Applications

### Indexing Large Human-Motion Databases

Emamon Keogh, Themistoklis Palpanas, Victor B. Zordan, Dimitrios Gunopulos (University of California, Riverside), Marc Cardle (University of Cambridge)

Data-driven animation has become the industry standard for computer games and many animated movies and special effects. In particular, motion capture data recorded from live actors, is the most promising approach offered thus far for animating realistic human characters. However, the manipulation of such data for general use and re-use is not yet a solved problem. Many of the existing techniques dealing with editing motion rely on indexing for annotation, segmentation, and re-ordering of the data. Euclidean distance is inappropriate for solving these indexing problems because of the inherent variability found in human motion. The limitations of Euclidean distance stems from the fact that it is very sensitive to distortions in the time axis. A partial solution to this problem, Dynamic Time Warping (DTW), aligns the time axis before calculating the Euclidean distance. However, DTW can only address the problem of local scaling. As we demonstrate in this paper, global or uniform scaling is just as important in the indexing of human motion. We propose a novel technique to speed up similarity search under uniform scaling, based on bounding envelopes. Our technique is intuitive and simple to implement. We describe algorithms that make use of this technique, we perform an experimental analysis with real datasets, and we evaluate it in the context of a motion capture processing system. The results demonstrate the utility of our approach, and show that we can achieve orders of magnitude of speedup over the brute force approach, the only alternative solution currently available.

### On The Marriage of Lp-norms and Edit Distance

Lei Chen (University of Waterloo), Raymond Ng (University of British Columbia)

Existing studies on time series are based on two categories of distance functions. The first category consists of the L $p$ -norms. They are metric distance functions but cannot support local time shifting. The second category consists of distance functions which are capable of handling local time shifting but are non-metric. The first contribution of this paper is the proposal of a new distance function, which we call ERP ("Edit distance with Real Penalty"). Representing a marriage of L1-norm and the edit distance, ERP can support local time shifting, and is a metric. The second contribution of the paper is the development of pruning strategies for large time series databases. Given that ERP is a metric, one way to prune is to apply the triangle inequality. Another way to prune is to develop a lower bound on the ERP distance. We propose such a lower bound, which has the nice computational property that it can be efficiently indexed with a standard B+-tree. Moreover, we show that these two ways of pruning can be used simultaneously for ERP distances. Specifically, the false positives

obtained from the B+-tree can be further minimized by applying the triangle inequality. Based on extensive experimentation with existing benchmarks and techniques, we show that this combination delivers superb pruning power and search time performance, and dominates all existing strategies.

#### **Approximate NN queries on Streams with Guaranteed Error/performance Bounds**

Nick Koudas (AT&T Labs-Research), Beng Chin Ooi, Kian-Lee Tan, Rui Zhang (National University of Singapore)

In data stream applications, data arrive continuously and can only be scanned once as the query processor has very limited memory (relative to the size of the stream) to work with. Hence, queries on data streams do not have access to the entire data set and query answers are typically approximate. While there have been many studies on the  $k$  Nearest Neighbors (kNN) problem in conventional multi-dimensional databases, the solutions cannot be directly applied to data streams for the above reasons. In this paper, we investigate the kNN problem over data streams. We first introduce the  $\epsilon$ -approximate kNN (ekNN) problem that finds the approximate kNN answers of a query point  $Q$  such that the absolute error of the  $k$ -th nearest neighbor distance is bounded by  $\epsilon$ . To support ekNN queries over streams, we propose a technique called DISC (Adaptive Indexing on Streams by space-filling Curves). DISC can adapt to different data distributions to either (a) optimize memory utilization to answer ekNN queries under certain accuracy requirements or (b) achieve the best accuracy under a given memory constraint. At the same time, DISC provides efficient updates and query processing which are important requirements in data stream applications. Extensive experiments were conducted using both synthetic and real data sets and the results confirm the effectiveness and efficiency of DISC.

#### **Object Fusion in Geographic Information Systems**

Catriel Beeri, Yaron Kanza, Eliyahu Safra, Yehoshua Sagiv (The Hebrew University)

Given two geographic databases, a fusion algorithm should produce all pairs of corresponding objects (i.e., objects that represent the same real-world entity). Four fusion algorithms, which only use locations of objects, are described and their performance is measured in terms of recall and precision. These algorithms are designed to work even when locations are imprecise and each database represents only some of the real-world entities. Results of extensive experimentation are presented and discussed. The tests show that the performance depends on the density of the data sources and the degree of overlap among them. All four algorithms are much better than the current state of the art (i.e., the one-sided nearest-neighbor join). One of these four algorithms is best in all cases, at a cost of a small increase in the running time compared to the other algorithms.

#### **Maintenance of Spatial Semijoin Queries on Moving Points**

Glenn Iwerks, Hanan Samet (University of Maryland at College Park), Kenneth Smith (MITRE Corp.)

In this paper, we address the maintenance of spatial semijoin queries over continuously moving points, where points are modeled as linear functions of time. This is analogous to the maintenance of a materialized view except, as time advances, the query result may change independently of updates. As in a materialized view, we assume there is no prior knowledge of updates before they occur. We present a new approach, continuous fuzzy sets (CFS), to maintain continuous spatial semijoins efficiently. CFS is compared experimentally to a simple scaling of previous work. The result is significantly better performance of CFS compared to previous work by up to an order of magnitude in some cases.

#### **Voronoi-Based K Nearest Neighbor Search for Spatial Network Databases**

Mohammad Kolahdouzan, Cyrus Shahabi (University of Southern California)

A frequent type of query in spatial networks (e.g., road networks) is to find the  $K$  nearest neighbors (KNN) of a given query object. With these networks, the distances between objects depend on their network connectivity and it is computationally expensive to compute the distances (e.g., shortest paths) between objects. In this paper, we propose a novel approach to efficiently and accurately evaluate KNN queries in spatial network databases using first order Voronoi diagram. This approach is based on partitioning a large network to small Voronoi regions, and then pre-computing distances

both within and across the regions. By localizing the pre-computation within the regions, we save on both storage and computation and by performing across-the-network computation for only the border points of the neighboring regions, we avoid global pre-computation between every node-pair. Our empirical experiments with several real-world data sets show that our proposed solution outperforms approaches that are based on on-line distance computation by up to one order of magnitude, and provides a factor of four improvement in the selectivity of the filter step as compared to the index-based approaches.

#### **A Framework for Projected Clustering of High Dimensional Data Streams**

Charu Aggarwal (T.J. Watson Res. Ctr.), Jiawei Han, Jianyong Wang (UIUC), Philip Yu (T.J. Watson Res. Ctr.)

The data stream problem has been studied extensively in recent years, because of the great ease in collection of stream data. The nature of stream data makes it essential to use algorithms which require only one pass over the data. Recently, single-scan, stream analysis methods have been proposed in this context. However, a lot of stream data is high-dimensional in nature. High-dimensional data is inherently more complex in clustering, classification, and similarity search. Recent research discusses methods for projected clustering over high-dimensional data sets. This method is however difficult to generalize to data streams because of the complexity of the method and the large volume of the data streams. In this paper, we propose a new, high-dimensional, projected data stream clustering method, called HPStream. The method incorporates a fading cluster structure, and the projection based clustering methodology. It is incrementally updatable and is highly scalable on both the number of dimensions and the size of the data streams, and it achieves better clustering quality in comparison with the previous stream clustering methods. Our performance study with both real and synthetic data sets demonstrates the efficiency and effectiveness of our proposed framework and implementation methods.

### **Research Session 22: Query Processing**

#### **Efficient Query Evaluation on Probabilistic Databases**

Nilesh Dalvi, Dan Suciu (University of Washington)

We describe a system that supports arbitrarily complex SQL queries on probabilistic databases. The query semantics is based on a probabilistic model and the results are ranked, much like in Information Retrieval. Our main focus is efficient query evaluation, a problem that has not received attention in the past. We describe an optimization algorithm that can compute efficiently most queries. We show, however, that the data complexity of some queries is  $\#P$ -complete, which implies that these queries do not admit any efficient evaluation methods. For these queries we describe both an approximation algorithm and a Monte-Carlo simulation algorithm.

#### **Efficient Indexing Methods for Probabilistic Threshold Queries over Uncertain Data**

Reynold Cheng, Yuni Xia, Sunil Prabhakar, Rahul Shah, Jeffrey S. Vitter (Purdue University)

It is infeasible for a sensor database to contain the exact value of each sensor at all points in time. This uncertainty is inherent in these systems due to measurement and sampling errors, and resource limitations. In order to avoid drawing erroneous conclusions based upon stale data, the use of uncertainty intervals that model each data item as a range and associated probability density function (pdf) rather than a single value has recently been proposed. Querying these uncertain data introduces imprecision into answers, in the form of probability values that specify the likeliness the answer satisfies the query. These queries are more expensive to evaluate than their traditional counterparts but are guaranteed to be correct and more informative due to the probabilities accompanying the answers. Although the answer probabilities are useful, for many applications, it is only necessary to know whether the probability exceeds a given threshold -- we term these Probabilistic Threshold Queries (PTQ). In this paper we address the efficient computation of these types of queries. In particular, we develop two index structures and associated algorithms to efficiently answer PTQs. The first index scheme is based on the idea of augmenting uncertainty information to an R-tree. We establish the difficulty of this problem by mapping one-dimensional intervals to a two-dimensional space, and show that the problem of interval indexing with probabilities is significantly harder than

interval indexing which is considered a well-solved problem. To overcome the limitations of this R-tree based structure, we apply a technique called variance based clustering, where data points with similar degrees of uncertainty are clustered together. Our extensive index structure can answer the queries for various kinds of uncertainty pdfs, in an almost optimal sense. Extensive experimental evaluations are carried out to compare the effectiveness of various techniques.

#### **Probabilistic Ranking of Database Query Results**

Surajit Chaudhuri, Gautam Das (Microsoft Research), Vagelis Hristidis (Florida International University), Gerhard Weikum (MPI Informatik)

We investigate the problem of ranking answers to a database query when many tuples are returned. We adapt and apply principles of probabilistic models from Information Retrieval for structured data. Our proposed solution is domain-independent. It leverages data and workload statistics and correlations. Our ranking functions can be further customized for different applications. We present results of preliminary experiments which demonstrate the efficiency as well as the quality of our ranking system.

### **Research Session 23: Novel Models**

#### **An Annotation Management System for Relational Databases**

Deepavali Bhagwat, Laura Chiticariu, Wang-Chiew Tan, Gaurav Vijayvargiya (University of California, Santa Cruz)

We present an annotation management system for relational databases. In this system, every piece of data in a relation is assumed to have zero or more annotations associated with it and annotations are propagated along, from the source to the output, as data is being transformed through a query. Such an annotation management system is important for understanding the provenance and quality of data, especially in applications that deal with integration of scientific and biological data. We present an extension, pSQL, of a fragment of SQL that has three different types of annotation propagation schemes, each useful for different purposes. The default scheme propagates annotations according to where data is copied from. The default-all scheme propagates annotations according to where data is copied from among all equivalent formulations of a given query. The custom scheme allows a user to specify how annotations should propagate. We present a storage scheme for the annotations and describe algorithms for translating a pSQL query under each propagation scheme into one or more SQL queries that would correctly retrieve the relevant annotations according to the specified propagation scheme. For the default-all scheme, we also show how we generate finitely many queries that can simulate the annotation propagation behavior of the set of all equivalent queries, which is possibly infinite. The algorithms are implemented and the feasibility of the system is demonstrated by a set of experiments that we have conducted.

#### **Symmetric Relations and Cardinality-Bounded Multisets in Database Systems**

Kenneth Ross, Julia Stoyanovich (Columbia University)

In a binary symmetric relationship, A is related to B if and only if B is related to A. Symmetric relationships between  $k$  participating entities can be represented as multisets of cardinality  $k$ . Cardinality-bounded multisets are natural in several real-world applications. Conventional representations in relational databases suffer from several consistency and performance problems. We argue that the database system itself should provide native support for cardinality-bounded multisets. We provide techniques to be implemented by the database engine that avoid the drawbacks, and allow a schema designer to simply declare a table to be symmetric in certain attributes. We describe a compact data structure, and update methods for the structure. We describe an algebraic symmetric closure operator, and show how it can be moved around in a query plan during query optimization in order to improve performance. We describe indexing methods that allow efficient lookups on the symmetric columns. We show how to perform database normalization in the presence of symmetric relations. We provide techniques for inferring that a view is symmetric. We also describe a syntactic SQL extension that allows the succinct formulation of queries over symmetric relations.

### Algebraic Manipulation of Scientific Datasets

Bill Howe, David Maier (Oregon Health & Science University)

We investigate algebraic processing strategies for large numeric datasets equipped with a possibly irregular grid structure. Such datasets arise, for example, in computational simulations, observation networks, medical imaging, and 2-D and 3-D rendering. Existing approaches for manipulating these datasets are incomplete: The performance of SQL queries for manipulating large numeric datasets is not competitive with specialized tools. Database extensions for processing multidimensional discrete data can only model regular, rectilinear grids. Visualization software libraries are designed to process gridded datasets efficiently, but no algebra has been developed to simplify their use and afford optimization. Further, these libraries are data dependent --physical changes to data representation or organization break user programs. In this paper, we present an algebra of grid fields for manipulating both regular and irregular gridded datasets, algebraic optimization techniques, and an implementation backed by experimental results. We compare our techniques to those of spatial databases and visualization software libraries, using real examples from an Environmental Observation and Forecasting System. We find that our approach can express optimized plans inaccessible to other techniques, resulting in improved performance with reduced programming effort.

### Research Session 24: Query Processing And Optimization

#### Multi-objective Query Processing for Database Systems

Wolf-Tilo Balke (University of California, Berkeley), Ulrich Gützer (University of Tübingen)

Query processing in database systems has developed beyond mere exact matching of attribute values. Scoring database objects and retrieving only the top k matches or Pareto-optimal result sets (skyline queries) are already common for a variety of applications. Specialized algorithms using either paradigm can avoid naïve linear database scans and thus improve scalability. However, these paradigms are only two extreme cases of exploring viable compromises for each user's objectives. To find the correct result set for arbitrary cases of multi-objective query processing in databases we will present a novel algorithm for computing sets of objects that are non-dominated with respect to a set of monotonic objective functions. Naturally containing top k and skyline retrieval paradigms as special cases, this algorithm maintains scalability also for all cases in between. Moreover, we will show the algorithm's correctness and instance-optimality in terms of necessary object accesses and how the response behavior can be improved by progressively producing result objects as quickly as possible, while the algorithm is still running.

#### Lifting the Burden of History from Adaptive Query Processing

Amol Deshpande (University of California, Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research, Berkeley)

Adaptive query processing schemes attempt to reoptimize query plans during the course of query execution. A variety of techniques for adaptive query processing have been proposed, varying in the granularity at which they can make decisions. The eddy is the most aggressive of these techniques, with the flexibility to choose tuple-by-tuple how to order the application of operators. In this paper we identify and address a fundamental limitation of the original eddies proposal: the "burden of history" in routing. We observe that routing decisions have long-term effects on the state of operators in the query, and can severely constrain the ability of the eddy to adapt over time. We then propose a mechanism we call STAIRs that allows the query engine to manipulate the state stored inside the operators and undo the effects of past routing decisions. We demonstrate that eddies with STAIRs achieve both high adaptively and good performance in the face of uncertainty, outperforming prior eddy proposals by orders of magnitude.

#### A Combined Framework for Grouping and Order Optimization

Thomas Neumann, Guido Moerkotte (University of Mannheim)

Since the introduction of cost-based query optimization by Selinger et al. in their seminal paper, the performance-critical role of interesting orders has been recognized. Some algebraic operators change



interesting orders (e.g. sort and select), while others exploit them (e.g. merge join). Likewise, Wang and Cherniack (VLDB 2003) showed that existing groupings should be exploited to avoid redundant grouping operations. Ideally, the reasoning about interesting orderings and groupings should be integrated into one framework. So far, no complete, correct, and efficient algorithm for ordering and grouping inference has been proposed. We fill this gap by proposing a general two-phase approach that efficiently integrates the reasoning about orderings and groupings. Our experimental results show that with a modest increase of the time and space requirements of the preprocessing phase both orderings and groupings can be handled at the same time. More importantly, there is no additional cost for the second phase during which the plan generator changes and exploits orderings and groupings by adding operators to subplans.

#### **The Case for Precision Sharing**

Sailesh Krishnamurthy, Michael J. Franklin (University of California, Berkeley), Joseph M. Hellerstein (University of California, Berkeley and Intel Research Berkeley), Garrett Jacobson (University of California, Berkeley)

Sharing has emerged as a key idea of static and adaptive stream query processing systems. Inherent in these systems is a tension between sharing common work and avoiding unnecessary work. Increased sharing has generally led to more unnecessary work. Our approach of precision sharing aims to share aggressively without unnecessary work. We show why "adaptive" tuple lineage is more generally applicable and use it for precisely shared static dataflows. We also show how "static" ordering constraints can be used for precision sharing in adaptive systems. Finally, we report an experimental study of precision sharing.

### **Industrial Session 1: Novel SQL Extensions**

#### **Returning Modified Rows - SELECT Statements with Side Effects**

Andreas Behm, Serge Rielau, Richard Swagerman (IBM Toronto Lab)

SQL in the IBM® DB2® Universal Database™ for Linux®, UNIX®, and Windows® (DB2 UDB) database management product has been extended to support nested INSERT, UPDATE, and DELETE operations in SELECT statements. This allows database applications additional processing on modified rows. Within a single unit of work, applications can retrieve a result set containing the modified rows from a table or view modified by an SQL data-change operation. This eliminates the need to select the row after an INSERT or UPDATE, or before a DELETE statement. As a result, fewer network round trips, less server CPU time, fewer cursors, and less server memory are required. In addition, deadlocks can be avoided. The proposed approach is integrated with the set semantics of SQL, and does not require any procedural logic or modifications on the underlying relational data model. Pipelining multiple update, insert and delete operations using the same source data provides a very efficient way for multi-table data-change statements typically found in ETL (extraction, transformation, load) applications. We demonstrate significant performance benefit with our experiences in the TPC-C benchmark. Experimental results show that the new SQL is more efficient in query execution compared to classic SQL.

#### **PIVOT and UNPIVOT: Optimization and Execution Strategies in an RDBMS and Scalable Pull-Based Caching for Erratic Data Sources**

Conor Cunningham, César Galindo-Legaria, Goetz Graefe (Microsoft Corp.)

PIVOT and UNPIVOT, two operators on tabular data that exchange rows and columns, enable data transformations useful in data modeling, data analysis, and data presentation. They can quite easily be implemented inside a query processor, much like select, project, and join. Such a design provides opportunities for better performance, both during query optimization and query execution. We discuss query optimization and execution implications of this integrated design and evaluate the performance of this approach using a prototype implementation in Microsoft SQL Server.

## **A Multi-Purpose Implementation of Mandatory Access Control in Relational Database Management Systems**

Walid Rjaibi, Paul Bird (IBM Toronto Lab)

Mandatory Access Control (MAC) implementations in Relational Database Management Systems (RDBMS) have focused solely on Multilevel Security (MLS). MLS has posed a number of challenging problems to the database research community, and there has been an abundance of research work to address those problems. Unfortunately, the use of MLS RDBMS has been restricted to a few government organizations where MLS is of paramount importance such as the intelligence community and the Department of Defense. The implication of this is that the investment of building an MLS RDBMS cannot be leveraged to serve the needs of application domains where there is a desire to control access to objects based on the label associated with that object and the label associated with the subject accessing that object, but where the label access rules and the label structure do not necessarily match the MLS two security rules and the MLS label structure. This paper introduces a flexible and generic implementation of MAC in RDBMS that can be used to address the requirements from a variety of application domains, as well as to allow an RDBMS to efficiently take part in an end-to-end MAC enterprise solution. The paper also discusses the extensions made to the SQL compiler component of an RDBMS to incorporate the label access rules in the access plan it generates for an SQL query, and to prevent unauthorized leakage of data that could occur as a result of traditional optimization techniques performed by SQL compilers.

## **Industrial Session 2: New DBMS Architectures and Performance**

### **Hardware Acceleration in Commercial Databases: A Case Study of Spatial Operations**

Nag Nagender Bandi (University of California, Santa Barbara), Chengyu Sun (California State University, Los Angeles), Divyakant Agrawal, Amr El Abbadi (University of California, Santa Barbara)

Traditional databases have focused on the issue of reducing I/O cost as it is the bottleneck in many operations. As databases become increasingly accepted in areas such as Geographic Information Systems (GIS) and Bio-informatics, commercial DBMS need to support data types for complex data such as spatial geometries and protein structures. These non-conventional data types and their associated operations present new challenges. In particular, the computational cost of some spatial operations can be orders of magnitude higher than the I/O cost. In order to improve the performance of spatial query processing, innovative solutions for reducing this computational cost are beginning to emerge. Recently, it has been proposed that hardware acceleration of an off-the-shelf graphics card can be used to reduce the computational cost of spatial operations. However, this proposal is preliminary in that it establishes the feasibility of the hardware assisted approach in a stand-alone setting but not in a real-world commercial database. In this paper we present an architecture to show how hardware acceleration of an off-the-shelf graphics card can be integrated into a popular commercial database to speed up spatial queries. Extensive experimentation with real-world datasets shows that significant improvement in the performance of spatial operations can be achieved with this integration. The viability of this approach underscores the significance of a tighter integration of hardware acceleration into commercial databases for spatial applications.

### **P\*TIME: Highly Scalable OLTP DBMS for Managing Update-Intensive Stream Workload**

Sang K. Cha (Transact In Memory, Inc.), Changbin Song (Seoul National University)

Over the past thirty years since the system R and Ingres projects started to lay the foundation for today's RDBMS implementations, the underlying hardware and software platforms have changed dramatically. However, the fundamental RDBMS architecture, especially, the storage engine architecture, largely remains unchanged. While this conventional architecture may suffice for satisfying most of today's applications, its deliverable performance range is far from meeting the so-called growing "real-time enterprise" demand of acquiring and querying high-volume update data streams cost-effectively. P\*TIME is a new, memory-centric light-weight OLTP RDBMS designed and built from scratch to deliver orders of magnitude higher scalability on commodity SMP hardware than existing RDBMS implementations, not only in search but also in update performance. Its storage engine layer incorporates our previous innovations for exploiting engine-level micro-

parallelism such as differential logging and optimistic latch-free index traversal concurrency control protocol. This paper presents the architecture and performance of P\*TIME and reports our experience of deploying P\*TIME as the stock market database server at one of the largest on-line brokerage firms.

### **Generating Thousand Benchmark Queries in Seconds**

Meikel Poess (Oracle Corp.), John M. Stephens, Jr. (Gradient Systems)

The combination of an exponential growth in the amount of data managed by a typical business intelligence system and the increased competitiveness of a global economy has propelled decision support systems (DSS) from the role of exploratory tools employed by a few visionary companies to become a core requirement for a competitive enterprise. That same maturation has often resulted in a selection process that requires an ever more critical system evaluation and selection to be completed in an increasingly short period of time. While there have been some advances in the generation of data sets for system evaluation (see [3]), the quantification of query performance has often relied on models and methodologies that were developed for systems that were more simplistic, less dynamic, and less central to a successful business. In this paper we present QGEN, a flexible, high-level query generator optimized for decision support system evaluation. QGEN is able to generate arbitrary query sets, which conform to a selected statistical profile without requiring that the queries be statically defined or disclosed prior to testing. Its novel design links query syntax with abstracted data distributions, enabling users to parameterize their query workload to match an emerging access pattern or data set modification. This results in query sets that retain comparability for system comparisons while reflecting the inherent dynamism of operational systems, and which provide a broad range of syntactic and semantic coverage, while remaining focused on appropriate commonalities within a particular evaluation process or business segment.

## **Industrial Session 3: Semantic Query Approaches**

### **Supporting Ontology-Based Semantic Matching in RDBMS**

Souripriya Das, Eugene Chong, George Eadon, Jagannathan Srinivasan (Oracle Corp.)

Ontologies are increasingly being used to build applications that utilize domain-specific knowledge. This paper addresses the problem of supporting ontology-based semantic matching in RDBMS. Specifically, 1) A set of SQL operators, namely `ONT_RELATED`, `ONT_EXPAND`, `ONT_DISTANCE`, and `ONT_PATH`, are introduced to perform ontology-based semantic matching, 2) A new indexing scheme `ONT_INDEXTYPE` is introduced to speed up ontology-based semantic matching operations, and 3) System-defined tables are provided for storing ontologies specified in OWL. Our approach enables users to reference ontology data directly from SQL using the semantic match operators, thereby opening up possibilities of combining with other operations such as joins as well as making the ontology-driven applications easy to develop and efficient. In contrast, other approaches use RDBMS only for storage of ontologies and querying of ontology data is typically done via APIs. This paper presents the ontology-related functionality including inferencing, discusses how it is implemented on top of Oracle RDBMS, and illustrates the usage with several database applications.

### **BioPatentMiner: An Information Retrieval System for BioMedical Patents**

Sougata Mukherjea, Bhuvan Bamba (IBM India Research Lab)

Before undertaking new biomedical research, identifying concepts that have already been patented is essential. Traditional keyword based search on patent databases may not be sufficient to retrieve all the relevant information, especially for the biomedical domain. More sophisticated retrieval techniques are required. This paper presents BioPatentMiner, a system that facilitates information retrieval from biomedical patents. It integrates information from the patents with knowledge from biomedical ontologies to create a Semantic Web. Besides keyword search and queries linking the properties specified by one or more RDF triples, the system can discover Semantic Associations between the resources. The system also determines the importance of the resources to rank the results of a search and prevent information overload while determining the Semantic Associations.

### **Flexible String Matching Against Large Databases in Practice**

Nick Koudas, Amit Marathe, Divesh Srivastava (AT&T Labs–Research)

Data Cleaning is an important process that has been at the center of research interest in recent years. Poor data quality is the result of a variety of reasons, including data entry errors and multiple conventions for recording database fields, and has a significant impact on a variety of business issues. Hence, there is a pressing need for technologies that enable flexible (fuzzy) matching of string information in a database. Cosine similarity with tf-idf is a well-established metric for comparing text, and recent proposals have adapted this similarity measure for flexibly matching a query string with values in a single attribute of a relation. In deploying tf-idf based flexible string matching against real AT&T databases, we observed that this technique needed to be enhanced in many ways. First, along the functionality dimension, where there was a need to flexibly match along multiple string-valued attributes, and also take advantage of known semantic equivalences. Second, we identified various performance enhancements to speed up the matching process, potentially trading off a small degree of accuracy for substantial performance gains. In this paper, we report on our techniques and experience in dealing with flexible string matching against real AT&T databases.

### **Industrial Session 4: Automatic Tuning in Commercial DVMSS**

#### **DB2 Design Advisor: Integrated Automatic Physical Database Design**

Daniel Zilio (IBM Toronto Lab), Jun Rao (IBM Almaden), Sam Lightstone (IBM Toronto Lab), Guy Lohman (IBM Almaden), Adam Storm, Christian Garcia-Arellano (IBM Toronto Lab), Scott Fadden (IBM Portland)

The DB2 Design Advisor in IBM DB2 Universal Database (DB2 UDB) Version 8.2 for Linux, UNIX and Windows is a tool that, for a given workload, automatically recommends physical design features that are any subset of indexes, materialized query tables (also called materialized views), shared-nothing database partitionings, and multidimensional clustering of tables. Our work is the very first industrial-strength tool that covers the design of as many as four different features, a significant advance to existing tools, which support no more than just indexes and materialized views. Building such a tool is challenging, because of not only the large search space introduced by the interactions among features, but also the extensibility needed by the tool to support additional features in the future. We adopt a novel "hybrid" approach in the Design Advisor that allows us to take important interdependencies into account as well as to encapsulate design features as separate components to lower the reengineering cost. The Design Advisor also features a built-in module that automatically reduces the given workload, and therefore provides great scalability for the tool. Our experimental results demonstrate that our tool can quickly provide good physical design recommendations that satisfy users' requirements.

#### **Automatic SQL Tuning in Oracle 10g**

Benoit Dageville, Dinesh Das, Karl Dias, Khaled Yagoub, Mohamed Zait, Mohamed Ziauddin (Oracle Corp.)

SQL tuning is a very critical aspect of database performance tuning. It is an inherently complex activity requiring a high level of expertise in several domains: query optimization, to improve the execution plan selected by the query optimizer; access design, to identify missing access structures; and SQL design, to restructure and simplify the text of a badly written SQL statement. Furthermore, SQL tuning is a time consuming task due to the large volume and evolving nature of the SQL workload and its underlying data. In this paper we present the new Automatic SQL Tuning feature of Oracle 10g. This technology is implemented as a core enhancement of the Oracle query optimizer and offers a comprehensive solution to the SQL tuning challenges mentioned above. Automatic SQL Tuning introduces the concept of SQL profiling to transparently improve execution plans. It also generates SQL tuning recommendations by performing cost-based access path and SQL structure "what-if" analyses. This feature is exposed to the user through both graphical and command line interfaces. The Automatic SQL Tuning is an integral part of the Oracle's framework for self-managing databases. The superiority of this new technology is demonstrated by comparing the results of Automatic SQL Tuning to manual tuning using a real customer workload.

**Database Tuning Advisor for Microsoft SQL Server 2005**

Sanjay Agrawal, Surajit Chaudhuri, Lubor Kollar, Arun Marathe, Vivek Narasayya, Manoj Syamala (Microsoft Corp.)

The Database Tuning Advisor (DTA) that is part of Microsoft SQL Server 2005 is an automated physical database design tool that significantly advances the state-of-the-art in several ways. First, DTA is capable to providing an integrated physical design recommendation for horizontal partitioning, indexes, and materialized views. Second, unlike today's physical design tools that focus solely on performance, DTA also supports the capability for a database administrator (DBA) to specify manageability requirements while optimizing for performance. Third, DTA is able to scale to large databases and workloads using several novel techniques including: (a) workload compression (b) reduced statistics creation and (c) exploiting test server to reduce load on production server. Finally, DTA greatly enhances scriptability and customization through the use of a public XML schema for input and output. This paper provides an overview of DTA's novel functionality, the rationale for its architecture, and demonstrates DTA's quality and scalability on large customer workloads.

**Industrial Session 5: XML Implementations Automatic Physical Design And Indexing****Query Rewrite for XML in Oracle XML DB**

Muralidhar Krishnaprasad, Zhen Liu, Anand Manikutty, James W. Warner, Vikas Arora, Susan Kotsovolos (Oracle Corp.)

Oracle XML DB integrates XML storage and querying using the Oracle relational and object relational framework. It has the capability to physically store XML documents by shredding them as relational or object relational data, and creating logical XML documents using SQL/XML publishing functions. However, querying XML in a relational or object relational database poses several challenges. The biggest challenge is to efficiently process queries against XML in a database whose fundamental storage is table-based and whose fundamental query engine is tuple-oriented. In this paper, we present the 'XML Query Rewrite' technique used in Oracle XML DB. This technique integrates querying XML using XPath embedded inside SQL operators and SQL/XML publishing functions with the object relational and relational algebra. A common set of algebraic rules is used to reduce both XML and object queries into their relational equivalent. This enables a large class of XML queries over XML type tables and views to be transformed into their semantically equivalent relational or object relational queries. These queries are then amenable to classical relational optimizations yielding XML query performance comparable to relational. Furthermore, this rewrite technique lays out a foundation to enable rewrite of XQuery over XML.

**Indexing XML Data Stored in a Relational Database**

Shankar Pal, Istvan Cseri, Gideon Schaller, Oliver Seeliger, Leo Giakoumakis, Vasili Zolotov (Microsoft Corp.)

As XML usage grows for both data-centric and document-centric applications, introducing native support for XML data in relational databases brings significant benefits. It provides a more mature platform for the XML data model and serves as the basis for interoperability between relational and XML data. Whereas query processing on XML data shredded into one or more relational tables is well understood, it provides limited support for the XML data model. XML data can be persisted as a byte sequence (BLOB) in columns of tables to support the XML model more faithfully. This introduces new challenges for query processing such as the ability to index the XML blob for good query performance. This paper reports novel techniques for indexing XML data in the upcoming version of Microsoft® SQL Server™, and how it ties into the relational framework for query processing.

**Automated Statistics Collection in DB2 UDB**

Ashraf Abouinaga, Peter Haas (IBM Almaden), Mokhtar Kandil, Sam Lightstone (IBM Toronto Development Lab), Guy Lohman, Volker Markl (IBM Almaden), Ivan Popivanov (IBM Toronto Development Lab), Vijayshankar Raman (IBM Almaden)

The use of inaccurate or outdated database statistics by the query optimizer in a relational DBMS often results in a poor choice of query execution plans and hence unacceptably long query processing times. Configuration and maintenance of these statistics has traditionally been a time-consuming manual operation, requiring that the database administrator (DBA) continually monitor query performance and data changes in order to determine when to refresh the statistics values and when and how to adjust the set of statistics that the DBMS maintains. In this paper we describe the new Automated Statistics Collection (ASC) component of DB2 Stinger. This autonomic technology frees the DBA from the tedious task of manually supervising the collection and maintenance of database statistics. ASC monitors both the update-delete-insert (UDI) activities on the data as well as query feedback (QF), i.e., the results of the queries that are executed on the data. ASC uses these two sources of information to automatically decide which statistics to collect and when to collect them. This combination of UDI-driven and QF-driven autonomic processes ensures that the system can handle unforeseen queries while also ensuring good performance for frequent and important queries. We present the basic concepts, architecture, and key implementation details of ASC in DB2 Stinger, and present a case study showing how the use of ASC can speed up a query workload by orders of magnitude without requiring any DBA intervention.

**High Performance Index Build Algorithms for Intranet Search Engines**

Marcus Fontoura, Eugene Shekita, Jason Zien, Sridhar Rajagopalan, Andreas Neumann (IBM Almaden)

There has been a substantial amount of research on high-performance algorithms for constructing an inverted text index. However, constructing the inverted index in a intranet search engine is only the final step in a more complicated index build process. Among other things, this process requires an analysis of all the data being indexed to compute measures like PageRank. The time to perform this global analysis step is significant compared to the time to construct the inverted index, yet it has not received much attention in the research literature. In this paper, we describe how the use of slightly outdated information from global analysis and a fast index construction algorithm based on radix sorting can be combined in a novel way to significantly speed up the index build process without sacrificing search quality.

**Automating the Design of Multi-dimensional Clustering Tables in Relational Databases**

Sam Lightstone (IBM Toronto Lab), Bishwaranjan Bhattacharjee (IBM T.J. Watson Research Center)

The ability to physically cluster a database table on multiple dimensions is a powerful technique that offers significant performance benefits in many OLAP, warehousing, and decision-support systems. An industrial implementation of this technique for the DB2® Universal Database™ (DB2 UDB) product, called multidimensional clustering (MDC), which co-exists with other classical forms of data storage and indexing methods, was described in VLDB 2003. This paper describes the first published model for automating the selection of clustering keys in single-dimensional and multidimensional relational databases that use a cell/block storage structure for MDC. For any significant dimensionality (3 or more), the possible solution space is combinatorially complex. The automated MDC design model is based on what-if query cost modeling, data sampling, and a search algorithm for evaluating a large constellation of possible combinations. The model is effective at trading the benefits of potential combinations of clustering keys against data scarcity and performance. It also effectively selects the granularity at which dimensions should be used for clustering (such as week of year versus month of year). We show results from experiments indicating that the model provides design recommendations of comparable quality to those made by human experts. The model has been implemented in the IBM® DB2 UDB for Linux®, UNIX® and Windows® Version 8.2 release.

## Industrial Session 6: Data Management with RFIDs and Ease of Use Territories

### Integrating Automatic Data Acquisition with Business Processes Experiences with SAP's Auto-ID Infrastructure

Christof Bornhövd, Tao Lin, Stephan Haller, Joachim Schaper (SAP Research)

Smart item technologies, like RFID and sensor networks, are considered to be the next big step in business process automation. Through automatic and real-time data acquisition, these technologies can benefit a great variety of industries by improving the efficiency of their operations. SAP's Auto-ID infrastructure enables the integration of RFID and sensor technologies with existing business processes. In this paper we give an overview of the existing infrastructure, discuss lessons learned from successful customer pilots, and point out some of the open research issues.

### Managing RFID Data

Sudarshan S. Chawathe (University of Maryland), Venkat Krishnamurthy, Sridhar Ramachandran (OAT Systems), and Sanjay Sarma (MIT)

Radio-Frequency Identification (RFID) technology enables sensors to efficiently and inexpensively track merchandise and other objects. The vast amount of data resulting from the proliferation of RFID readers and tags poses some interesting challenges for data management. We present a brief introduction to RFID technology and highlight a few of the data management challenges.

### Production Database Systems: Making Them Easy is Hard Work

David Campbell (Microsoft Corporation)

Enterprise capable database products have evolved into incredibly complex systems, some of which present hundreds of configuration parameters to the system administrator. So, while the processing and storage costs for maintaining large volumes of data have plummeted, the human costs associated with maintaining the data have continued to rise. In this presentation, we discuss the framework and approach used by the team who took Microsoft SQL Server from a state where it had several hundred configuration parameters to a system that can configure itself and respond to changes in workload and environment with little human intervention.

## Industrial Session 7: Data Management Challenges in Life Sciences and Email Systems

### Managing Data from High-Throughput Genomic Processing: A Case Study (invited)

Toby Bloom, Ted Sharpe (Broad Institute of MIT and Harvard)

Genomic data has become the canonical example of very large, very complex data sets. As such, there has been significant interest in ways to provide targeted database support to address issues that arise in genomic processing. Whether genomic data is truly a special case, or just another application area exhibiting problems common to other domains, is an as yet unanswered question. In this abstract, we explore the structure and processing requirements of a large-scale genome sequencing center, as a case study of the issues that arise in genomic data managements, and as a means to compare those issues with those that arise in other domains.

### Database Challenges in the Integration of Biomedical Data Sets (invited)

Rakesh Nagarajan (Washington University, School of Medicine), Mushtaq Ahmed (Persistent Systems Pvt. Ltd.), Aditya Phatak (Persistent Systems Pvt. Ltd.)

The clinical and basic science research domains present exciting and difficult data integration issues. Solving these problems is crucial as current research efforts in the field of biomedicine heavily depend upon integrated storage, querying, analysis, and visualization of clinicopathology information, genomic annotation, and large scale functional genomic research data sets. Such large scale experimental analyses are essential to decipher the pathophysiological processes occurring in most human diseases so that they may be effectively treated. In this paper, we discuss the challenges of integration of multiple biomedical data sets not only at the university level but also at the national level and present the data warehousing based solution we have employed at Washington University

School of Medicine. We also describe the tools we have developed to store, query, analyze, and visualize these data sets together.

#### **The Bloomba Personal Content Database (invited)**

Raymie Stata (Stata Labs and University of California, Santa Cruz), Patrick Hunt (Stata Labs), Thiruvalluvan M. G. (iSoftTech Ltd.)

We believe continued growth in the volume of personal content, together with a shift to a multi-device personal computing environment, will inevitably lead to the development of Personal Content Databases (PCDBs). These databases will make it easier for users to find, use, and replicate large, heterogeneous repositories of personal content. In this paper, we describe the PCDB used to power Bloomba, a commercial personal information manager in broad use. We highlight areas where the special requirements of personal content and personal platforms have influenced the design and implementation of our PCDB. We also discuss what we have and have not been able to leverage from the database community and suggest a few lines of research that would be useful to builders of PCDBs.

### **Industrial Session 8: Issues in Data Warehousing**

#### **Trends in data warehousing: A Practitioner's View (invited)**

William O'Connell (IBM Toronto Lab)

This talk will present emerging data warehousing reference architectures, and focus on trends and directions that are shaping these enterprise installations. Implications will be highlighted, including both of new and old technology. Stack seamless integration is also pivotal to success, which also has significant implications on things such as Metadata.

#### **Technology Challenges in a Data Warehouse (invited)**

Ramesh Bhashyam (Teradata Corporation)

This presentation will discuss several database technology challenges that are faced when building a data warehouse. It will touch on the challenges posed by high capacity drives and the mechanisms in Teradata DBMS to address that. It will consider the features and capabilities required of a database in a mixed application environment of a warehouse and some solutions to address that.

#### **Sybase IQ Multiplex – Designed For Analytics (invited)**

Roger MacNicol, Blaine French (Sybase Inc.)

The internal design of database systems has traditionally given primacy to the needs of transactional data. A radical re-evaluation of the internal design giving primacy to the needs of complex analytics shows clear benefits in large databases for both single servers and in multi-node shared-disk grid computing. This design supports the trend to keep more years of more finely grained data online by ameliorating the data explosion problem.

### **Demo Group 1: Data Mining**

#### **GPX: Interactive Mining of Gene Expression Data**

Daxin Jiang, Jian Pei, Aidong Zhang (State University of New York at Buffalo, USA)

Discovering co-expressed genes and coherent expression patterns in gene expression data is an important data analysis task in bioinformatics research and biomedical applications. Although various clustering methods have been proposed, two tough challenges still remain on how to integrate the users' domain knowledge and how to handle the high connectivity in the data. Recently, we have systematically studied the problem and proposed an effective approach. In this paper, we describe a demonstration of GPX (for Gene Pattern eXplorer), an integrated environment for interactive exploration of coherent expression patterns and co-expressed genes in gene expression data. GPX integrates several novel techniques, including the coherent pattern index graph, a gene annotation panel, and a graphical interface, to adopt users' domain knowledge and support explorative operations in the clustering procedure. The GPX system as well as its techniques will be



showcased, and the progress of GPX will be exemplified using several real-world gene expression data sets.

### **Computing Frequent Itemsets Inside Oracle 10G**

Wei Li, Ari Mozes (Oracle Corporation, USA)

Frequent itemset counting is the first step for most association rule algorithms and some classification algorithms. It is the process of counting the number of occurrences of a set of items that happen across many transactions. The goal is to find those items which occur together most often. Expressing this functionality in RDBMS engines is difficult for two reasons. First, it leads to extremely inefficient execution when using existing RDBMS operations since they are not designed to handle this type of workload. Second, it is difficult to express the special output type of itemsets. In Oracle 10G, we introduce a new SQL table function which encapsulates the work of frequent itemset counting. It accepts the input dataset along with some user-configurable information, and it directly produces the frequent itemset results. We present examples of typical computations with frequent itemset counting inside Oracle 10G. We also describe how Oracle dynamically adapts during frequent itemset execution as a result of changes in the nature of the data as well as changes in the available system resources.

### **StreamMiner: A Classifier Ensemble-based Engine to Mine Concept-drifting Data Streams**

Wei Fan (IBM T.J. Watson Research, USA)

We demonstrate StreamMiner, a random decision-tree ensemble based engine to mine data streams. A fundamental challenge in data stream mining applications (e.g., credit card transaction authorization, security buy-sell transaction, and phone call records, etc) is concept-drift or the discrepancy between the previously learned model and the true model in the new data. The basic problem is the ability to judiciously select data and adapt the old model to accurately match the changed concept of the data stream. StreamMiner uses several techniques to support mining over data streams with possible concept-drifts. We demonstrate the following two key functionalities of StreamMiner: Detecting possible concept-drift on the fly when the trained streaming model is used to classify incoming data streams without knowing the ground truth. Systematic data selection of old data and new data chunks to compute the optimal model that best fits on the changing data streams.

### **Semantic Mining and Analysis of Gene Expression Data**

Xin Xu, Gao Cong, Beng Chin Ooi, Kian-Lee Tan, Anthony K.H.Tung (National University of Singapore)

Association rules can reveal biological relevant relationship between genes and environments / categories. However, most existing association rule mining algorithms are rendered impractical on gene expression data, which typically contains thousands or tens of thousands of columns (gene expression levels), but only tens of rows (samples). The main problem is that these algorithms have an exponential dependence on the number of columns. Another shortcoming is evident that too many associations are generated from such kind of data. To this end, we have developed a novel depth-first row-wise algorithm FARMER that is specially designed to efficiently discover and cluster association rules into interesting rule groups (IRGs) that satisfy user-specified minimum support, confidence and chi-square value thresholds on biological datasets as opposed to finding association rules individually. Based on FARMER, we have developed a prototype system that integrates semantic mining and visual analysis of IRGs mined from gene expression data.

### **HOS-Miner: A System for Detecting Outlying Subspaces of High-dimensional Data**

Ji Zhang, Meng Lou (University of Toronto, Canada), Tok Wang Ling (National University of Singapore), Hai Wang (University of Toronto, Canada)

We identify a new and interesting high-dimensional outlier detection problem in this paper, that is, detecting the subspaces in which given data points are outliers. We call the subspaces in which a data point is an outlier as its Outlying Subspaces. In this paper, we will propose the prototype of a dynamic subspace search system, called HOS-Miner (HOS stands for High-dimensional Outlying Subspaces), that utilizes a sample-based learning process to effectively identify the outlying subspaces of a given point.

### **VizTree: a Tool for Visually Mining and Monitoring Massive Time Series Databases**

Jessica Lin, Eamonn Keogh, Stefano Lonardi (University of California, Riverside, USA), Jeffrey P. Lankford, Donna M. Nystrom (The Aerospace Corporation, USA)

Moments before the launch of every space vehicle, engineering discipline specialists must make a critical go/no-go decision. The cost of a false positive, allowing a launch in spite of a fault, or a false negative, stopping a potentially successful launch, can be measured in the tens of millions of dollars, not including the cost in morale and other more intangible detriments. The Aerospace Corporation is responsible for providing engineering assessments critical to the go/no-go decision for every Department of Defense (DoD) launch vehicle. These assessments are made by constantly monitoring streaming telemetry data in the hours before launch. For this demonstration, we will introduce VizTree, a novel time-series visualization tool to aid the Aerospace analysts who must make these engineering assessments. VizTree was developed at the University of California, Riverside and is unique in that the same tool is used for mining archival data and monitoring incoming live telemetry. Unlike other time series visualization tools, VizTree can scale to very large databases, giving it the potential to be a generally useful data mining and database tool.

### **Demo Group 2: Distributed Databases**

#### **An Electronic Patient Record "on Steroids": Distributed, Peer-to-Peer, Secure and Privacy-conscious**

Serge Abiteboul (INRIA, France), Bogdan Alexe (Bell Labs, Lucent Technologies, USA), Omar Benjelloun, Bogdan Cautis (INRIA, France), Irini Fundulaki (Bell Labs, Lucent Technologies, USA), Tova Milo (INRIA, France, and Tel-Aviv University, Israel), Arnaud Sahuguet (Tel-Aviv University, Israel)

We demonstrate a novel solution for the privacy-conscious integration of distributed data, that relies on the combination of two key technologies: Active XML, that provides a highly flexible peer-to-peer paradigm for data integration, and GUPster, that unifies the enforcement of access control and source descriptions. While each of these two technologies has been demonstrated separately before, we show here that their synergy yields a powerful generic platform that seamlessly handles data integration and privacy enforcement tasks in a highly distributed setting.

#### **Queries and Updates in the coDB Peer to Peer Database System**

Enrico Franconi (Free University of Bozen-Bolzano, Italy), Gabriel Kuper (University of Trento, Italy), Andrei Lopatenko (Free University of Bozen-Bolzano, Italy, and University of Manchester, UK), and Ilya Zaihrayeu (University of Trento, Italy)

In this short paper we present the coDB P2P DB system. A network of databases, possibly with different schemas, are interconnected by means of GLAV coordination rules, which are inclusions of conjunctive queries, with possibly existential variables in the head; coordination rules may be cyclic. Each node can be queried in its schema for data, which the node can fetch from its neighbors, if a coordination rule is involved.

#### **A-ToPSS: A Publish/Subscribe System Supporting Imperfect Information Processing**

Haifeng Liu, Hans-Arno Jacobsen (University of Toronto, Canada)

In the publish/subscribe paradigm, information providers disseminate publications to all consumers who have expressed interest by registering subscriptions with the publish/subscribe system. This paradigm has found wide-spread applications, ranging from selective information dissemination to network management. In all existing publish/subscribe systems, neither subscriptions nor publications can capture uncertainty inherent to the information underlying the application domain. However, in many situations, exact knowledge of either specific subscriptions, or publications, is not available. To address this problem, we demonstrate a new publish/subscribe model based on possibility theory and fuzzy set theory to process imperfect information for, either expressing subscriptions, or publications, or both combined. Furthermore, a solid system is developed to experiment the new model.

**Efficient Constraint Processing for Highly Personalized Location Based Services**

Zhengdao Xu, Hans-Arno Jacobsen (University of Toronto, Canada)

Recently, with the advances in wireless communications and location positioning technology, the potential for tracking, correlating, and filtering information about moving entities has greatly increased. In this demonstration, we define two types of the location constraints, n-body and n-body (static) constraints, which model the correlation among a set of moving objects, and demonstrate an algorithm that efficiently evaluates those location constraints. We will show that our algorithm using Kd-tree indexing outperform the naïve approach even with little overhead for the partition update.

**LH\*RS: A Highly Available Distributed Storage System**

Witold Litwin, Rim Moussa (Universit'e Paris Dauphine, France), Thomas Schwarz (Santa Clara University, USA)

The ideal storage system is always available and incrementally expandable. Existing storage systems fall far from this ideal. Affordable computers and high-speed networks allow us to investigate storage architectures closer to the ideal. Our demo, present a prototype implementation of LH\*RS: a highly available scalable and distributed data structure.

**Demo Group 3: XML****Semantic Query Optimization in an Automata-Algebra Combined XQuery Engine over XML Streams**

Hong Su, Elke A. Rundensteiner and Murali Mani (Worcester Polytechnic Institute, USA)

Our raindrop system tackles challenges of stream processing that are particular to XML. This demo demonstrates two major aspects. First, it demonstrates the modeling issues, i.e., how to take advantage of both token-based automata paradigm and tuple-based algebraic paradigm for processing XQuery over XML streams. Second, it demonstrates the optimization issues, i.e., how to use schema knowledge to optimize an XQuery over XML streams based on our modeling.

**ShreX: Managing XML Documents in Relational Databases**

Fang Du (OGI/OHSU, USA), Sihem Amer-Yahia (AT&amp;T Labs, USA), Juliana Freire (OGI/OHSU, USA)

We describe ShreX, a freely-available system for shredding, loading and querying XML documents in relational databases. ShreX supports all mapping strategies proposed in the literature as well as strategies available in commercial RDBMSs. It provides generic (mapping-independent) functions for loading shredded documents into relations and for translating XML queries into SQL. ShreX is portable and can be used with any relational database backend.

**A Uniform System for Publishing and Maintaining XML Data**

Byron Choi, Wenfei Fan, Xibei Jia (University of Edinburgh, UK), Arek Kasprzyk (European Bioinformatics Institute)

Taking real-life data from Gene Ontology (GO), this demonstration shows how a uniform system can efficiently publish the GO data in XML with respect to predefined recursive target DTDs, and how it incrementally updates the target XML data in response to changes to the underlying GO database.

**An Injection of Tree Awareness: Adding Staircase Join to PostgreSQL**

Sabine Mayer, Torsten Grust (University of Konstanz, Germany), Maurice van Keulen (University of Twente, The Netherlands), Jens Teubner (University of Konstanz, Germany)

The syntactic well formedness constraints of XML (opening and closing tags nest properly) imply that XML processors face the challenge to efficiently handle data that takes the shape of ordered, unranked trees. Although RDBMSs have originally been designed to manage table-shaped data, we propose their use as XML and XPath processors. In our setup, the database system employs a relational XML document encoding, the XPath accelerator, which maps information about the XML node hierarchy to a table, thus making it possible to evaluate XPath expressions on SQL hosts.

Conventional RDBMSs, nevertheless, remain ignorant of many interesting properties of the encoded tree data, and were thus found to make no or poor use of these properties. This is why we devised a new join algorithm, staircase join, which incorporates the tree-specific knowledge required for an efficient SQL-based evaluation of XPath expressions. In a sense, this demonstration delivers the promise we have made at VLDB 2003: a notion of tree awareness can be injected into a conventional disk-based RDBMS kernel in terms of staircase join. The demonstration features a side-by-side comparison of both, an original and a staircase-join enhanced instance of PostgreSQL. The required changes to PostgreSQL were local, the achieved effect, however, is significant: the demonstration proves that this injection of tree awareness turns PostgreSQL into a high-performance XML processor that closely adheres to the XPath semantics.

#### **FluXQuery: An Optimizing XQuery Processor for Streaming XML Data**

C. Koch (TU Wien, Austria), S. Scherzinger (University of Passau, Germany), N. Schweikardt (Humboldt University Berlin, Germany), B. Stegmaier (University of Passau, Germany)

We present the FluXQuery engine, the first optimizing XQuery engine for XML streams. Optimization in FluXQuery is based on a new internal query language called FluX which slightly extends XQuery by a construct for event-based query processing. By allowing for the conscious use of main memory buffers, it supports reasoning over the employment of buffers during query optimization and evaluation.

#### **Demo Group 4: Web and Web-Services**

##### **COMPASS: A Concept-based Web Search Engine for HTML, XML, and Deep Web Data**

Jens Graupmann, Michael Biber, Christian Zimmer, Patrick Zimmer, Matthias Bender, Martin Theobald, Gerhard Weikum (Max-Planck Institute for Computer Science, Saarbruecken, Germany)

Today's web search engines are still following the paradigm of keyword-based search. Although this is the best choice for large scale search engines in terms of throughput and scalability, it inherently limits the ability to accomplish more meaningful query tasks. XML query engines (e.g., based on XQuery or XPath), on the other hand, have powerful query capabilities; but at the same time their dedication to XML data with a global schema is their weakness, because most web information is still stored in diverse formats and does not conform to common schemas. Typical web formats include static HTML pages or pages that are generated dynamically from underlying database systems, accessible only through portal interfaces. We have developed an expressive style of concept-based and context-aware querying with relevance ranking that encompasses different, non-schematic data formats and integrates Web Services as well as Deep Web sources. Coined COMPASS (Context-Oriented Multi-Format Portal-Aware Search System), our system features this new language that combines the simplicity of web search engines with the expressiveness of (simple forms of) XML query languages.

##### **Discovering and Ranking Semantic Associations over a Large RDF Metabase**

Chris Halaschek, Boanerges Aleman-Meza, I. Budak Arpinar, and Amit P. Sheth (University of Georgia, Athens, USA)

Information retrieval over semantic metadata has recently received a great amount of interest in both industry and academia. In particular, discovering complex and meaningful relationships among this data is becoming an active research topic. Just as ranking of documents is a critical component of today's search engines, the ranking of relationships will be essential in tomorrow's analytics engines. Building upon our recent work on specifying these semantic relationships, which we refer to as Semantic Associations, we demonstrate a system where these associations are discovered among a large semantic metabase represented in RDF. Additionally we employ ranking techniques to provide users with the most interesting and relevant results.

**An Automatic Data Grabber for Large Web Sites**

Valter Crescenzi (University of Roma, Italy), Giansalvatore Mecca (University of Potenza, Italy), Paolo Merialdo, Paolo Missier (University of Roma, Italy)

We demonstrate a system to automatically grab data from data intensive web sites. The system first infers a model that describes at the intentional level the web site as a collection of classes; each class represents a set of structurally homogeneous pages, and it is associated with a small set of representative pages. Based on the model a library of wrappers, one per class, is then inferred, with the help an external wrapper generator. The model, together with the library of wrappers, can thus be used to navigate the site and extract the data.

**WS-CatalogNet: An Infrastructure for Creating, Peering, and Querying e-Catalog Communities**

Karim Baina, Boualem Benatallah, Hye-young Paik (University of New South Wales, Australia), Farouk Toumani, Christophe Rey (ISIMA, France), Agnieszka Rutkowska, Bryan Harianto (University of New South Wales, Australia)

Given that e-catalogs are often autonomous and heterogeneous, effectively integrating and querying them is a delicate and time-consuming task. More importantly, the number of e-catalogs to be integrated and queried may be large and continuously changing. We present the WS-CatalogNet system: a Web services based information sharing middleware infrastructure whose aims is to enhance the potential of e-catalogs by focusing on scalability and flexible aspects of their sharing and access. WS-CatalogNet builds upon lessons learned in our previous work on Web information integration in order to provide a peer-to-peer and semantic-based e-catalogs sharing middleware infrastructure through which: (i) e-catalogs can be categorized and described using domain ontologies, (ii) heterogeneous service ontologies can be linked via peer relationships, (iii) e-catalogs selection caters for flexible matching between user queries and e-catalogs descriptions.

**Trust-Serv: A Lightweight Trust Negotiation Service**

Harvard Skogsrud, Boualem Benatallah (University of New South Wales, Australia), Fabio Casati (Hewlett-Packard Labs, USA), Manh Q. Dinh (University of Technology, Sydney, Australia)

Trust negotiation is an approach to access control that does not require the provider to know the requester a priori. Instead, trust is established on-the-fly in a negotiation. However, specifying and maintaining policies is hard in existing systems. The reason is usually that policy specification requires time-consuming and error-prone low-level programming. We present Trust-Serv, a model-driven trust negotiation system for Web services. Our system addresses the policy management problem by modeling trust negotiation policies as extended state machines. Trust-Serv has been implemented as a middleware that is transparent to the Web services and to the developers of Web services, simplifying Web service development and management as well as enabling scalable deployments. In our demo, we will demonstrate policy specification and management, as well as how trust negotiations are automated using our system.

**Demo Group 5: Query Optimization and Database Integrity****Green Query Optimization: Taming Query Optimization Overheads through Plan Recycling**

Parag Sarda, Jayant R. Haritsa (Indian Institute of Science, Bangalore, India)

PLASTIC is a recently-proposed tool to help query optimizers significantly amortize optimization overheads through a technique of plan recycling. The tool groups similar queries into clusters and uses the optimizer-generated plan for the cluster representative to execute all future queries assigned to the cluster. An earlier demo had presented a basic prototype implementation of PLASTIC. We have now significantly extended the scope, usability, and efficiency of PLASTIC, by incorporating a variety of new features, including an enhanced query feature vector, variable-sized clustering and a decision-tree-based query classifier. The demo of the upgraded PLASTIC tool is shown on commercial database platforms (IBM DB2 and Oracle 9i).

**Progressive Optimization in Action**

Vijayshankar Raman, Volker Markl, David Simmen, Guy Lohman, Hamid Pirahesh (IBM Almaden Research Center, USA)

Progressive Optimization (POP) is a technique to make query plans robust, and minimize need for DBA intervention, by repeatedly re-optimizing a query during runtime if the cardinalities estimated during optimization prove to be significantly incorrect. POP works by carefully calculating validity ranges for each plan operator under which the overall plan can be optimal. POP then instruments the query plan with checkpoints that validate at runtime that cardinalities do lie within validity ranges, and re-optimizes the query otherwise. In this demonstration we showcase POP implemented for a research prototype version of IBM a mix of real-world and synthetic benchmark databases and workloads. For selected queries of the workload we display the query plans with validity ranges as well as the placement of the various kinds of CHECK operators using the DB2 graphical plan explain tool. We also execute the queries, showing how and where re-optimization is triggered through the CHECK operators, the new plan generated upon re-optimization, and the extent to which previously computed intermediate results are reused.

**CORDS: Automatic Generation of Correlation Statistics in DB2**

Ihab F. Ilyas (Purdue University, USA), Volker Markl, Peter J. Haas, Paul G. Brown, Ashraf Aboulnaga (IBM Almaden Research Center, USA)

When query optimizers erroneously assume that database columns are statistically independent, they can underestimate the selectivities of conjunctive predicates by orders of magnitude. Such underestimation often leads to drastically suboptimal query execution plans. We demonstrate cords, an efficient and scalable tool for automatic discovery of correlations and soft functional dependencies between column pairs. We apply cords to real, synthetic, and TPC-H benchmark data, and show that cords discovers correlations in an efficient and scalable manner. The out-put of cords can be visualized graphically, making cords a useful mining and analysis tool for database administrators. cords ranks the discovered correlated column pairs and recommends to the optimizer a set of statistics to collect for the "most important" of the pairs. Use of these statistics speeds up processing times by orders of magnitude for a wide range of queries.

**CHICAGO: A Test and Evaluation Environment for Coarse-Grained Optimization**

Tobias Kraft, Holger Schwarz (University of Stuttgart, Germany)

Relational OLAP tools and other database applications generate sequences of SQL statements that are sent to the database server as result of a single information request issued by a user. Coarse-Grained Optimization is a practical approach for the optimization of such statement sequences based on rewrite rules. In this demonstration we present the CHICAGO test and evaluation environment that allows to assess the effectiveness of rewrite rules and control strategies. It includes a lightweight heuristic optimizer that modifies a given statement sequence using a small and variable set of rewrite rules.

**SVT: Schema Validation Tool for Microsoft SQL-Server**

Enrnest Teniente, Carles Farré, Toni Urpí, Carlos Beltrán, and David Gañán (University of Catalunya, Spain)

We present SVT, a tool for validating database schemas in SQL Server. This is done by means of testing desirable properties that a database schema should satisfy. To our knowledge, no commercial relational DBMS provides yet a tool able to perform such kind of validation.

**Demo Group 6: Streaming Data and Spatial and Multimedia Databases****CAPE: Continuous Query Engine with Heterogeneous-Grained Adaptivity**

Elke A. Rundensteiner, Luping Ding, Timothy Sutherland, Yali Zhu, Brad Pilech, Nishant Mehta

In the demonstration, we will present our continuous query engine named CAPE which is designed to function effectively in highly dynamic stream environments. The core foundation of CAPE is an optimization framework with heterogeneous-grained adaptivity. In particular, the following aspects of the CAPE system will be shown: (1) highly reactive and communicative query operators, (2)

adaptive operator scheduling, (3) runtime plan reoptimization and migration even between sub-plans containing stateful operators, and (4) adaptive query plan distribution across multiple machines. In summary, CAPE is shown to be truly reactive and adaptive in a rich diversity of stream environments.

#### **HiFi: A Unified Architecture for High Fan-in Systems**

Owen Cooper, Anil Edakkunni, Michael J. Franklin (Berkeley, USA), Wei Hong (Intel Research Berkeley), Shawn R. Jeffery, Sailesh Krishnamurthy, Fredrick Reiss, Shariq Rizvi, Eugene Wu (Berkeley, USA)

Advances in data acquisition and sensor technologies are leading towards the development of "High Fan-in" architectures: widely distributed systems whose edges consist of numerous receptors such as sensor networks and RFID readers and whose interior nodes consist of traditional host computers organized using the principle of successive aggregation. Such architectures pose significant new data management challenges. The HiFi system, under development at University of California, Berkeley, is aimed at addressing these challenges. We demonstrate an initial prototype of HiFi that uses data stream query processing to acquire, filter, and aggregate data from multiple devices including sensor motes, RFID readers, and low power single-board computer gateways organized as a High Fan-in system.

#### **An Integration Framework for Sensor Networks and Data Stream Management Systems**

Daniel J. Abadi, Wolfgang Lindner, Samuel Madden (MIT, USA), Jörg Schuler (Tufts University, USA)

This demonstration shows an integrated query processing environment where users can seamlessly query both a data stream management system and a sensor network with one query expression. By integrating the two query processing systems, the optimization goals of the sensor network (primarily power) and server network (primarily latency and quality) can be unified into one quality of service metric. The demo shows various steps of the unified optimization process for a sample query where the effects of each step that the optimizer takes can be directly viewed using a quality of service monitor. Our demo includes sensors deployed in the demo area in a tiny mockup of a factory application.

#### **QStream: Deterministic Querying of Data Streams**

Sven Schmidt, Henrike Berthold, Wolfgang Lehner (Dresden University of Technology, Germany)

Current developments in processing data streams are based on the best-effort principle and therefore not adequate for many application areas. When sensor data is gathered by interface hardware and is used for triggering data-dependent actions, the data has to be queried and processed not only in an efficient but also in a deterministic way. Our streaming system prototype embodies novel data processing techniques. It is based on an operator component model and runs on top of a real-time capable environment. This enables us to provide real Quality-of-Service for data stream queries.

#### **AIDA: an Adaptive Immersive Data Analyzer**

Mehdi Sharifzadeh, Cyrus Shahabi, Bahareh Navai, Farid Parvini, Albert A. Rizzo (University of Southern California, USA)

In this demonstration, we show various querying capabilities of an application called AIDA. AIDA is developed to help the study of attention disorder in kids. In a different study [1], we collected several immersive sensory data streams from kids monitored in an immersive application called the virtual classroom. This dataset, termed immersidata is used to analyze the behavior of kids in the virtual classroom environment. AIDA's database stores all the geometry of the objects in the virtual classroom environment and their spatio-temporal behavior. In addition, it stores all the immersidata collected from the kids experimenting with the application. AIDA's graphical user interface then supports various spatio-temporal queries on these datasets. Moreover, AIDA replays the immersidata streams as if they are collected in real-time and on which supports various continuous queries. This demonstration is a proof-of-concept prototype of a typical design and development of a domain-specific query and analysis application on the users' interaction data with immersive environments.

**BilVideo Video Database Management System**

Özgür Ulusoy, Uğur Güdükbay, Mehmet Emin Dönderler, Ediz Şaykol, Cemil Alper (Bilkent University, Turkey)

A prototype video database management system, which we call BilVideo, is presented. BilVideo provides an integrated support for queries on spatio-temporal, semantic and low-level features (color, shape, and texture) on video data. BilVideo does not target a specific application, and thus, it can be used to support any application with video data. An example application, news archives search system, is presented with some sample queries.

**PLACE: A Query Processor for Handling Real-time Spatio-temporal Data Streams**

Mohamed F. Mokbel, Xiaopeng Xiong, Walid G. Aref, Susanne E. Hambrusch, Sunil Prabhakar, Moustafa A. Hammad (Purdue University, USA)

The emergence of location-aware services calls for new real-time spatio-temporal query processing algorithms that deal with large numbers of mobile objects and queries. In this demo, we present PLACE (Pervasive Location-Aware Computing Environments); a scalable location-aware database server developed at Purdue University. The PLACE server addresses scalability by adopting an incremental evaluation mechanism for answering concurrently executing continuous spatiotemporal queries. The PLACE server supports a wide variety of stationery and moving continuous spatio-temporal queries through a set of pipelined spatio-temporal operators. The large numbers of moving objects generate real-time spatio-temporal data streams.



## CONFERENCE CO-LOCATED WORKSHOPS

### **Second International XML Database Symposium (XSym 2004)**

<http://www.xsym.org/>

**August 29th – 30th, 2004**

**Location: Manitoba**

The theme of the XML Database Symposium (XSym) is everything in the intersection of Database and XML Technologies. Today, we see growing interest in using these technologies together for many web-based and database-centric applications. XML is being used to publish data from database systems to the Web by providing input to content generators for Web pages, and database systems are increasingly used to store and query XML data, often by handling queries issued over the Internet. As database systems increasingly start talking to each other over the Web, there is a fast growing interest in using XML as the standard exchange format for distributed query processing. As a result, many relational database systems export data as XML documents and import data from XML documents and provide query and update capabilities for XML data. In addition, so-called native XML database and integration systems are appearing on the database market, whose claim is to be especially tailored to store, maintain and easily access XML-documents. The goal of this symposium is to bring together academics, practitioners, users and vendors to discuss the use and synergy between the above-mentioned technologies. Many commercial systems built today are increasingly using these technologies together and it is important to understand the various research and practical issues. The wide range of participants will help the various communities understand both specific and common problems. This symposium will provide the opportunity for all involved to debate new issues and directions for research and development work in the future.

### **Workshop on Information Integration on the Web (IIWeb)**

<http://cips.eas.asu.edu/iiweb.htm>

**August 30th, 2004**

**Location: Quebec**

The explosive growth and popularity of the world-wide web has resulted in a huge number of information sources on the Internet and the promise of unprecedented information-gathering capabilities to lay users. Unfortunately, the promise has not yet been transformed into reality. While there are sources relevant to virtually any user's query, the morass of sources presents a formidable hurdle to effectively accessing the information. One way of alleviating this problem is to develop web-based information integration systems or agents, which take a user's query or request (e.g., monitor a site), and access the relevant sources or services to efficiently support the request. The purpose of this workshop is to bring together researchers that are working in a variety of areas that are all related to the larger problem of integrating information on the Web. This includes research in the areas of machine learning, data mining, automatic planning, constraint reasoning, databases, view integration, information extraction, semantic web, web services, and other related areas.

### **Second Workshop on Spatio-Temporal Database Management (STDBM)**

<http://db.cs.ualberta.ca/stdbm04/>

**August 30th, 2004**

**Location: Library**

The one-day Second Workshop on Spatio-Temporal Database Management will bring together leading researchers and developers in the area of spatio-temporal databases in order to discuss state-of-the-art as well as novel research in spatio-temporal databases. In the spirit of a workshop, submission of novel ideas and positions that can spark discussion among the attendees are strongly encouraged. The workshop is intended to serve as a forum for disseminating research and experience in this rapidly growing area.

**Secure Data Management in a Connected World (SDM)**<http://www.extra.research.philips.com/sdm-workshop/>

August 30th, 2004

Location: Confederation 3

Concepts like ubiquitous computing and ambient intelligence that exploit increasingly interconnected networks and mobility put new requirements on data management. An important element in the connected world is that data will be accessible anytime anywhere. This also has its downside in that it becomes easier to get unauthorized data access. Furthermore, it will become easier to collect, store, and search personal information and endanger peoples' privacy. As a result security and privacy of data becomes more and more of an issue. Therefore, secure data management, which is also privacy-enhanced, turns out to be a challenging goal that will also seriously influence the acceptance of ubiquitous computing and ambient intelligence concepts by society. Aim of the workshop is to bring together people from the security research community and data management research community in order to exchange ideas on the secure management of data in the context of emerging networked services and applications. The workshop will provide forum for discussing practical experiences and theoretical research efforts that can help in solving these critical problems in secure data management. Authors from both academia and industry are invited to submit papers presenting novel research on the topics of interest.

**International Workshop on Data Management for Sensor Networks (DMSN)**<http://db.cs.pitt.edu/dmsn04/>

August 30th, 2004

Location: York

The workshop will focus on the challenges of data processing and management in networks of remote, wireless, battery-powered sensing devices (sensor networks). The power-constrained, lossy, noisy, distributed, and remote nature of such networks means that traditional data management techniques often cannot be applied without significant re-tooling. Furthermore, new challenges associated with acquisition and processing of live sensor data mean that completely new database techniques must also be developed.

**Second International Workshop On Databases, Information Systems and Peer-to-Peer Computing (DBISP2P)**<http://horizonless.ddns.comp.nus.edu.sg/dbisp2p04/>

August 29th – 30th, 2004

Location: Algonquin

The aim of this second workshop is to continue exploring the promise of P2P to offer exciting new possibilities in distributed information processing and database technologies. The realization of this promise lies fundamentally in the availability of enhanced services such as structured ways for classifying and registering shared information, verification and certification of information, content distributed schemes and quality of content, security features, information discovery and accessibility, interoperation and composition of active information services, and finally market-based mechanisms to allow cooperative and non cooperative information exchanges. The P2P paradigm lends itself to constructing large scale complex, adaptive, autonomous and heterogeneous database and information systems, endowed with clearly specified and differential capabilities to negotiate, bargain, coordinate and self-organize the information exchanges in large scale networks. This vision will have a radical impact on the structure of complex organizations (business, scientific or otherwise) and on the emergence and the formation of social communities, and on how the information is organized and processed. The P2P information paradigm naturally encompasses static and wireless connectivity, and static and mobile architectures. Wireless connectivity combined with the increasingly small and powerful mobile devices and sensors pose new challenges as well as opportunities to the database community. Information becomes ubiquitous, highly distributed and accessible anywhere and at any

time over highly dynamic, unstable networks with very severe constraints on the information management and processing capabilities. What techniques and data models may be appropriate for this environment, and yet guarantee or approach the performance, versatility and capability that users and developers come to enjoy in traditional static, centralized and distributed database environment? Is there a need to define new notions of consistency and durability, and completeness, for example? The proposed workshop will build on the success of the first one. It will concentrate on exploring the synergies between current database research and P2P computing. It is our belief that database research has much to contribute to the P2P grand challenge through its wealth of techniques for sophisticated semantics-based data models, new indexing algorithms and efficient data placement, query processing techniques and transaction processing. Database technologies in the new information age will form the crucial components of the first generation of complex adaptive P2P information systems, which will be characterized by their ability to continuously self-organize, adapt to new circumstances, promote emergence as an inherent property, optimize locally but not necessarily globally, deal with approximation and incompleteness. This workshop will also concentrate on the impact of complex adaptive information systems on current database technologies and their relation to emerging industrial technologies such as IBM's autonomic computing initiative.

## **2nd International Workshop on Semantic Web and Databases (SWBD)**

<http://swdb.semanticweb.org/>

**August 29th – 30th, 2004**

**Location: Territories**

The Semantic Web is a key initiative being promoted by the World Wide Web Consortium (W3C) as the next generation of the current web. The objective of this workshop is to discuss database and information system applications that use semantics and to gain insight into the evolution of semantic web technology. Early commercial applications that make use of machine-understandable metadata range from information retrieval to Web-enabling of old-tech IBM 3270 sessions. Current developments include metadata-based Enterprise Application Integration (EAI) systems, data modeling solutions, and wireless applications. Machine-understandable metadata is emerging as a new foundation for component-based approaches to application development. Within the context of reusable distributed components, Web services represent the latest architectural advancement. These two concepts can be synthesized providing powerful new mechanisms for quickly modeling, creating and deploying complex applications that readily adapt to real world need. We welcome submissions that describe applications which are trying to achieve the above goal through the use of specifications such as WSDL, UDDI, SOAP, ebXML, XML/RDF, possibly in conjunction with transformation languages such as XSLT, RuleML, etc. We invite submissions that discuss real applications (piloted, deployed, robust prototypes) attacking real world problems in the Semantic Web arena today instead of the future. This potentially means very pragmatic decisions (e.g., executable agents that run in the context of graph-based application models) caused by the limitations of today's technology (e.g., are there robust/scalable XML/RDF parsers, stores, and inference engines?). Submissions that also address issues related to the business aspects, e.g., is there a market?, is there a sustainable value proposition?, what is the business model?, etc., will also be considered.

**5th VLDB Workshop on Technologies for E-Services (TES)**<http://conductor.commerceone.com/VLDB-TES04.html>**August 29th – 30th, 2004****Location: British Columbia**

VLDB-TES'2004 is the fifth workshop in a successful series of annual workshops on technologies for E-Services, which have been held in conjunction with the International Conference on Very Large Data Bases. E-Services and Web Services have emerged as an effective technology for the automation of application integration across networks and organizations. They are supported by a large number of EAI and B2B automation platforms and are being heavily used in many corporations. Despite the early success, there are still many issues that need to be addressed to make integration easier, faster, cheaper, and more manageable. For examples, many standards are still immature or lack the necessary consensus to be widely adopted, development and runtime support is still far away from the level we are used to in traditional middleware, security standard is still evolving, and the service management angle is an area yet to be explored and exploited. The objective of VLDB-TES'04 is to bring together researchers, practitioners, and users to exchange new ideas, developments and experiences on issues related to E-Services. Unlike previous editions of TES, this year's edition is specifically focused on (a) innovative research work that may still be early in results, and (b) practical experiences in applying E-Service or Web Service technologies. Although we do encourage submissions that report on research results, we would also like to stimulate submissions that describe work in progress or ideas that have not yet been fully developed, but that the author feel have potential to make an impact. To this end, we would also like to solicit submissions of short papers (about 1500 words, or about 5 pages) that succinctly portray the issues, approaches, or practical experiences.

## CONFERENCE OFFICERS

### Conference Organizers

#### General Chair

John Mylopoulos, University of Toronto, Canada

#### VLDB Endowment Liaison

Kyu-Young Whang, KAIST, Korea

#### Area Coordinator, North America

Philip Bernstein, Microsoft Research, USA

#### Area Coordinator, South America

Alberto Laender, Federal University of Minas Gerais, Brazil

#### Area Coordinator, Europe, Mideast & Africa

Avigdor Gal, Technion, Israel

#### Area Coordinator, Far East & Australia

Hongjun Lu, Hong Kong University of Science and Technology, China

#### Publicity and Publications Chair

Mariano Consens, University of Toronto, Canada

#### Web Coordinator

Manuel Kolp, Catholic University of Louvain, Belgium

### Technical Program

#### Technical Program Chair

M. Tamer Özsu, University of Waterloo, Canada

#### Core Database Technology Program Chair

Donald Kossmann, University of Heidelberg, Germany

#### Infrastructure for Information Systems Program Chair

Renée J. Miller, University of Toronto, Canada

#### Industrial and Applications Program Chairs

Berni Schiefer, IBM Laboratories, Canada

José A. Blakeley, Microsoft Corporation, USA

#### Tutorial Program Co-Chairs

Raymond Ng, University of British Columbia, Canada

Matthias Jarke, RWTH Aachen, Germany

#### Panel Program Co-Chairs

Jarek Gryz, York University, Canada

Fred Lochovsky, HKUST, Hong Kong

#### Demonstrations Chairs

Bettina Kemme, McGill University, Canada

David Toman, University of Waterloo, Canada

#### Proceedings Editor

Mario A. Nascimento, University of Alberta, Canada

#### eProceedings Editor

Patricia Rodriguez-Gianolli, University of Toronto, Canada

#### Workshops Co-Chairs

Alberto Mendelzon, University of Toronto, Canada

S. Sudarshan, Indian Institute of Technology, Bombay, India

### Local Organization

#### Local Arrangement Chair

Nick Koudas, AT&T Labs Research, USA

#### Exhibits Chair

H.-Arno Jacobsen, University of Toronto, Canada

#### Treasurer

Kenneth Barker, University of Calgary, Canada

#### Fund Raising Chair

Victor DiCiccio, University of Waterloo, Canada

#### Registration Chair

Grant Weddell, University of Waterloo, Canada

## Program Committee Members

### Core Database Technology

Karl Aberer, EPFL, Switzerland  
 Anastassia Ailamaki, CMU, USA  
 Demet Aksoy, UC Davis, USA  
 Paul Aoki, PARC, USA  
 Walid Aref, Purdue University, USA  
 Remzi Arpacı-Dusseau, Univ. of Wisconsin-Madison, USA  
 Philip Bohannon, Bell Labs, USA  
 Christian Boehm, University of Munich, Germany  
 Philippe Bonnet, University of Copenhagen, Denmark  
 Stephane Bressan, National U. of Singapore, Singapore  
 Mike Carey, BEA Systems, USA  
 Stefano Ceri, Politecnico Milan, Italy  
 Surajit Chaudhuri, Microsoft Research, USA  
 Arbee Chen, National Dong Hwa University, Taiwan  
 Bobby Cochrane, IBM Almaden, USA  
 Laurent Daynes, Sun Microsystems, France  
 Alin Deutsch, UC San Diego, USA  
 Klaus Dittrich, University of Zurich, USA  
 Wenfei Fan, Bell Labs, USA  
 Mary Fernandez, ATT Research, USA  
 Elena Ferrari, University of Milano, Italy  
 Daniela Florescu, BEA Systems, USA  
 Mike Franklin, UC Berkeley, USA  
 Johann-Christoph Freytag, Humboldt-Universität, Germany  
 Johannes Gehrke, Cornell University, USA  
 Goetz Graefe, Microsoft, USA  
 Ralf Güting, Fernuni Hagen, Germany  
 Peter Haas, IBM Almaden, USA  
 Alon Halevy, University of Washington, USA  
 Sven Helmer, University of Mannheim, Germany  
 Alexander Hinneburg, University of Halle, Germany  
 Wie Hong, Intel Research Lab, USA  
 H.V. Jagadish, University of Michigan, USA  
 Christian S. Jensen, Aalborg University, Denmark  
 Björn Þór Jónsson, Reykjavik University, Iceland  
 Alfons Kemper, University of Passau, Germany  
 Martin Kersten, CWI, Netherlands  
 Masaru Kitsuregawa, KITSuregawa Lab, Japan  
 George Kollios, Boston University, USA  
 Hank Korth, Lehigh University, USA  
 Alexandros Labrinidis, University of Pittsburgh, USA  
 Paul Larson, Microsoft Research, USA  
 David Lomet, Microsoft Research, USA  
 Hongjun Lu, HKUST, China  
 Ioana Manolescu, INRIA Futurs, France  
 Volker Markl, IBM Almaden, USA  
 Guido Moerkotte, University of Mannheim, Germany  
 Mario A. Nascimento, University of Alberta, Canada  
 Chris Olston, Carnegie Mellon University, USA  
 Fatma Özcan, IBM Almaden, USA  
 Dimitris Papadias, HKUST, China  
 Yannis Papakonstantinou, UC San Diego, USA  
 Krithi Ramamritham, IIT Bombay, India  
 Tore Risch, Uppsala University, Sweden  
 Michael Rys, Microsoft, USA  
 Ken Salem, University of Waterloo, Canada  
 Sunita Sarawagi, IIT Bombay, India  
 Oded Shmueli, Technion, Israel  
 Ioana Stanoi, IBM Almaden, USA  
 S. Sudarshan, IIT Bombay, India  
 Kian-Lee Tan, National University of Singapore, Singapore  
 Dirk van Gucht, University of Indiana, USA  
 Jun Yang, Duke University, USA  
 Clement Yu, University of Illinois, USA  
 Masatoshi Yoshikawa, Nagoya University, Japan  
 Carlo Zaniolo, UCLA, USA  
 Justin Zobel, RMIT, Australia

### Infrastructure for Information Systems

Sihem Amer-Yahia, AT&T Labs Research, USA  
 Paolo Atzeni, Università di Roma Tre, Italy  
 Phil Bernstein, Microsoft Research, USA  
 Leo Bertossi, Carleton University, Canada  
 Peter Buneman, University of Edinburgh, UK  
 Tiziana Catarci, Università di Roma La Sapienza, Italy  
 Soumen Chakrabarti, Indian Inst. of Tech., Bombay, India  
 Mitch Cherniack, Brandeis University, USA  
 Panos Chrysanthis, University of Pittsburgh, USA  
 Mariano Consens, University of Waterloo, Canada  
 Isabel Cruz, University of Illinois at Chicago, USA  
 Tamraparni Dasu, AT&T Labs-Research, USA  
 Asuman Dogac, Middle East Tech Univ., Turkey  
 Lise Getoor, University of Maryland at College Park, USA  
 Dimitrios Gunopoulos, Univ. of California, Riverside, USA  
 Laura Haas, IBM Silicon Valley Lab, USA  
 Jayant Haritsa, Indian Inst. of Science, Bangalore, India  
 Wynne Hsu, National University of Singapore, Singapore  
 Carlos Hurtado, Universidad de Chile, Santiago, Chile  
 Chris Jermaine, University of Florida, USA  
 Christoph Koch, University of Edinburgh, UK  
 Nick Koudas, AT&T Labs Research, USA  
 Hans-Peter Kriegel, University of Munich, Germany  
 Mong Li Lee, National University of Singapore, Singapore  
 Maurizio Lenzerini, Università di Roma La Sapienza, Italy  
 Qiong Luo, Hong Kong Univ. Science & Tech., China  
 Pat Martin, Queens University, Canada  
 David Maier, Oregon Health & Science University, USA  
 Alberto Mendelzon, University of Toronto, Canada  
 Tova Milo, Tel Aviv University, Israel  
 Jeff Naughton, University of Wisconsin-Madison, USA  
 Raymond Ng, UBC, Canada  
 Joann Ordille, Avaya, USA  
 Themis Palpanas, University of California, Riverside, USA  
 Jignesh Patel, University of Michigan, USA  
 Lucian Popa, IBM Almaden Research Center, USA  
 Alexandra Poulouvasilis, Univ. of London, UK  
 Davood Rafiei, University of Alberta, Canada  
 Raghu Ramakrishnan, University of Wisconsin-Madison, USA  
 Rajeev Rastogi, Bell Labs, USA  
 Prasan Roy, Indian Institute of Technology, Bombay, India  
 Elke Rundensteiner, Worcester Polytechnic Institute, USA  
 Timos Sellis, National Technical Univ. of Athens, Greece  
 Dennis Shasha, New York University, USA  
 Kyuseok Shim, Seoul National Univ., South Korea  
 Eric Simon, Inria Rocquencourt, France  
 Avi Silberschatz, Yale University, USA  
 Divesh Srivastava, AT&T Labs Research, USA  
 Wang-Chiew Tan, University of California, Santa Cruz, USA  
 David Toman, University of Waterloo, Canada  
 Alejandro Vaisman, Universidad de Buenos Aires, Argentina  
 Victor Vianu, Univ. California, San Diego, USA  
 Limsoon Wong, Institute for Infocomm Research, Singapore

**INDUSTRIAL AND APPLICATION PROGRAM**

Achim Kraiss, SAP, Germany  
 Alejandro P. Buchmann, Darmstadt University of  
 Technology, Germany  
 Anand Deshpande, Persistent Systems, India  
 Anil K. Nori, Microsoft Corporation, USA  
 Erhard Rahm, Universität Leipzig, Germany  
 Francois J. Raab, TPC Certified Auditor, USA  
 Glenn Paulley, iAnywhere Solutions, Canada  
 Guy M. Lohman, IBM Almaden Research Center, USA  
 Jacob Slonim, Dalhousie University, Canada  
 Joseph M. Hellerstein, UC Berkeley, USA  
 Josep-L. Larriba-Pey, Univ. Politècnica de Catalunya, Spain  
 Michael L. Brodie, Verizon, USA  
 Srinivasan Seshadri, Strand Genomics, India  
 Seckin Unlu, Intel Corporation, USA  
 Sophie Cluet, INRIA, France  
 Umeshwar Dayal, Hewlett-Packard Laboratories, USA  
 Vishu Krishnamurthy, Oracle Corporation, USA

**WORKSHOP SELECTION**

USA Nicolas Bruno, Microsoft, USA  
 Gianni Mecca, Università della Basilicata, Italy  
 Alberto Mendelzon, University of Toronto, Canada (co-chair)  
 Beng Chin Ooi, National University of Singapore  
 S. Sudarshan, IIT Bombay, India (co-chair)  
 Janet Wiener, Hewlett-Packard, USA  
 Peter Wood, Birkbeck College, UK  
 Calisto Zuzarte, IBM Toronto, Canada

# University of Waterloo



## Database Group

<http://db.waterloo.ca>

---

*The University of Waterloo Database Group is one of the largest database groups with eight core faculty (Ashraf Aboulnaga, Edward Chen, Ihab Ilyas, M. Tamer Özsu, Ken Salem, David Toman, Frank Tompa, Grant Weddell), four associated faculty, and about forty graduate students. The research is broad-based, and addresses data management issues as they apply to complex objects, geometric and spatial data, video and audio data, historical data, structured documents, semistructured data and text, real-time data, streaming data, and Web data as well as to traditional relational database systems. This wide coverage mirrors the expanding role of database systems in managing very large amounts of diverse data. The research also has wide coverage of many aspects of database system functionality. Consequently, there are ongoing projects that explore storage structures, indexing methods, query languages and optimizers, and transaction management. This broad coverage in terms of application domains and database functionality is a unique and identifying characteristic of the group. More information about UW Database Group and its activities can be obtained from its web site.*

---



**NOTES**

A large, empty rectangular box with a thin black border, intended for the user to write their notes. It occupies the majority of the page's vertical space.

**NOTES**

A large, empty rectangular box with a thin black border, occupying most of the page below the 'NOTES' header. It is intended for the user to write their notes.

**NOTES**

A large, empty rectangular box with a thin black border, intended for taking notes. It occupies the majority of the page's vertical space below the header.

**NOTES**

A large, empty rectangular box with a thin black border, intended for the user to write their notes. The box occupies most of the page's vertical space below the header.

**NOTES**

A large, empty rectangular box with a thin black border, occupying most of the page below the 'NOTES' header. It is intended for the user to write their notes.

**NOTES**

A large, empty rectangular box with a thin black border, occupying most of the page below the 'NOTES' header. It is intended for the user to write their notes.

**NOTES**

A large, empty rectangular box with a thin black border, occupying most of the page below the 'NOTES' header. It is intended for the user to write their notes.

**NOTES**

A large, empty rectangular box with a thin black border, intended for the user to write their notes. It occupies the central portion of the page, below the 'NOTES' header and above the footer.



**NOTES**

A large, empty rectangular box with a thin black border, occupying most of the page below the 'NOTES' header. It is intended for the user to write their notes.

**NOTES**

A large, empty rectangular box with a thin black border, occupying most of the page below the 'NOTES' header. It is intended for the user to write their notes.





# TORONTO 2004

30<sup>th</sup> INTERNATIONAL CONFERENCE

August 29<sup>th</sup> - September 3<sup>rd</sup>, 2004

Royal York Hotel, Toronto, Canada

[www.vldb04.org](http://www.vldb04.org)

## PREMIER SPONSOR



**DB2.**

Information Management Software

## PLATINUM SPONSOR

**ORACLE®**

## GOLD SPONSORS

Microsoft  
**SQL Server™ 2005**

**sgi**

**SYBASE®**  
The Enterprise. Unwired.

## SILVER SPONSORS



**intel®**

InterSystems  
**CACHE**  
Make Applications Faster

## ACADEMIC SPONSORS



UNIVERSITY OF  
**TORONTO**

